NOKIA

7210 SERVICE ACCESS SWITCH

7210 SAS OS Services Guide
7210 SAS-D
7210 SAS-E
7210 SAS-K2F2T1C
7210 SAS-K2F4T6C
Release 9.0.R8

3HE11478AAAJTQZZA

Issue: 01

November 2017

Nokia — Proprietary and confidential.
Use pursuant to applicable agreements.

Nokia is a registered trademark of Nokia Corporation. Other products and company names mentioned herein may be trademarks or trade-names of their respective owners.

The information presented is subject to change without notice. No responsibility is assumed for inaccuracies contained herein.

© 2016, 2017 Nokia.

Contains proprietary/trade secret information which is the property of Nokia and must not be made available to, or copied or used by anyone outside Nokia without its written authorization. Not to be used or disclosed except in accordance with applicable agreements.

TABLE OF CONTENTS

Preface 17
About This Guide 17
Audience 17
List of Technical Publications 18

Getting Started
In This Chapter 21
Alcatel-Lucent 7210 SAS Services Configuration Process 21

Services Overview
In This Section 23
Introduction 25
Service Types on 7210 SAS-D and 7210 SAS-E 26
Service Types on 7210 SAS-K 2F2T1C 26
Service Types on 7210 SAS-K 2F4T6C 26
Service Policies on 7210 SAS-D and 7210 SAS-E 27
Service Policies on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C 29
Nokia Service Model on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C 30
Alcatel-Lucent Service Model on 7210 SAS-K 2F4T6C 31
Service Entities on 7210 SAS-D, 7210 SAS-E, and 7210 SAS- K 2F2T1C 32
Service Entities on 7210 SAS-K 2F4T6C 33
Customers 33
Service Access Points (SAPs) on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C 34
Service Access Points (SAPs) on 7210 SAS-K 2F4T6C 35
SAP Encapsulation Types and Identifiers 37
Ethernet Encapsulations supported on 7210 SAS-E 37
Ethernet Encapsulations supported on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 38
Services and SAP Encapsulations on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 39
Default SAP on a Dot1q Port on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 39
Default SAPs on a QinQ Port on a 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 40
Configuration Notes for use of Default QinQ SAPs for transit service in a ring deployment on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 43
SAP Configuration Considerations applicable for SAPs configured on access-uplink ports on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 43
Configuration Guidelines for 7210 SAS-E 45
Configuration Guidelines for 7210 SAS-D 46
Configuration Guidelines for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C 48
Service Distribution Points (SDPs) on 7210 SAS-K 2F4T6C 50
SDP Binding 50
Spoke and Mesh SDPs 51
SDP Using BGP Route Tunnel 51

SDP Keepalives 52
SDP Administrative Groups 53
Mixed-LSP Mode of Operation 54
G.8032/Ethernet Ring Protection Switching 56
Overview of G.8032 Operation 56
G8302- Ethernet Ring Sub-Rings 62
G8302- Virtual and Non-Virtual Channel 64
G8302- Ethernet Ring Sub Ring using non-virtual-link 66
G8302- OAM Considerations 69
G8302- QoS Considerations 69
G8302- Service and Solution Combinations 69
G8302- Configuration guidelines 70
Layer 2 Control Processing (L2CP) 71
Service Creation Process Overview on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 73
Deploying and Provisioning Services 74
Phase 1: Core Network Construction 74
Phase 2: Service Administration 74
Phase 3: Service Provisioning 74
Configuration Notes 75
General 75
Configuring Global Service Entities with CLI 77
Service Model Entities 77
Basic Configuration 78
Common Configuration Tasks 79
Configuring Customers 79
Customer Information 79
Service Creation Process Overview on 7210 SAS-K 2F4T6C 81
Deploying and Provisioning Services - 7210 SAS-K 2F4T6C 82
Phase 1: Core Network Construction 82
Phase 2: Service Administration 82
Phase 3: Service Provisioning 82
Configuration Notes 83
General 83
Configuring Global Service Entities with CLI 83
Service Model Entities 83
Basic Configuration 84
Common Configuration Tasks 85
Configuring Customers 85
Customer Information 85
Configuring an SDP 86
SDP Configuration Tasks 86
Configuring an SDP 87
Configuring a Mixed-LSP SDP 88
Ethernet Connectivity Fault Management (ETH-CFM) 89
Common Actionable Failures 93
MEP and MIP Support 94
ETH-CFM Support for 7210 SAS-D 94
ETH-CFM Support for 7210 SAS-E 95

ETH-CFM Support for 7210 SAS-K 2F2T1C 95
ETH-CFM Support for 7210 SAS-K 2F4T6C 96
Configuring ETH-CFM Parameters 98
Service Management Tasks 100
Modifying Customer Accounts 100
Deleting Customers 101
Global Services Command Reference 103
Command Hierarchies 103
Show Commands 109
Tools Perform Commands for 7210 SAS-K 2F4T6C 109

VLL Services

In This Chapter 145
Ethernet Pipe (Epipe) Services 146
Epipe Service Overview 147
Epipe Oper State decoupling on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 149
Dynamic Multi-Segment Pseudowire Routing 151
Dynamic Multi-Segment Pseudowire Routing - Overview 151
Multi-Segment Pseudowire Routing - Pseudowire Routing 156
Multi-Segment Pseudowire Routing - Configuring VLLs using Dynamic MS-PWs 158
Multi-Segment Pseudowire Routing - Pseudowire Redundancy 160
Multi-Segment Pseudowire Routing - VCCV OAM for Dynamic MS-PWs 163
Multi-Segment Pseudowire Routing - VCCV-Ping on Dynamic MS-PWs 163
Multi-Segment Pseudowire Routing - VCCV-Trace on Dynamic MS-PWs 164
Multi-Segment Pseudowire Routing - Example Dynamic MS-PW Configuration 165
PW Redundancy - Master-Slave Operation 167
VLL Resilience for a Switched Pseudowire Path 172
Access Node Resilience Using MC-LAG and Pseudowire Redundancy 176
VLAN Range for SAPs (or Dot1q range SAPs) in an Epipe Service 178
Processing behavior for Dot1q range SAPs in access-uplink mode on 7210 SAS-D and 7210 SAS-K 2F2T1C 178
VLAN Range SAPs feature Support and Restrictions 179
Processing behavior for Dot1q range SAPs on SAS-K 2F4T6C 180
Epipe Pseudowire switching in 7210 SAS-K 2F4T6C 182
Pseudo-wire Switching with Protection 183
Pseudo-wire Switching Behavior 185
Pseudo-wire Redundancy 187
VLL service resiliency with Pseudowire on 7210 SAS-K 2F4T6C 188
VLL Resilience with Two Destination PE Nodes 188
VLL Resilience - Operation of Master-Slave Pseudo-wire Redundancy with Existing Scenarios 191
VLL Resilience for a Switched Pseudo-wire Path 192
VLL Resilience - Pseudo-wire Redundancy Service Models 193
Processing Endpoint SAP Active/Standby Status Bits 195
Processing and Merging 196
MPLS Entropy Label and Hash Label 197
VLL Service Considerations 198
SDPs 198
SAP Encapsulations 198

QoS Policies 199
Filter Policies 199
MAC Resources 199
Configuring a VLL Service with CLI 201
Basic Configurations 202
Common Configuration Tasks 202
Creating an Epipe Service for 7210 SAS-E 203
Creating an Epipe Service for 7210 SAS-D 203
Creating an Epipe Service with range SAPs 204
Creating an Epipe Service for 7210 SAS-K 2F2T1C 204
Creating an Epipe Service for 7210 SAS-K 2F2T1C with range SAPs 205
Configuring Default QinQ SAPs for Epipe Transit Traffic in a Ring Scenario in access-uplink mode 217
Service Management Tasks 218
Modifying Epipe Service Parameters 219
Disabling an Epipe Service 219
Re-Enabling an Epipe Service 220
Deleting an Epipe Service 220
VLL Services Command Reference 221

Virtual Private LAN Service

In This Chapter 277
VPLS Service Overview 279
VPLS Packet Walk-through for a Local service 280
VPLS Packet Walk-through for a Distributed service 283
VPLS Features 286
VPLS Enhancements on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C 286
VPLS Enhancements on 7210 SAS-K 2F4T6C 286
VPLS over MPLS - supported on 7210 SAS-K 2F4T6C 287
VPLS over SAPs on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F4T1C, and 7210 SAS-K 2F4T6C . . 288
VPLS MAC Learning and Packet Forwarding 288
IGMP Snooping on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C 289
IGMP Snooping on 7210 SAS-K 2F4T6C 290
Multicast VLAN Registration (MVR) support 292
L2 Forwarding Table Management 294
FIB Size on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C 294
FIB Size on 7210 SAS-K 2F4T6C 294
FIB Size Alarms 295
Local Aging Timers on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C 296
Local and Remote Aging Timers on 7210 SAS-K 2F4T6C 296
Disable MAC Aging 297
Disable MAC Learning 297
Unknown MAC Discard 297
VPLS and Rate Limiting 297
MAC Move 298
Split Horizon SAP Groups on 7210 SAS-K 298
VPLS and Spanning Tree Protocol 299
Spanning Tree Operating Modes 299
Multiple Spanning Tree 300

MSTP for QinQ SAPs 302
Enhancements to the Spanning Tree Protocol in MPLS-based VPLS service (7210 SAS-K 2F4T6C only) 304
VPLS Access Redundancy on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 307
STP-Based Redundant Access to VPLS 307
VPLS Redundancy on 7210 SAS-K 2F4T6C using MPLS uplinks 308
Spoke SDP Redundancy (on 7210 SAS-K 2F4T6C) 308
Spoke SDP Based Redundant Access (on 7210 SAS-K 2F4T6C) 309
Inter-Domain VPLS Resiliency Using Multi-Chassis Endpoints (on 7210 SAS-K 2F4T6C) 310
VPLS Access Redundancy when using MPLS uplinks on 7210 SAS-K 2F4T6C 311
STP-Based Redundant Access to VPLS using PWs on 7210 SAS-K 2F4T6C 311
Redundant Access to VPLS Without STP and using MPLS uplinks (on 7210 SAS-K 2F4T6C) . . 312
MAC Flush Message Processing in VPLS services with MPLS uplinks (on 7210 SAS-K 2F4T6C) 313
Dual Homing to a VPLS Service (on 7210 SAS-K 2F4T6C) 314
VPLS Service Considerations 316
SAP Encapsulations 316
VLAN Processing 316
IP Interface in a VPLS Service on 7210 SAS-E 317
BGP-AD VPLS 318
BGP-AD and Target LDP (T-LDP) Interaction on 7210 SAS-K 2F4T6C 318
BGP-AD VPLS - SDP Usage 319
BGP-AD VPLS - Automatic Creation of SDPs 319
BGP-AD VPLS - Manually Provisioned SDP 321
BGP-AD VPLS - Automatic Instantiation of Pseudowires (SDP Bindings) 321
BGP-AD VPLS - Mixing Statically Configured and Auto-Discovered Pseudowires in a VPLS service 322
BGP-AD VPLS - Resiliency Schemes 322
Routed VPLS on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 323
Routed VPLS - IES IP Interface Binding- 7210 SAS-D and 7210 SAS-K 2F2T1C 323
Routed VPLS - IES or VPRN IP Interface Binding 324
Routed VPLS - Assigning a Service Name to a VPLS Service 324
Routed VPLS - Service Binding Requirements 324
Routed VPLS - Bound Service Name Assignment 325
Routed VPLS - Binding a Service Name to an IP Interface 325
Routed VPLS - IP Interface Attached VPLS Service Constraints 325
Routed VPLS - IP Interface and VPLS Operational State Coordination 326
Routed VPLS - IP Interface MTU and Fragmentation on 7210 SAS-D 326
Routed VPLS - IP Interface MTU and Fragmentation on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C 326
Routed VPLS - ARP and VPLS FIB Interactions 327
Routed VPLS - Specific ARP Cache Behavior 328
Routed VPLS - The allow-ip-int-binding VPLS Flag 329
Routed VPLS - SAPs Only Supported on Standard Ethernet Ports 330
Routed VPLS - LAG Port Membership Constraints 330
Routed VPLS - VPLS Feature Support and Restrictions 330
Routed VPLS - VPLS SAP Ingress IP Filter Override on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 331

Routed VPLS - QoS Support for VPLS SAPs and IP interface 335
Routed VPLS - Routing Related Protocols on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 335
Routed VPLS - Spanning Tree and Split Horizon 336
Routed VPLS support available and Caveats 337
Epipe Emulation using Dot1q VLAN range SAP in VPLS with G.8032 338
Epipe Emulation using Dot1q VLAN range SAP in VPLS with G.8032 -Configuration guidelines and restrictions 340
Configuring a VPLS Service with CLI 343
Basic Configuration 344
Common Configuration Tasks 346
Configuring VPLS Components 346
Creating a VPLS Service 347
Configuring a VPLS SAP 356
Configuring VPLS Redundancy 369
Creating a Management VPLS for SAP Protection 369
Configuring Load Balancing with Management VPLS 371
Configuring a BGP-Auto-Discovery on 7210 SAS-K 2F4T6C 373
Configuration Steps 374
BGP- AD- LDP Signaling 377
BGP-AD- Pseudowire Template 378
Configuring a VPLS Management Interface 380
Creating a Management VPLS for SAP Protection, using MPLS uplinks (on 7210 SAS-K 2F2T1C) 381
Creating a Management VPLS for Spoke SDP Protection, using MPLS uplinks (7210 SAS-K 2F4T6C) 382
Configuring Load Balancing with Management VPLS using MPLS uplinks (on 7210 SAS-K 2F4T6C) 383
Configuring Selective MAC Flush, when using MPLS uplinks (on 7210 SAS-K 2F4T6C) 388
Configuring AS Pseudo-wire in VPLS, when using MPLS uplinks (on 7210 SAS-K 2F4T6C) 388
Service Management Tasks 390
Modifying VPLS Service Parameters 390
Modifying Management VPLS Parameters 391
Deleting a Management VPLS 391
Disabling a Management VPLS 392
Deleting a VPLS Service 393
Disabling a VPLS Service 393
Re-Enabling a VPLS Service 393
VPLS Services Command Reference 395
Command Hierarchies 395
DHCP Commands 420
VPLS Commands 463

Internet Enhanced Service on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C
In This Chapter 493
IES Service Overview 494
IES Features 496
IP Interfaces 496
IPv6 support for IES IP interfaces associated with Access-Uplink SAPs (on 7210 SAS-D) 497

SAPs 498
Encapsulations 498
CPE Connectivity Check 498
QoS Policies 499
CPU QoS for IES interfaces in access-uplink mode 499
Filter Policies 500
Configuring an IES Service with CLI 501
Basic Configuration 502
Common Configuration Tasks 503
Configuring IES Components 504
Configuring an IES Service 504
Configuring IES Interface Parameters 504
Configuring IES SAP Parameters 505
Configuring SAP Parameters 506
Service Management Tasks 507
Modifying IES Service Parameters 507
Deleting an IES Service 508
Disabling an IES Service 508
Re-Enabling an IES Service 508
IES Services Command Reference 509
Command Hierarchies 509

Internet Enhanced Service on 7210 SAS-K 2F4T6C

In This Chapter 539
IES Service Overview 540
IES Features 541
IP Interfaces 541
SAPs 542
Encapsulations 542
CPE Connectivity Check 542
QoS Policies 543
CPU QoS for IES interfaces 543
Filter Policies 543
Configuring an IES Service with CLI 545
Basic Configuration 545
Common Configuration Tasks 545
Configuring IES Components 546
Service Management Tasks 548
Modifying IES Service Parameters 549
Deleting an IES Service 549
Disabling an IES Service 550
Re-Enabling an IES Service 550
IES Services Command Reference 551
Command Hierarchies 551

Virtual Private Network Service on 7210 SAS-K 2F4T6C

In This Chapter 581
VPRN Service Overview 583
Routing Prerequisites 584

BGP Support 584
Route Distinguisher 585
Route Reflector 586
CE to PE Route Exchange 586
Route Redistribution 586
CPE Connectivity Check 587
VPRN Features 589
IP Interfaces 589
SAPs 589
Encapsulations 589
QoS Policies for 7210 SAS-K 2F4T6C 590
Filter Policies 590
CPU QoS and DSCP Marking for Self-Generated Traffic 590
CE to PE Routing Protocols 590
PE to PE Tunneling Mechanisms 591
Per VRF Route Limiting 591
Spoke SDP termination 591
Using OSPF in IP-VPNs 591
Service Label Mode of a VPRN 592
Inter-AS VPRNs 592
Configuring a VPRN Service with CLI 595
Basic Configuration 595
Common Configuration Tasks 596
Configuring VPRN Components 597
Configuring VPRN Protocols - OSPF 603
VPRN OSPF CLI Syntax 603
Service Management Tasks 604
Modifying VPRN Service Parameters 604
Deleting a VPRN Service 605
Disabling a VPRN Service 606
Re-enabling a VPRN Service 607
VPRN Services Command Reference for 7210 SAS-K 2F4T6C 609
Command Hierarchies 609

Service Global Commands
In This Chapter 825
Show, Clear, Debug, Commands 826

Appendix: Port-Based Split Horizon
In This Chapter 931
Overview 932
Topology 932
Configuration Guidelines 934
Verification 936

Appendix: DHCP Management
In This Chapter 937
DHCP Principles 938
DHCP Features 940

Using Option 82 Field 940
Trusted and Untrusted 941
DHCP Snooping 941
Common Configuration Guidelines 943
Configuration Guidelines for DHCP relay and snooping 943
Configuring Option 82 Handling 943

Common CLI Command Descriptions
In This Chapter 945
Common Service Commands 946

Table of Contents

List of Tables

Getting Started

Table 1: Configuration Process 21

Services Overview

Table 2: Service and Encapsulation for access port 39
Table 3: Port Type and Encapsulation for access uplink ports 39
Table 4: SAP and Service Combinations for 7210 SAS-E 46
Table 5: SAP and Service Combinations for 7210 SAS-D 47
Table 6: SAP and Service Combinations for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C 49
Table 7: L2CP for 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K5 and 7210 SAS-12 Platforms 72
Table 8: ETH-CC defect condition groups 93
Table 9: ETH-CFM Support Matrix for 7210 SAS-D 94
Table 10: ETH-CFM Support Matrix for 7210 SAS-E 95
Table 11: ETH-CFM Support Matrix for 7210 SAS-K 2F2T1C 95
Table 12: ETH-CFM Support Matrix for 7210 SAS-K 2F4T6C 96

VLL Services

Table 13: Final Disposition of the packet based on per FC and per SAP policer or meter. 252
Table 14: Table for Dot1p values extracted from the packet on SAP (ingress) when Dot1p-inner and Dot1p-outer CLI command are configured:. 254
Table 15: Table for Dot1p values marked in the packet on SAP egress when Dot1p inner and Dot1p-outer CLI command are configured:. 255

Virtual Private LAN Service

Table 16: Routing behavior in RVPLS and interaction ARP Cache and MAC FIB 328
Table 17: ACL Lookup behavior with Ingress Override filter attached to an IES interface in a R-VPLS service. 332
Table 18: ACL Lookup behavior without Ingress Override filter attached to an IES interface in a R-VPLS service. 333
Table 19: Routing Protocols on IP interfaces bound to a VPLS service. 335
Table 20: SAP BPDU Encapsulation States 365
Table 21: Final Disposition of the packet based on per FC and per SAP policer or meter. 472

List of Figures

Services Overview

Figure 1: Multiple SAPs in a service using QinQ uplinks in 7210 SAS configured in access-uplink mode 34
Figure 2: Service Access Point (SAP) for 7210 SAS configured with MPLS uplinks 35
Figure 3: Multiple SAPs in a service using QinQ uplinks in 7210 SAS configured in access-uplink mode 36
Figure 4: Multiple SAPs on a Single Port 37
Figure 5: Multiple SAPs on a Single Port 38
Figure 6: MPLS Service Distribution Point (SDP) pointing from ALA-A to ALA-B 51
Figure 7: G.8032 Ring in the Initial State 57
Figure 8: G.8032 Ring in the Protecting State 58
Figure 9: 0-3 Ring Example 60
Figure 10: G.8032 Sub-Ring 63
Figure 11: 0-6 Sub-Ring Homed to VPLS 66
Figure 12: Service Creation and Implementation Flow 73
Figure 13: Service Creation and Implementation flow 81
Figure 14: Ethernet OAM Model for Ethernet Access - Business 91
Figure 15: Ethernet OAM Model for Ethernet Access - Wholesale 91

VLL Services

Figure 16: Local Epipe/VLL Service using Access or Access-uplink ports on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C 147
Figure 17: Distributed Epipe/ VLL service using Network ports on 7210 SAS-K 2F4T6C 148
Figure 18: Dynamic MS-PW Overview 151
Figure 19: Figure 2 MS-PW Addressing using FEC129 AII Type 2 152
Figure 20: Advertisement of PE Addresses by PW Routing 153
Figure 21: Signaling of Dynamic MS-PWs using T-LDP 154
Figure 22: Mapping of AII to SAP 154
Figure 23: VLL Using Dynamic MS-PWs, Inter-AS Scenario 155
Figure 24: Pseudowire Redundancy 161
Figure 25: Dynamic MS-PW Example 165
Figure 26: Master-Slave Pseudowire Redundancy 168
Figure 27: VLL Resilience 171
Figure 28: VLL Resilience with Pseudowire Redundancy and Switching 173
Figure 29: VLL Resilience with Pseudowire Switching 175
Figure 30: Access Node Resilience 177
Figure 31: VLL Resilience with Pseudowire Redundancy and Switching 184
Figure 32: Pseudo-wire Switching TLV 186
Figure 33: VLL Resilience 188
Figure 34: Master-Slave Pseudo-wire Redundancy 190
Figure 35: VLL Resilience 191
Figure 36: VLL Resilience with Pseudo-wire Switching 192
Figure 37: Redundant VLL Endpoint Objects 194
Figure 38: SDPs — Unidirectional Tunnels 210
Figure 39: VLL Resilience 214

Figure 40: VLL Resilience with Pseudo-wire Switching 215
Figure 41: Default QinQ SAP for Transit Traffic in a Ring Scenario 217

Virtual Private LAN Service

Figure 42: VPLS Service Architecture 280
Figure 43: Access Port Ingress Packet Format and Lookup 281
Figure 44: Network Port Egress Packet Format and Flooding 281
Figure 45: VPLS Service Architecture 283
Figure 46: Access Port Ingress Packet Format and Lookup 284
Figure 47: Network Port Egress Packet Format and Flooding 285
Figure 48: MVR and MVR by Proxy 293
Figure 49: Access Resiliency 302
Figure 50: Dual Homed 7210 SAS E Acting as MTU-s in Two-Tier Hierarchy H-VPLS 307
Figure 51: HVPLS with spoke redundancy 309
Figure 52: HVPLS Resiliency based on AS Pseudowires 310
Figure 53: Dual Homed MTUs in Two-Tier Hierarchy H-VPLS 311
Figure 54: Dual Homed CE Connection to VPLS 314
Figure 55: BGP-AD and T-LDP Interaction 318
Figure 56: Epipe Emulation in a ring using VPLS with G.8032 338
Figure 57: Example Configuration for Protected VPLS SAP 370
Figure 58: Example Configuration for Load Balancing Across with Management VPLS 371
Figure 59: BGP AD Configuration Example 375
Figure 60: BGP-AD CLI Command Tree 375
Figure 61: BGP AD Triggering LDP Functions 377
Figure 62: Show Router LDP Session Output 378
Figure 63: Show Router LDP Bindings FEC-Type Services 378
Figure 64: PW-Template-Binding CLI Syntax 380
Figure 65: Sample Topology-AS Pseudo-wire in VPLS 388

Internet Enhanced Service on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

Figure 66: Internet Enhanced Service 494

Internet Enhanced Service on 7210 SAS-K 2F4T6C

Figure 67: Internet Enhanced Service 540

Virtual Private Network Service on 7210 SAS-K 2F4T6C

Figure 68: Virtual Private Routed Network 584
Figure 69: Route Distinguisher 585
Figure 70: Directly connected to IP target 587
Figure 71: Multiple Hops to IP target 587
Figure 72: Inter-AS Option-A: VRF-to-VRF model 593
Figure 73: Inter-AS Option-B 593
Figure 74: Option C Example 594
Figure 75: OSPF Areas 733
Figure 76: Split Horizon Group Example 932
Figure 77: IP Address Assignment with DHCP 938

Preface

About This Guide

This guide describes system concepts and configuration examples to provision logical Input/ Output modules (IOM) cards, Media Dependent Adapters (MDAs), and Ethernet ports on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K2F2T1C, and 7210 SAS-K2F4T6C platforms.

On 7210 SAS devices, not all the Command Line Interface (CLI) commands are supported on all the platforms and in all the modes. In most cases, the CLI commands explicitly mention the list of supported platforms in this guide. In a few cases, it is implied and easy to know the CLIs not supported on a particular platform.

NOTES:

- 7210 SAS-K5 stands for 7210 SAS-K 2F2T1C and 7210 SAS-K12 stands for 7210 SAS-K 2F4T6C platforms.
- 7210 SAS-E, 7210 SAS-D, and 7210 SAS-K 2F2T1C operate in access-uplink mode by default. There is no need of an explicit user configuration needed for this.
  7210 SAS-K 2F4T6C operates in Access-uplink mode and Network mode. There is no explicit BOF configuration required for it.

This document is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as CLI syntax and command usage.

Audience

This manual is intended for network administrators who are responsible for configuring the 7210 SAS-Series routers. It is assumed that the network administrators have an understanding of networking principles and configurations, routing processes, and protocols and standards, including:

- CLI concepts
- MDA and port configuration
- QoS policies
- Services

List of Technical Publications

The 7210 SAS-D, 7210 SAS-E, 7210 SAS-K2F2T1C, and 7210 SAS-K2F4T6C OS documentation set is composed of the following books:

- 7210 SAS-D, 7210 SAS-E, 7210 SAS-K2F2T1C, and 7210 SAS-K2F4T6C OS Basic System Configuration Guide

  This guide describes basic system configurations and operations.
- 7210 SAS-D, 7210 SAS-E, 7210 SAS-K2F2T1C, and 7210 SAS-K2F4T6C OS System Management Guide

  This guide describes system security and access configurations as well as event logging and accounting logs.
- 7210 SAS-D, 7210 SAS-E, 7210 SAS-K2F2T1C, and 7210 SAS-K2F4T6C OS Interface Configuration Guide

  This guide describes card, Media Dependent Adapter (MDA), link aggregation group (LAG) and port provisioning.
- 7210 SAS-D, 7210 SAS-E, 7210 SAS-K2F2T1C, and 7210 SAS-K2F4T6C OS Router Configuration Guide

  This guide describes logical IP routing interfaces and associated attributes such as an IP address, port, as well as IP and MAC-based filtering.
- 7210 SAS-K2F4T6C OS MPLS Guide

  This guide describes how to configure Multi-protocol Label Switching (MPLS) and Label Distribution Protocol (LDP).
- 7210 SAS-D, 7210 SAS-E, 7210 SAS-K2F2T1C, and 7210 SAS-K2F4T6C OS OS Services Guide

  This guide describes how to configure service parameters such as customer information and user services.
- 7210 SAS-D, 7210 SAS-E, 7210 SAS-K2F2T1C, and 7210 SAS-K2F4T6C OS OAM and Diagnostic Guide

  This guide describes how to configure features such as service mirroring and Operations, Administration and Management (OAM) tools.
- 7210 SAS-D and 7210 SAS-E OS OS Quality of Service Guide

  This guide describes how to configure Quality of Service (QoS) policy management.
- 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C Quality of Service Guide

  This guide describes how to configure Quality of Service (QoS) policy management.
- 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C7210 OS Routing Protocols Guide

This guide provides an overview of routing concepts and provides configuration examples for OSPF, IS-IS and route policies.

Getting Started

In This Chapter

This book provides process flow information to configure provision services.

Alcatel-Lucent 7210 SAS Services Configuration Process

Table 1 lists the tasks necessary to configure subscriber services and configure mirroring. This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 1: Configuration Process

Area	Task	Chapter
Subscribers	Subscriber services	
	Global entities	Configuring Global Service Entities with CLI on page 77
	VLL services	VLL Service Considerations on page 198
	VPLS service	Virtual Private LAN Service on page 277
	IES service	Internet Enhanced Service
	VPRN service	Virtual Private Network Service on 7210 SAS-K 2F4T6C on page 581
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and Protocol Support on page 947

SERVICES OVERVIEW

In This Section

This section provides an overview of the 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K-Series subscriber services, service model, and service entities. Additional details on the individual subscriber services supported on different 7210 platforms and their configuration options can be found in subsequent chapters.

Topics in this chapter include:

- Introduction on page 25
  - Service Types on 7210 SAS-D and 7210 SAS-E on page 26
  - Service Types on 7210 SAS-K 2F2T1C on page 26
  - Service Types on 7210 SAS-K 2F4T6C on page 26
  - Service Policies on 7210 SAS-D and 7210 SAS-E on page 27
  - Service Policies on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C on page 29
- Nokia Service Model on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C on page 30
- Alcatel-Lucent Service Model on 7210 SAS-K 2F4T6C on page 31
- Service Entities on 7210 SAS-D, 7210 SAS-E, and 7210 SAS- K 2F2T1C on page 32
- Service Entities on 7210 SAS-K 2F4T6C on page 33
- Customers on page 33
- Service Access Points (SAPs) on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C on page 34
- Service Access Points (SAPs) on 7210 SAS-K 2F4T6C on page 35

In This Section

- SAP Encapsulation Types and Identifiers on page 37
  - Ethernet Encapsulations supported on 7210 SAS-E on page 37
  - Ethernet Encapsulations supported on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C on page 38
  - Services and SAP Encapsulations on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C on page 39
  - Default SAP on a Dot1q Port on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C on page 39
  - Default SAPs on a QinQ Port on a 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C on page 40
  - SAP Configuration Considerations applicable for SAPs configured on access-uplink ports on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C on page 43
  - QinQ SAP Configuration restrictions for 7210 SAS in Network mode only on page 54
- G.8032/Ethernet Ring Protection Switching on page 56
- Service Creation Process Overview on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C on page 73
  - Overview of G.8032 Operation on page 56
  - G8302- Ethernet Ring Sub-Rings on page 62
  - G8302- Virtual and Non-Virtual Channel on page 64
  - G8302- OAM Considerations on page 69
  - G8302- QoS Considerations on page 69
  - G8302- Service and Solution Combinations on page 69
  - G8302- Configuration guidelines on page 70
  - Deploying and Provisioning Services on page 74
  - Configuration Notes on page 75

Introduction

A service is a globally unique entity that refers to a type of connectivity service for either Internet or VPN connectivity. Each service is uniquely identified by a service ID and an optional service within a service area. The 7210 SAS-Series service model uses logical service entities to construct a service. In the service model, logical service entities provide a uniform, service-centric configuration, management, and billing model for service provisioning.

On the 7210 SAS-E and 7210 SAS-D platforms, services can provide Layer 2/bridged service between a service access point (SAP) and another service access point (a SAP is where traffic enters and exits the service) on the same (local) router. It cannot support distributed services using MPLS uplinks.

On the 7210 SAS-K 2F2T1C platforms, services can provide Layer 2/bridged service between a service access point (SAP) and another service access point (a SAP is where traffic enters and exits the service) on the same (local) router. It cannot support distributed services using MPLS uplinks.

On the 7210 SAS-K 2F4T6C platforms, services can provide Layer 2/bridged service or Layer 3/ IP routed connectivity between a service access point (SAP) on one router and another service access point (a SAP is where traffic enters and exits the service) on the same (local) router or another router (distributed). It supports the use of either MPLS uplinks or Ethernet uplinks. 7210 SAS-K 2F4T6C supports both local and distributed service. A distributed service spans more than one router. Distributed services use service distribution points (SDPs) to direct traffic to another Nokia router through a service tunnel. SDPs are created on each participating router, specifying the origination address (the router participating in the service communication) and the destination address of another router. SDPs are then bound to a specific customer service. Without the binding process, far-end router is not able to participate in the service (there is no service without associating an SDP with a service).

Service Types on 7210 SAS-D and 7210 SAS-E

The 7210 SAS-D and 7210 SAS-E offers the following types of subscriber services described in more detail in the referenced chapters:

- Virtual Leased Line (VLL) services:
  - Ethernet pipe (Epipe) — A Layer 2 point-to-point VLL service for Ethernet frames. See Ethernet Pipe (Epipe) Services on page 146.
- Virtual Private LAN Service (VPLS) — A Layer 2 multipoint-to-multipoint VPN bridging service or VPN (using QinQ uplinks).See Virtual Private LAN Service on page 277.

Service Types on 7210 SAS-K 2F2T1C

The 7210 SAS-K 2F2T1C offers the following types of subscriber services described in more detail in the referenced chapters:

- Virtual Leased Line (VLL) services:
  - Ethernet pipe (Epipe) — A Layer 2 point-to-point VLL service for Ethernet frames. See Ethernet Pipe (Epipe) Services on page 146.
- Virtual Private LAN Service (VPLS) — A Layer 2 multipoint-to-multipoint VPN bridging service or VPN (using QinQ uplinks). See Virtual Private LAN Service on page 277.

Service Types on 7210 SAS-K 2F4T6C

The 7210 SAS-K 2F4T6C offers the following types of subscriber services described in more detail in the referenced chapters:

- Virtual Leased Line (VLL) services:
  - Ethernet pipe (Epipe) — A Layer 2 point-to-point VLL service for Ethernet frames. See Ethernet Pipe (Epipe) Services on page 146.
- Virtual Private LAN Service (VPLS) — A Layer 2 multipoint-to-multipoint VPN bridging service or VPN (using QinQ uplinks). See Virtual Private LAN Service on page 277.

- Routed VPLS service (RVPLS) - A service that integrates routing and bridging into a single service. It provides the capability to provide Layer-2 multipoint-to-multipoint services locally on the node, typically using an IP interface for uplink connectivity. See Virtual Private LAN Service on page 277.
- Internet Enhanced Services (IES) - A direct Internet access service where the customer is assigned an IP interface for Internet connectivity. See the Internet Enhanced Service on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C on page 493 for more information about IES.
- Virtual Private Routed Network (VPRN) - A Layer 3 IP multipoint-to-multipoint VPN service as defined in RFC 2547bis. See the Virtual Private LAN Service on page 277 for more information about VPRN services.

Service Policies on 7210 SAS-D and 7210 SAS-E

Common to connectivity services on both 7210 SAS-D and 7210 SAS-E platforms are policies assigned to the service. Policies are defined at a global level and then applied to a service on the router. Policies are used to define 7210 SAS-Series service enhancements. The types of policies common to all 7210 SAS-Series connectivity services are:

- SAP Quality of Service (QoS) policies allow different classes of traffic within a service at SAP ingress. Access egress QoS policies allow differential treatment of various traffic classes within a service (SAPs) which exists in an egress port.

  QoS ingress and egress policies determine the QoS characteristics for a SAP. A QoS ingress policy applied to a SAP specifies the number of meters, meter characteristics (such as forwarding class, committed, and peak information rates, etc.) and the mapping of traffic to a forwarding class. A QoS egress policy defines the queue characteristics (such as CBS, CIR, PIR). A QoS policy must be created before it can be applied to a SAP. A single ingress QoS policy can be associated with a SAP. A single access egress QoS policy can be associated with a port.
- Filter policies allow selective blocking of traffic matching criteria from ingressing or egressing a SAP.

  Filter policies, also referred to as access control lists (ACLs), control the traffic allowed in or out of a SAP based on MAC or IP match criteria. Associating a filter policy on a SAP is optional. Filter policies are identified by a unique filter policy ID. A filter policy must be created before it can be applied to a SAP. A single ingress and single egress filter policy can be associated with a SAP.
- Scheduler policies define the operating parameters (such as scheduling algorithm, weights per priority). Depending on the platform, these are either associated with SAPs or physical ports.
- Accounting policies define how to count the traffic usage for a service for billing purposes.

The routers provide a comprehensive set of service-related counters. Accounting data can be collected on a per-service, per-forwarding class basis that enables network operators to accurately measure network usage and bill each customer for each individual service using any of several billing models.

Service Policies on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

Common to connectivity services on both 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, are policies assigned to the service. Policies are defined at a global level and then applied to a service on the router. Policies are used to define 7210 SAS-Series service enhancements. The types of policies common to all 7210 SAS-Series connectivity services are:

- SAP Quality of Service (QoS) policies allow different classes of traffic within a service at SAP ingress and at SAP egress.

  QoS ingress and egress policies determine the QoS characteristics for a SAP. A QoS ingress and egress policies applied to a SAP specifies the number of queue, queue characteristics (such as forwarding class, committed, and peak information rates, etc.) and the mapping of traffic to a forwarding class. A QoS policy must be created before it can be applied to a SAP. A single ingress and egress QoS policy can be associated with a SAP.
- Filter policies allow selective blocking of traffic matching criteria from ingressing and egressing a SAP.

  Filter policies, also referred to as access control lists (ACLs), control the traffic allowed in or out of a SAP based on MAC or IP match criteria. Associating a filter policy on a SAP is optional. Filter policies are identified by a unique filter policy ID. A filter policy must be created before it can be applied to a SAP. A single ingress and single egress filter policy can be associated with a SAP.
- Accounting policies define how to count the traffic usage for a service for billing purposes.

  The routers provide a comprehensive set of service-related counters. Accounting data can be collected on a per-service, per-forwarding class basis that enables network operators to accurately measure network usage and bill each customer for each individual service using any of several billing models.

Nokia Service Model on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

In the Alcatel-Lucent service model, the service edge routers are deployed at the provider edge. Services are provisioned on the service routers and transported across an IP and/or IP/MPLS provider core network in encapsulation tunnels created using generic router encapsulation MPLS label switched paths (LSPs). 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C support only QinQ and Dot1q Layer 2 uplinks, which are used to transport the services to the provider edge in a hierarchal configuration. The 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C do not support transport tunnels that use MPLS LSPs or GRE SDPs. 7210 SAS-K 2F4T6Csupports the use of only MPLS LSPs (it does not support the use of GRE SDPs).

The service model uses logical service entities to construct a service. The logical service entities are designed to provide a uniform, service-centric configuration, management, and billing model for service provisioning. Some benefits of this service-centric design include:

- Many services can be bound to a single customer.
- QoS policies, filter policies, and accounting policies are applied to each service instead of correlating parameters and statistics from ports to customers to services.

Service provisioning uses logical entities to provision a service where additional properties can be configured for bandwidth provisioning, QoS, security filtering, accounting/billing to the appropriate entity.

Alcatel-Lucent Service Model on 7210 SAS-K 2F4T6C

In the Alcatel-Lucent service model, the service edge routers are deployed at the provider edge. Services are provisioned on the service routers and transported across an IP and/or IP/MPLS provider core network in encapsulation tunnels created using generic router encapsulation MPLS label switched paths (LSPs). 7210 SAS-K 2F4T6C support transport tunnels that use MPLS LSPs or QinQ/Dot1q Layer 2 uplinks. These tunnels are used to transport the services to the provider edge in a hierarchical configuration.

The service model uses logical service entities to construct a service. The logical service entities are designed to provide a uniform, service-centric configuration, management, and billing model for service provisioning. Some benefits of this service-centric design include:

- Many services can be bound to a single customer.
- Many services can be bound to a single tunnel.
- Tunnel configurations are independent of the services they carry.
- Changes are made to a single logical entity rather than multiple ports on multiple devices. It is easier to change one tunnel rather than several services.
- The operational integrity of a logical entity (such as a service tunnel and service end points) can be verified rather than dozens of individual services improving management scaling and performance.
- On 7210 SAS platforms, a failure in the network core can be correlated to specific subscribers and services.
- QoS policies, filter policies, and accounting policies are applied to each service instead of correlating parameters and statistics from ports to customers to services.

Service provisioning uses logical entities to provision a service where additional properties can be configured for bandwidth provisioning, QoS, security filtering, accounting/billing to the appropriate entity.

Service Entities on 7210 SAS-D, 7210 SAS-E, and 7210 SAS- K 2F2T1C

The basic logical entities in the service model used to construct a service are:

- Customers (see page 33)
- Service Access Points (SAPs) on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C (see page 34)

Service Entities on 7210 SAS-K 2F4T6C

The basic logical entities in the service model used to construct a service are:

- Customers on page 33
- Service Access Points (SAPs) on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C on page 34
- Service Access Points (SAPs) on 7210 SAS-K 2F4T6C on page 35
- SAP Encapsulation Types and Identifiers on page 37
- Service Distribution Points (SDPs) on 7210 SAS-K 2F4T6C on page 50
- G.8032/Ethernet Ring Protection Switching on page 61
- Layer 2 Control Processing (L2CP) on page 71
- Service Creation Process Overview on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C on page 73

Customers

The terms customers and subscribers are used synonymously. The most basic required entity is the customer ID value assigned when the customer account is created. To provision a service, a customer ID must be associated with the service at the time of service creation.

Service Access Points (SAPs) on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

Each subscriber service type is configured with at least one service access point (SAP). A SAP identifies the customer interface point for a service on an Alcatel-Lucent 7210 SAS router (Figure 1). The SAP configuration requires that slot, MDA, and port information be specified. The slot, MDA, and port parameters must be configured prior to provisioning a service (see the Cards, MDAs, and Ports sections of the 7210 SAS OS Interface Configuration Guide).

A SAP is a local entity to the router and is uniquely identified by:

- The physical Ethernet port
- The encapsulation type
- The encapsulation identifier (ID)

Depending on the encapsulation, a physical port can have more than one SAP associated with it. SAPs can only be created on ports designated as “access” or “access uplink” in the physical port configuration. SAPs can be created on ports designated as core facing “access uplink” ports. These ports have a different set of features enabled in software.

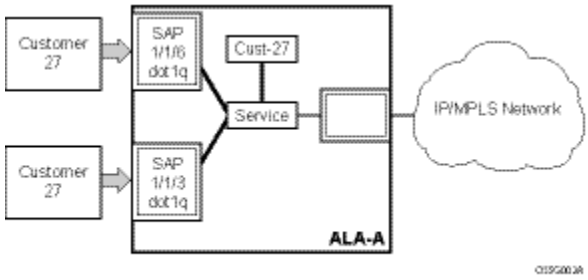

Figure 1: Multiple SAPs in a service using QinQ uplinks in 7210 SAS configured in access-uplink mode

The Figure 1 shows a SAP used for customer service delivery with access-uplink SAPs (also known as, QinQ SAPs) used for service transport on 7210 SAS devices that support only L2 uplinks (also known as, access-uplink mode platforms).

Service Access Points (SAPs) on 7210 SAS-K 2F4T6C

Each subscriber service type is configured with at least one service access point (SAP). A SAP identifies the customer interface point for a service on an Alcatel-Lucent 7210 SAS router (Figure 2). The SAP configuration requires that slot, MDA, and port information be specified. The slot, MDA, and port parameters must be configured prior to provisioning a service (see the Cards, MDAs, and Ports sections of the 7210 SAS OS Interface Configuration Guide).

A SAP is a local entity to the router and is uniquely identified by:

- The physical Ethernet port
- The encapsulation type
- The encapsulation identifier (ID)

Depending on the encapsulation, a physical port can have more than one SAP associated with it. SAPs can only be created on ports designated as “access” in the physical port configuration.

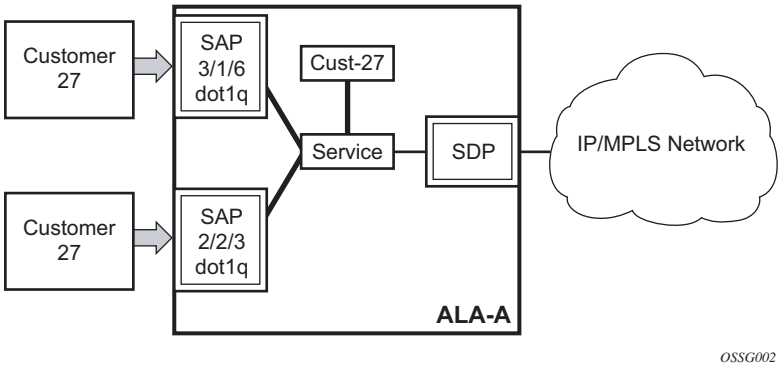

Figure 2: Service Access Point (SAP) for 7210 SAS configured with MPLS uplinks

The Figure 2 shows a SAP used for customer service delivery with SDP used for service transport on 7210 SAS devices that support MPLS uplinks.

SAPs can be created on ports designated as core facing “access uplink” ports when using QinQ uplinks. Access uplink ports have a different set of features enabled in software.

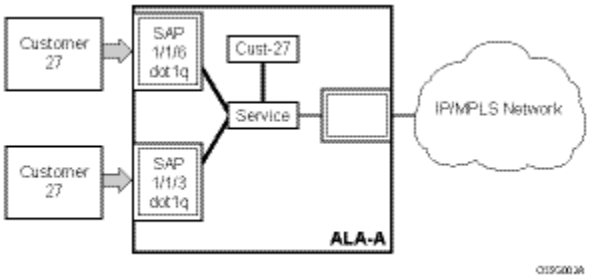

Figure 3: Multiple SAPs in a service using QinQ uplinks in 7210 SAS configured in access-uplink mode

The Figure 3 shows a SAP used for customer service delivery with access-uplink SAPs (also known as, QinQ SAPs) used for service transport on 7210 SAS devices when using L2 uplinks.

SAP Encapsulation Types and Identifiers

The encapsulation type is an access property of a service Ethernet port. The appropriate encapsulation type for the port depends on the requirements to support multiple services on a single port on the associated SAP and the capabilities of the downstream equipment connected to the port. For example, a port can be tagged with IEEE 802.1Q (referred to as dot1q) encapsulation in which each individual tag can be identified with a service. A SAP is created on a given port by identifying the service with a specific encapsulation ID.

Ethernet Encapsulations supported on 7210 SAS-E

The following lists encapsulation service options on Ethernet ports:

- Null — Supports a single service on the port. For example, where a single customer with a single service customer edge (CE) device is attached to the port. The encapsulation ID is always 0 (zero).
- Dot1q — Supports multiple services for one customer or services for multiple customers (Figure 4). For example, the port is connected to a customer who wants multiple services. The encapsulation ID used to distinguish an individual service is the VLAN ID in the IEEE 802.1Q header.

The following lists encapsulation service options on Ethernet access uplink ports:

- QinQ — The QinQ encapsulation type adds a IEEE 802.1Q tag to the 802.1Q tagged packets entering the network to expand the VLAN space by tagging tagged packets, producing a double tagged frame. On the 7210 SAS-E, QinQ encapsulation is supported only on access uplink ports.

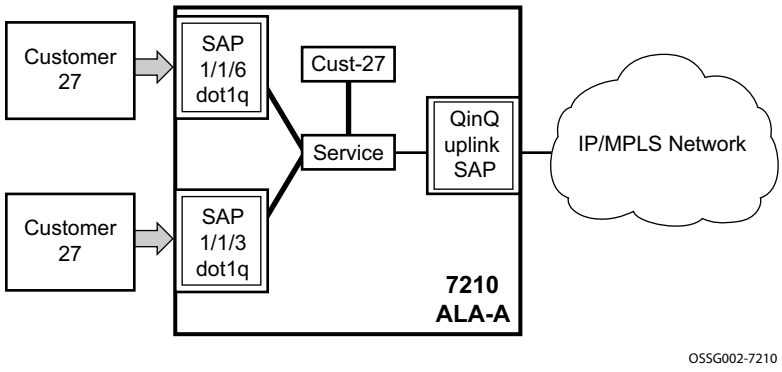

Figure 4: Multiple SAPs on a Single Port

Ethernet Encapsulations supported on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

The following lists encapsulation service options on Ethernet access ports:

- Null - Supports a single service on the port. For example, where a single customer with a single service customer edge (CE) device is attached to the port. The encapsulation ID is always 0 (zero).
- Dot1q - Supports multiple services for one customer or services for multiple customers (Figure 5). For example, the port is connected to a customer who wants multiple services. The encapsulation ID used to distinguish an individual service is the VLAN ID in the IEEE 802.1Q header.
- QinQ - The QinQ encapsulation type adds a IEEE 802.1Q tag to the 802.1Q tagged packets entering the network to expand the VLAN space by tagging tagged packets, producing a double tagged frame.

The following lists encapsulation service options on Ethernet access uplink ports:

- QinQ - The QinQ encapsulation type adds a IEEE 802.1Q tag to the 802.1Q tagged packets entering the network to expand the VLAN space by tagging tagged packets, producing a double tagged frame.

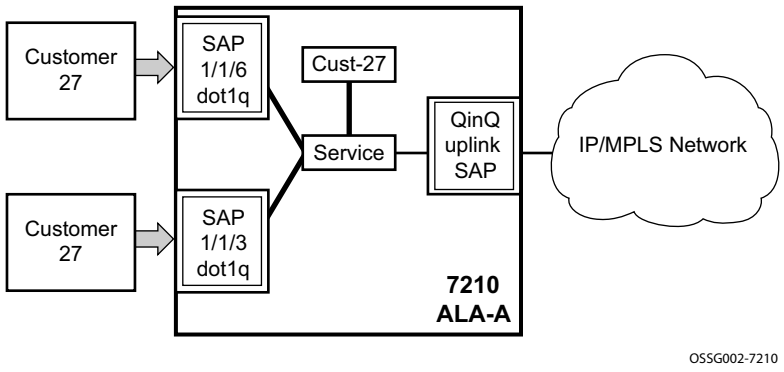

Figure 5: Multiple SAPs on a Single Port

Services and SAP Encapsulations on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

Table 3 lists the service and SAP Encapsulation information for Ethernet ports:

Table 2: Service and Encapsulation for access port

Port Type	Encapsulation	7210 SAS Platforms Support
Ethernet	Null	7210 SAS-E, 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C
Ethernet	Dot1q	7210 SAS-E, 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C
Ethernet	QinQ	7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

Table 3 lists the service and SAP Encapsulation information for Ethernet access uplink ports.

Table 3: Port Type and Encapsulation for access uplink ports

Port Type	Encapsulation	7210 Platforms
Ethernet access uplink	QinQ	7210 SAS-E, 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

Default SAP on a Dot1q Port on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This feature introduces default SAP functionality on Dot1q-encapsulated ports. On a dot1q-encapsulated port where a default SAP is configured, all packets with q-tags not matching any explicitly defined SAPs will be assigned to this SAP. SAPs with default Dot1q encapsulation are supported in VPLS and Epipe services. In this context, the character "*" indicates default which means allow through. The default SAP also accepts untagged or priority tagged packets. A default SAP must be configured explicitly. When a default SAP is not configured explicitly, packets not matching any explicitly defined SAPs will be dropped.

One of the applications where this feature can be applicable is an access connection of a customer who uses the whole port to access Layer 2 services. The internal VLAN tags are transparent to the service provider. This can be provided by a null encapsulated port.

In this type of environment, logically two SAPs exist, a management SAP and a service SAP. The management SAP can be created by specifying a VLAN tag which is reserved to manage the CPE. The service SAP covers all other VLANs and behaves as a SAP on a null-encapsulated port.

There a few constraints related for the use of default SAP on a Dot1q-encapsulated port:

- This type of SAP is supported only on VPLS and Epipe services and cannot be created in IES services as it cannot preserve VLAN tag markings.
- For VPLS SAPs with STP enabled, STP listens to untagged and null-tagged BPDUs only. All other tagged BPDUs are forwarded like other customer packets. This is the same behavior as null-encapsulated ports.
- IGMP snooping is not supported on a default SAP. This would require remembering VLAN tags per hosts. By not allowing IGMP snooping of this SAP, all IGMP packets will be transparently forwarded.

Default SAPs on a QinQ Port on a 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

Default QinQ SAPs (notation - *.*) are used in ring ports to avoid the need to configure services on all the intermediate nodes in the ring which are transiting the service. Default QinQ SAPs matches all VLAN tagged traffic which is not classified into any other SAP configured on the same port. Only one EPIPE service with default QinQ SAPs is needed for transit service traffic on access-uplink ports. Default QinQ SAPs are allowed only on access-uplink ports and access ports. It can co-exist with 0.* SAP on an access-uplink or access port. A default QinQ SAP accepts only tagged packets. Untagged packets or priority tagged packets are not accepted on Default QinQ SAPs. 7210 SAS-K, accepts untagged and tagged packets on a default QinQ SAP.

When an EPIPE service with default QinQ SAPs on the ring ports is used for transit traffic in a ring deployment, no protection mechanism (example: STP or G.8032) is supported for Default QinQ SAPs. The upstream or head-end node on which the service originates must ensure the correct path on the ring is selected using either G.8032 or STP. When a VPLS service with default QinQ SAPs on the ring ports is used for transit traffic in a ring deployment, users can use either G.8032 or M-VPLS with xSTP for ring protection. When using G.8032, the state of the default QinQ SAPs in the VPLS service can be managed using a separate G.8032 control instance.
NOTE: G.8032 control instance cannot use Default QinQ SAPs.

Following features are available for use with Default QinQ SAPs configured in Epipe and VPLS service (unless explicitly specified, below listed features are applicable for both Epipe and VPLS service):

For Default QinQ SAPs on either access ports or access-uplink ports:

- On 7210 SAS-D and 7210 SAS-E, a default QinQ SAP accepts only tagged packets. Untagged packets or priority tagged packets are not accepted on Default QinQ SAPs.
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, a default QinQ SAP accepts both untagged and tagged packets.
- On 7210 SAS-D and 7210 SAS-E, Default QinQ SAP is available for use only in an Epipe and a VPLS service created with svc-sap-type parameter set to “null-star”. Default QinQ SAP can be configured along with other SAPs allowed in the same service (that is, service with svc-sap-type parameter set to “null-star”).
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, default QinQ SAP is available for use only in an Epipe and a VPLS service created with svc-sap-type parameter set to “any”. Default QinQ SAP can be configured along with other SAPs allowed in the same service (that is, service with svc-sap-type parameter set to “any”).
- MAC learning and aging is available for use in a VPLS service
- Per SAP MAC limit is available for use in a VPLS service
- Mac-move detection and Mac-pinning is available for use in a VPLS service
- Discard-unknown and discard-unknown-source is available for use in a VPLS service
- ETH-CFM and Y.1731 is not available for use
- STP (and all its different flavors) cannot be enabled in the service with Default QinQ SAPs
- MVPLS with xSTP can be used for loop prevention. The Default QinQ SAPs inherit the state from the associated MVPLS instance.
- G.8032 control instance cannot be configured in a service with Default QinQ SAP
- G.8032 can be used for loop prevention in ring deployments, where the Default QinQ SAPs are configured on the ring ports in a VPLS service. A separate G.8032 control instances needs to be configured for use on the ring ports and the service with Default QinQ ports needs to be associated with this G.8032 control instance
- IGMP snooping is not available for use
- L2PT and BPDU translation is not available for use
- On 7210 SAS-E and 7210 SAS-D, IP interface in a VPLS service is not supported in a service using this SAP. Note on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, IP interface in a VPLS service is not supported (irrespective of SAP encapsulations).

For Default QinQ SAPs created on Access-uplink Port:

- Ingress qos policy applied on an access uplink port is available for classification and policing on ingress.
- Egress qos policy applied on an access uplink port is available for egress queue shaping, scheduling and marking.
- SAP Ingress ACLs are available for use
- SAP Egress ACLs are not available for use
- SAP Ingress received count and SAP Egress forwarded count are available for use (appropriate accounting records can be used)

For Default QinQ SAPs created on access ports:

- SAP ingress qos policy is available for use.
- On 7210 SAS-D and 7210 SAS-E - Egress qos policy applied on an access port is available for egress shaping, scheduling and marking.
- On 7210 SAS-K 2F2T1C and 7210 SAS- K 2F4T6C - Egress qos policy applies on an access SAP is available for egress shaping, scheduling and marking.
- SAP Ingress ACLs are available for use
- SAP egress ACLs are not available for use
- On 7210 SAS-D and 7210 SAS-E - SAP Ingress Meter counters, SAP Ingress received count and SAP Egress forwarded counter are available for use (appropriate accounting records can be used).
- On 7210 SAS-K 2F2T1C and 7210 SAS- K 2F4T6C - SAP Ingress Queue counters, SAP Ingress received count and SAP Egress forwarded counter are available for use (appropriate accounting records can be used).

Configuration Notes for use of Default QinQ SAPs for transit service in a ring deployment on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

- If an Epipe service is used with Default QinQ SAPs on the ring ports for transit service in a ring deployment, no protection mechanism is available for the transit service (that is, Epipe service with the Default QinQ SAPs on ring ports). Both Epipe and VPLS services which are originating on different nodes in the ring can use the transit service. Protection/ Loop-detection mechanisms can be implemented for VPLS service configured in the ring nodes, by using MVPLS with xSTP on the nodes where the VPLS service is configured. No protection mechanisms are available for use with Epipe services on the node that originates the service.
- If a VPLS service is used with Default QinQ SAPs on the ring ports for transit service in a ring deployment, either MVPLS/xSTP or G.8032 can be used to protect the transit service (that is, VPLS service with the Default QinQ SAPs on ring ports). In this case, VPLS service which are originating on different nodes in the ring and use the transit VPLS service are also protected. Epipe services which are originating on different nodes in the ring cannot use the transit VPLS service.
- When using VPLS service with Default QinQ SAPs for transit service with either G.8032 or MVPLS with xSTP configured for protection, load-balancing of the traffic based on the VLAN IDs is not possible. If load-balancing is desired then it is better to use Epipe service with Default QinQ SAPs as the transit service.

SAP Configuration Considerations applicable for SAPs configured on access-uplink ports on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

NOTE: 7210 SAS-K 2F4T6C supports use of access-uplink mode and network mode simultaneously without requiring BOF configuration. In other words, on 7210 SAS-K 2F4T6C user can configure some ports as access-uplink ports and some ports as network ports simultaneously.

When configuring a SAP, consider the following:

- A SAP is a local entity and only locally unique to a given device. The same SAP ID value can be used on another 7210 SAS-Series.
- On 7210 SAS-D and 7210 SAS-E, a physical port can have only one SAP to be part of one service. Multiple SAPS can be defined over a physical port but each of these SAPs should belong to a different service.

- On 7210 SAS-K, multiple SAPs configured on the same port can be part of the same service.
- 7210 SAS-K supports use of Q1.0 SAP. This SAP matches packets received on a port with the outermost tag being Q1 and the inner tag being absent (that is, no tag) or the inner tag is a priority tag. It does not accept packets with any other VLAN tag value as the inner tag.
- There are no default SAPs configured on the node. All SAPs in subscriber services must be created.
- The default administrative state for a SAP at creation time is administratively enabled.
- When a SAP is deleted, all configuration parameters for the SAP will also be deleted.
- An access SAP is owned by and associated with the service in which it is created in each router.
- A port with a dot1q encapsulation type means the traffic for the SAP is identified based on a specific IEEE 802.1Q VLAN ID value. The VLAN ID is stripped off at SAP ingress and the appropriate VLAN ID placed on at SAP egress. As a result, VLAN IDs only have local significance, so the VLAN IDs for the SAPs for a service need not be the same at each SAP.

  NOTE: Some exceptions to this are dot1q range SAPs and Dot1q preserve SAPs configured on a port with dot1q encapsulation.

- If a port is administratively shutdown, all SAPs on that port will be operationally out of service.
- A SAP cannot be deleted until it has been administratively disabled (shutdown).
- Each SAP can have one each of the following policies assigned:
  - Ingress filter policy
  - Egress filter policy
  - Ingress QoS policy
  - Accounting policy
  - Egress QoS policy on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C
- An ingress QoS policy and accounting policy is assigned per access uplink port and cannot be assigned per access uplink SAP. In other words, all access-uplink SAPs on a given port share the QoS policy.
- Ingress filter policy and egress filter policy is available per access-uplink SAP.

- The `svc-sap-type' parameter value determines the type of SAPs provisioned in a service.
  - The Table 4, “SAP and Service Combinations for 7210 SAS-E,” on page 46 on page 36 provides details of SAP and service combinations allowed in access-uplink mode for 7210 SAS-E devices.
  - The Table 5, “SAP and Service Combinations for 7210 SAS-D,” on page 47 provides details of SAP and service combinations allowed in 7210 SAS-D devices.
  - The Table 6, “SAP and Service Combinations for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C,” on page 49 provides details of SAP and service combinations allowed in 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices.
- If a service’s sap-type is specified as dot1q-preserve, all the SAPs configured in the service must have the same VLAN ID. The outermost VLAN tag of the packets received on access port is not stripped, when svc-sap-type is set to dot1q-preserve.
- For dot1q range SAP, the outermost VLAN tag is compared against the range of VLAN IDs configured and mapped to the SAP if there is a match. The outermost tag is not stripped on SAP ingress and is forwarded through the other service endpoint.
- L2PT cannot be configured for use on all the configured SAPs simultaneously. The number of SAPs which can use this simultaneously is lesser than maximum amount of SAPs supported by the node.

Configuration Guidelines for 7210 SAS-E

1. For 7210 SAS-E devices, the traffic usage on a SAP can be monitored by enabling a counter. The counter collects the total number of packets or octets forwarded across the SAP. The user can configure the mode of the counter to collect data in either packets or octets. The mode of the counter can be specified for ingress SAP counters only. The egress SAP counters collect only the number of packets. The accounting records collect the count of packets or octets on a SAP using the accounting logs. The default mode of the ingress SAP counter is set to packets. The egress SAP counter is disabled by default.
2. The mode of the counter cannot be changed if an accounting policy is already associated with a SAP.
3. Ensure that egress SAP counters are enabled before associating accounting records that count egress forwarded packets.
4. Before modifying the counter, disable the account log generation, run the command:
   no collect-stats
5. Egress SAP statistics are not available on any of the SAPs of a port, on which a dot1q SAP and dot1q default SAP configuration are present at the same time. This limitation also applies for egress ACLs. In other words, egress ACLs are not supported on either dot1q SAP and Dot1q Default SAP when both these SAPs are configured on a port simultaneously.

6. Egress SAP statistics cannot be configured for use simultaneously on all the configured SAPs. The number of SAPs which can use this feature simultaneously, is less than the maximum amount of SAPs supported by the node.
7. There is no limit to the number of access ports allowed to be configured. In other words all ports can be configured as access ports. The number of access-uplink ports that can be configured is limited. In other words, only a subset of ports can be configured as access-uplink ports at a given time.
8. Service MTU is not supported on 7210 SAS-E.

The following table list the SAPs allowed with different value of svc-sap-type on 7210 SAS-E.

Table 4: SAP and Service Combinations for 7210 SAS-E

svc-sap-type	Access SAPs	Access Uplink SAPs
null-star	Null SAP, dot1q Default SAP	Q.* SAP, 0.* SAP
dot1q-preserve	dot1q SAP (dot1q VLAN tag not stripped on ingress)	Q1.Q2
dot1q	dot1q SAP, dot1q explicit null SAP	Q.* SAP, 0.* SAP

Configuration Guidelines for 7210 SAS-D

1. Ensure that egress SAP counters are enabled before associating accounting records that count egress forwarded packets.
2. Before modifying the counter, disable the account log generation, run the command:
   no collect-stats
3. Egress SAP statistics are not available on any of the SAPs of a port, on which a dot1q SAP and dot1q default SAP configuration are present at the same time. This limitation also applies for egress ACLs. In other words, egress ACLs are not supported on either dot1q SAP and Dot1q Default SAP when both these SAPs are configured on a port simultaneously.
4. Egress SAP statistics cannot be configured to use simultaneously on all the configured SAPs. The number of SAPs which can use this feature simultaneously is less than the maximum amount of SAPs supported by the node.
5. Service MTU is not supported on 7210 SAS-D.

The following table list the SAPs allowed with different value of svc-sap-type on 7210 SAS-D.

Table 5: SAP and Service Combinations for 7210 SAS-D

svc-sap-type	Access SAPs	Access Uplink SAPs
null-star	Null SAP, dot1q Default SAP Q.* SAP, 0.* SAP Default QinQ SAP (*.* SAP)	Q.* SAP, 0.* SAP Default QinQ SAP (*.* SAP)
dot1q-preserve	dot1q SAP (dot1q VLAN tag not stripped on ingress), Q1.Q2 SAP (Q2 tag VLAN must match the dot1q SAP VLAN ID)	Q1.Q2 SAP (Q2 tag VLAN ID must match the dot1q SAP VLAN ID)
any	null SAP, dot1q SAP, dot1q explicit null SAP Q1.Q2 SAP, Q.* SAP, 0.* SAP	Q1.Q2 SAP, Q.* SAP, 0.* SAP
dot1q-range	Dot1q SAP (dot1q VLAN tag not stripped on ingress), Q1.* SAP	Q1.* SAP

NOTES:

- Dot1q Default SAP cannot be configured when svc-sap-type is set to ‘any’.
- When svc-sap-type is set to ‘any’ for a NULL SAP, the system processes and forwards only packets with no VLAN tag (that is, untagged). All other packets with one or more VLAN tags (even those with priority tag only) are not processed and dropped. Users can use the service with svc-sap-type set to ‘null-star’, to process and forward packets with one or more tags (including priority tag) on a NULL SAP.
- The Default QinQ SAP processes only tagged packets received on a QinQ port. All tagged packets that do not match the specific SAP tags configured on the same port are processed by this SAP. The Default QinQ SAP cannot process untagged packets, even if 0.* SAP is not configured for use on that port.
- The Default QinQ SAPs is available for use with 0.* SAPs configured on the same port or in the same service. It is available for use with another default QinQ SAP configured in the same service (on a different port).
- In a VPLS service, the Default QinQ SAP is available for use with any other SAP type configured in a service configured with svc-sap-type parameter set to “null-star”.
- SAPs using connection-profile (to specify dot1q VLAN ranges) can be configured in a service only when svc-sap-type is set to 'dot1q-range'.

- When a service is configured to use svc-sap-type 'dot1q-range', the outermost V-LAN tag of the packets are not stripped when the packet is received on access port ingress. For more information on processing behavior for this type of service, see "Ethernet Pipe (Epipe) Services on page 146 section.
- Service MTU is not supported on 7210 SAS-D.
- The following counters are available on 7210 SAS-D devices:
  - Ingress and egress counters per SAP
  - Ingress policer counters per SAP
  - Egress queue counters per access port
  - Ingress and egress counters per access-uplink port
  - The number of counters available to count total received packets or octets on an access-uplink SAP (in a VPLS, VLL or IES service) ingress is limited and hence count of received packets or octets cannot be obtained for all the SAPs simultaneously. By default, these counts is not available.
  - The CLI command **config>service>epipe>sap>statistics>ingress** is available to associate a counter with the SAP and obtain the counts.
  - The number of counters available to count total forwarded packets or octets on an access SAP egress and access-uplink SAP egress is limited and hence count of received packets or octets cannot be obtained for all the SAPs simultaneously. By default, these counts is not available.
  - The CLI command *config>service>epipe>sap>statistics>ingress>forwarded-count* is available to associate a counter with the SAP and obtain the counts.

Configuration Guidelines for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

The 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C support the following types of SAPs:

  - **Access SAPs** – Null, Dot1q (Dot1q, Dot1q Default, Dot1q range, Dot1q explicit NULL), QinQ SAPs (Q1.Q2, Q1. *, Q1.0 QinQ Default SAP (that is, *.* SAP), 0.*) is supported.
  - **Access Uplink SAPs** – QinQ SAPs (various SAP types such as Q1.Q2, QinQ Default SAP (*.* SAP), 0.*, Q1.*, Q1.0) is supported
- 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C supports Q1.0 SAP. This SAP accepts a packet with the outermost tag as Q1 or a packet with outermost tag as Q1 and the following inner tag is a priority tag. Unlike 7x50, it does not accept packets with 2 tags with the outermost tag being Q1 and the inner tag being any tag other than priority tag.
- 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, allows any port to be configured in either access uplink mode or access mode. Additionally the ports can be in access uplink mode or access mode or they can be mix of the ports using either modes. There is no limit

to the number of access ports allowed to be configured. In other words all ports can be configured as access ports. The number of access-uplink ports depends on the number of QoS resources allocated per port. In other words, not all the ports can be configured as access-uplink ports at a given time.

- 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C support service MTU.
  - A received frame/packet length is checked against the configured service MTU after subtracting the length of the SAP encapsulation (including the L2 header) from the received frame length. The packet is further processed in the context of the service, if the computed length is less than equal to the configured service MTU or the packet is dropped.
  - The user must configure the correct service MTU across all the nodes, if support is available, through which the service is transported.
- 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C support only the following svc-sap type parameters:
  - **'any'** – A service configured with this value for svc-sap-type allows for configuration of all combination of access SAPs and access-uplink SAPs in the same service, except for dot1q range SAPs. A packet that is received with tags more than the number of SAP tags to which it is mapped to, is forwarded transparently in the service (the processing behavior is similar to any other packet mapped to the SAP).
  - **'dot1q-range'** – A service configured with this value for svc-sap-type allows for configuration of dot1q range SAPs and Q1.* access-uplink SAP in the same service.

The following table list the SAPs allowed on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C with different values of svc-sap-type.

Table 6: SAP and Service Combinations for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

svc-sap-type	Access SAPs	Access Uplink SAPs
any	null SAP, dot1q SAP, dot1q explicit null SAP dot1q default SAP Q1.Q2 SAP, Q.* SAP, Q1.0 SAP, 0.* SAP QinQ default SAP (*.* SAP)	Q1.Q2 SAP, Q.* SAP, Q1.0 SAP, 0.* SAP, QinQ default SAP (*.* SAP)
dot1q-range	Dot1q SAP (dot1q VLAN tag not stripped on ingress), Q1.* SAP	Q1.* SAP

Service Distribution Points (SDPs) on 7210 SAS-K 2F4T6C

NOTE: SDPs are supported only on 7210 SAS-K 2F4T6C. It is not supported on 7210 SAS-E, SAS-D, and SAS-K 2F2T1C.

A service distribution point (SDP) acts as a logical way to direct traffic from one router to another through a uni-directional (one-way) service tunnel. The SDP terminates at the far-end device which directs packets to the correct service egress SAPs on that device. A distributed service consists of a configuration with at least one SAP on a local node, one SAP on a remote node, and an SDP binding the service to the service tunnel.

An SDP has the following characteristics:

- An SDP is locally unique to a participating router. The same SDP ID can appear on other 7210 SAS-Series routers.
- An SDP uses the system IP address to identify the far-end edge router.
- An SDP is not specific to any one service or any type of service. Once an SDP is created, services are bound to the SDP. An SDP can also have more than one service type associated with it.
- All services mapped to an SDP use the same transport encapsulation type defined for the SDP.
- An SDP is a management entity. Even though the SDP configuration and the services carried within are independent, they are related objects. Operations on the SDP affect all the services associated with the SDP. For example, the operational and administrative state of an SDP controls the state of services bound to the SDP.

An SDP from the local device to a far-end router requires a return path SDP from the far-end 7210 SAS-Series back to the local router. Each device must have an SDP defined for every remote router to which it wants to provide service. SDPs must be created first, before a distributed service can be configured.

SDP Binding

To configure a distributed service from ALA-A to ALA-B, the SDP ID (1) must be specified in the service creation process to “bind” the service to the tunnel (the SDP). Otherwise, service traffic is not directed to a far-end and the far-end device(s) cannot participate in the service (there is no service). To configure a distributed service from ALA-B to ALA-A, the SDP ID (5) must be specified.

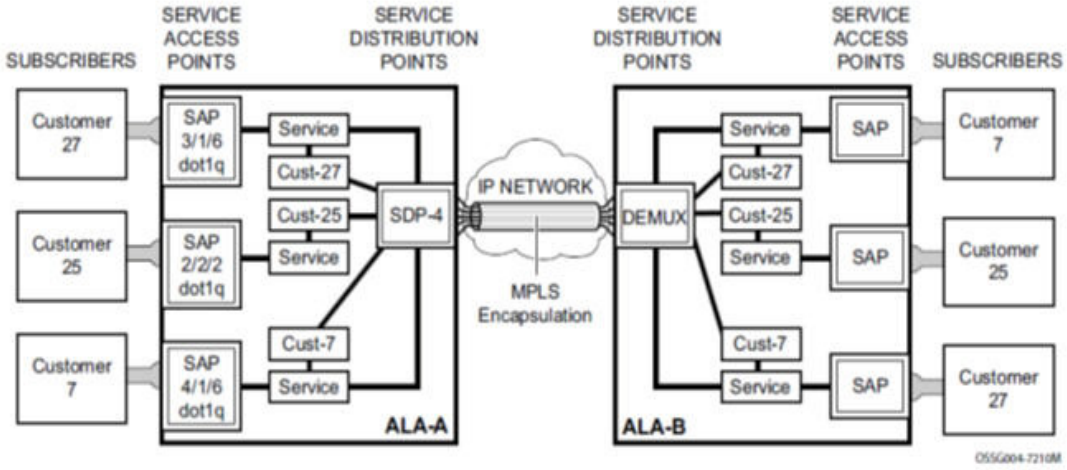

Figure 6: MPLS Service Distribution Point (SDP) pointing from ALA-A to ALA-B

Spoke and Mesh SDPs

When an SDP is bound to a service, it is bound as either a spoke SDP or a mesh SDP. The type of SDP indicates how flooded traffic is transmitted. The 7210 SAS Network Mode devices supports both spoke and mesh SDPs.

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” and not transmitted on the port it was received.

All 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.

SDP Using BGP Route Tunnel

SDP is enhanced to use BGP route tunnel to extend inter-AS support for L2 and L3 VPN services. An SDP can be configured to use MPLS transport method. MPLS SDP support is enhanced to allow a BGP route tunnel to reach the far-end PE. A single method of tunneling is allowed per

SDP (for example, LDP, RSVP-TE LSP or BGP route tunnel). BGP route tunnel method is excluded if multi-mode transport is enabled for an SDP.

For inter-AS far-end PE, next-hop for BGP route tunnel must be one of the local ASBR. The LSP type selected to reach the local ASBR (BGP labeled route next-hop) must be configured under the BGP global context. LDP must be supported to provide transport LSP to reach the BGP route tunnel next-hop.

Only BGP route labels can be used to transition from ASBR to the next-hop ASBR. The global BGP route tunnel transport configuration option must be entered to select an LSP to reach the PE node from ASBR node. On the last BGP segment, both "BGP+LDP" and LDP routes may be available to reach the far-end PE from the ASBR node. LDP LSP must be preferred due to higher protocol priority. This leads to just one label besides other labels in stack to identify VC/VPN at far-end PE nodes.

SDP Keepalives

SDP keepalives actively monitor the SDP operational state using periodic Alcatel-Lucent SDP ping echo request and echo reply messages. Alcatel-Lucent SDP ping is a part of Alcatel-Lucent's suite of service diagnostics built on an Alcatel-Lucent service-level OA&M protocol. When SDP ping is used in the SDP keepalive application, the SDP echo request and echo reply messages are a mechanism for exchanging far-end SDP status.

Configuring SDP keepalives on a given SDP is optional. SDP keepalives for a SDP have the following configurable parameters:

- Admin up/admin down state
- Hello time
- Message length
- Max drop count
- Hold down time

SDP keepalive echo request messages are only sent when the SDP is completely configured and administratively up and SDP keepalives is administratively up. If the SDP is administratively down, keepalives for the SDP are disabled.

SDP keepalive echo request messages are sent out periodically based on the configured Hello Time. An optional message length for the echo request can be configured. If max drop count echo request messages do not receive an echo reply, the SDP will immediately be brought operationally down.

If a keepalive response is received that indicates an error condition, the SDP will immediately be brought operationally down.

Once a response is received that indicates the error has cleared and the hold down time interval has expired, the SDP will be eligible to be put into the operationally up state. If no other condition prevents the operational change, the SDP will enter the operational state.

For information about configuring keepalive parameters, refer to Configuring an SDP on page 63.

SDP Administrative Groups

This feature introduces the support of SDP administrative groups, referred to as SDP admin groups. SDP admin groups provide a way for services using a PW template to automatically include or exclude specific provisioned SDPs. SDPs sharing a specific characteristic or attribute can be made members of the same admin group.

The user first creates the admin groups used by SDPs on this node:

config>service>sdp-group>group-name group-name value group-value create

A maximum of 32 admin groups can be created. The no option is only allowed if the group-name is not referenced in a pw-template or SDP.

The group value ranges from zero (0) to 31. It is uniquely associated with the group name at creation time. If the user attempts to configure another group name for a group value that is already assigned to an existing group name, the SDP admin group creation is failed. The same happens if the user attempts to configure an SDP admin group with a new name but associates it to a group value already assigned to an existing group name.

Next, the user configures the SDP membership in admin groups:

config>service>sdp>sdp-group *group-name*

The user can enter a maximum of one (1) admin group name at once. The user can execute the command multiple times to add membership to more than one admin group. The admin group name must have been configured or the command is failed. Admin groups are supported on an SDP and of type MPLS (BGP/RSVP/LDP). They are also supported on an SDP with the mixed-lsp-mode option enabled.

The user then selects which admin groups to include or exclude in a given PW template:

config>service>pw-template>sdp-include *group-name*

config>service>pw-template>sdp-exclude *group-name*

The admin group name must have been configured or the command is failed. The user can execute the command multiple times to include or exclude more than one admin group. The sdp-include and sdp-exclude commands can only be used with the use-provisioned-sdp option. If the same group name is included and excluded within the same PW template, only the exclude option will be enforced.

Any changes made to the admin group sdp-include and sdp-exclude constraints will only be reflected in existing spoke-sdps after the following command has been executed:

tools>perform>service>eval-pw-template>allow-service-impact

When the service is bound to the PW template, the SDP selection rules will enforce the admin group constraints specified in the sdp-include and sdp-exclude commands.

config>service>vpls>bgp>pw-template-binding policy-id

config>service>epipe>spoke-sdp-fec>pw-template-bind policy-id

Note that the group value is what is used to uniquely identify an SDP admin group throughout the network in the 5620 SAM. The node will send both the group name and value to 5620 SAM, or other SNMP device, at the creation of the SDP admin group. In all other operations in the node, such as adding an SDP to an admin group or including/excluding an SDP admin group in a service context, only the group name is sent to the 5620 SAM or the SNMP device.

SDP admin groups can be enabled on all 7210 services that make use of the PW template (that is, BGP-AD VPLS service.

Mixed-LSP Mode of Operation

The mixed LSP SDP allows for a maximum of two LSP types to be configured within an SDP. A primary LSP type and a backup LSP type. An RSVP primary LSP type can be backed up by an LDP LSP type.

An LDP LSP can be configured as a primary LSP type which can then be backed up by a BGP LSP type. At any given time, the service manager programs only one type of LSP in the line-card that will activate it to forward service packets according to the following priority order:

1. **RSVP LSP type**. One RSVP LSP can be configured per SDP. This is the highest priority LSP type.
2. **LDP LSP type**. One LDP FEC will be used per SDP. 7210 SAS does not support LDP ECMP.

3. **BGP LSP type**. One RFC 3107-labeled BGP prefix programmed by the service manager.

In the case of the RSVP/LDP SDP, the service manager will program the NHLFE(s) for the active LSP type preferring the RSVP LSP type over the LDP LSP type. If no RSVP LSP is configured or all configured RSVP LSPs go down, the service manager will re-program the line-card with the LDP LSP if available. If not, the SDP goes operationally down.

When a higher priority type LSP becomes available, the service manager reverts back to this LSP at the expiry of the revert-time timer or the failure of the currently active LSP, whichever comes first. The service manager then re-programs the line-card accordingly. If the infinite value is configured, then the SDP reverts to the highest priority type LSP only if the currently active LSP failed.

Note however, that LDP uses a tunnel down damp timer which is set to three seconds by default. When the LDP LSP fails, the SDP will revert to the RSVP LSP type after the expiry of this timer. For an immediate switchover this timer must be set to zero.

Use the **configure>router>ldp>tunnel-down-damp-time** command. For more information, see 7210 SAS OS MPLS User Guide.

If the value of the revert-time timer is changed, it will take effect only at the next use of the timer. Any timer which is outstanding at the time of the change will be restarted with the new value.

In the case of the LDP/BGP SDP, the service manager will prefer the LDP LSP type over the BGP LSP type. The service manager will re-program the line card with the BGP LSP if available otherwise it brings down the SDP operationally.

Also note the following difference in behavior of the LDP/BGP SDP compared to that of an RSVP/LDP SDP. For a given /32 prefix, only a single route will exist in the routing table: the IGP route or the BGP route. Thus, either the LDP FEC or the BGP label route is active at any given time. The impact of this is that the tunnel table needs to be re-programmed each time a route is deactivated and the other is activated. Furthermore, the SDP revert-time cannot be used as there is no situation where both LSP types are active for the same /32 prefix.

G.8032/Ethernet Ring Protection Switching

Ethernet ring protection switching offers ITU-T G.8032 specification compliance to achieve resiliency for Ethernet Layer 2 networks.G.8032 (Eth-ring) is built on Ethernet OAM and often referred to as Ring Automatic Protection Switching (R-APS).

Eth-rings are supported on VPLS SAPs. VPLS services supporting Rings SAPs can connect to other rings and Ethernet service using VPLS, SAPs. Eth-rings enables rings for core network or access network resiliency. A single point of interconnection to other services is supported. The Eth-ring service is a VLAN service providing protection for ring topologies and the ability to interact with other protection mechanisms for overall service protection. This ensures failures detected by Eth-ring only result in R-APS switchover when the lower layer cannot recover and that higher layers are isolated from the failure.

Rings are preferred in data networks where the native connectivity is laid out in a ring or there is a requirement for simple resilient LAN services. Due to the symmetry and the simple topology, rings are viewed a good solution for access and core networks where resilient LANS are required. The Alcatel-Lucent implementation can be used for interconnecting access rings and to provide traffic engineered backbone rings.

Eth-rings use one VID per control per ring instance and use one (typically) or multiple VIDs for data instances per control instance. A dedicated control VLAN (ERP VLAN) is used to run the protocol on the control VID. G.8032 controls the active state for the data VLANs (ring data instances) associated with a control instance. Multiple control instances allow logically separate rings on the same topology. The Alcatel-Lucent implementation supports dot1q, and qinq encapsulation for data ring instances. The control channel supports dot1q and qinq encapsulation.

Note: In 7210 SAS-D and 7210 SAS-K devices, CCMs used for G.8032 Ethernet ring protection service is implemented in hardware. In 7210 SAS-E, CCMs used for G.8032 Ethernet ring protection service is implemented in software.

Overview of G.8032 Operation

R-APS messages that carry the G.8032 protocol are sent on dedicated protocol VLAN called ERP VLAN (or Ring Control Instance). In a revertive case, G.8032 Protocol ensures that one Ring Protection Link (RPL) owner blocks the RPL link. R-APS messages are periodically sent around in both directions to inform other nodes in the Ring about the blocked port in the RPL owner node. In non-revertive mode any link may be the RPL link.Y.1731 Ethernet OAM CC is the basis of the RAPs messages. Y.1731 CC messages are typically used by nodes in the ring to monitor the health of each link in the ring in both directions. However, CC messages are not mandatory. Other link

layer mechanisms could be considered – for example LOS (Loss of Signal) when the nodes are directly connected.

Initially each Ring Node blocks one of its links and notifies other nodes in the ring about the blocked link. Once a ring node in the ring learns that another link is blocked, the node unblocks its blocked link possibly causing FDB flush in all links of the ring for the affected service VLANs, controlled by the ring control instance. This procedure results in unblocking all links but the one link and the ring normal (or idle) state is reached. In revertive mode the RPL link will be the link that is blocked when all links are operable after the revert time. In non-revertive mode the RPL link is no different that other ring links. Revertive mode offers predictability particularly when there are multiple ring instances and the operator can control which links are block on the different instances. Each time there is a topology change that affects Reachability, the nodes may flush the FDB and MAC learning takes place for the affected service VLANs, allowing forwarding of packets to continue. Figure 7 depicts this operational state:

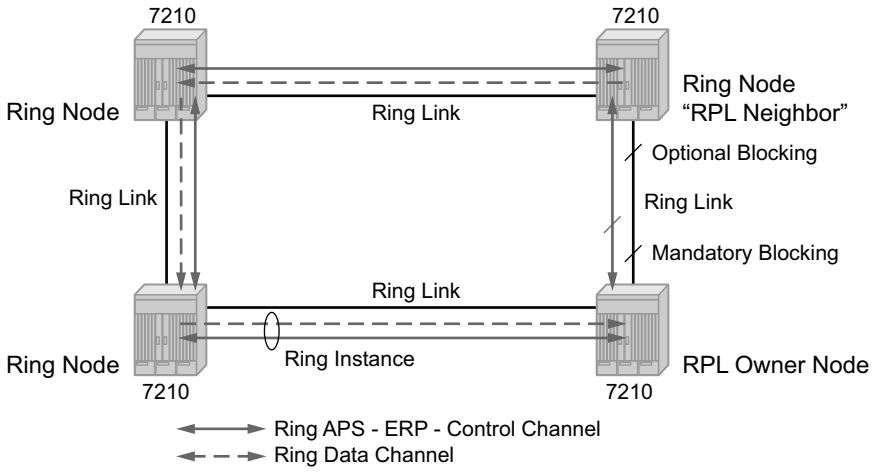

Figure 7: G.8032 Ring in the Initial State

When a ring failure occurs, a node or nodes detecting the failure (enabled by Y.1731 OAM CC monitoring) send R-APS message in both directions. This allows the nodes at both ends of the failed link to block forwarding to the failed link preventing it from becoming active. In revertive mode, the RPL Owner then unblocks the previously blocked RPL and triggers FDB flush for all nodes for the affected service instances. The ring is now in protecting state and full ring connectivity is restored. MAC learning takes place to allow Layer 2 packet forwarding on a ring. The following picture depicts the failed link scenario.

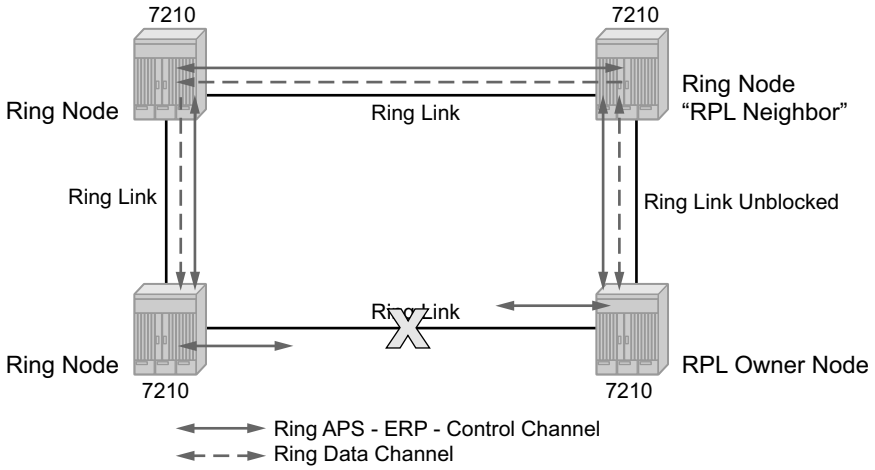

Figure 8: G.8032 Ring in the Protecting State

Once the failed link recovers, the nodes that blocked the link again send the R-APS messages indicating no failure this time. This in turn triggers RPL Owner to block the RPL link and indicate the Blocked RPL link the ring in R-APS message, which when received by the nodes at the recovered link cause them to unblock that link and restore connectivity (again all nodes in the ring perform FDB Flush and MAC learning takes place). The ring is back in the normal (or idle) state.

Within each path, Y.1731 Maintenance Entity Group (MEG) Endpoints (MEPs) are used to exchange R-APS specific information (specifically to co-ordinate switchovers) as well as optionally fast Continuity Check Messages (CCM) providing an inherent fault detection mechanism as part of the protocol. Failure detection of a ring path by one of the mechanisms triggers to activate the protection links. Upon failure, re-convergence times are a dependent on the failure detection mechanisms. In the case of Y.1731, the CCM transmit interval determines the response time. The 7210 SAS device supports 100ms (millisecond) message timers that allows for quicker restoration times. Alternatively, 802.3ah (Ethernet in the First Mile) or simple Loss of Signal can act as a trigger for a protection switch where appropriate. In case of direct connectivity between the nodes, there is no need to use Ethernet CC messaging for liveliness detection.

Revertive and non-revertive behaviors are supported. The Ring protection link (RPL) is configured and Eth-rings can be configured to revert to the RPL upon recovery.

G.8032 supports multiple data channels (VIDs) or instances per ring control instance (R-APS tag). G.8032 also supports multiple control instances such that each instance can support RPLs on different links providing for a load balancing capability however once services have been assigned to one instance the rest of the services that need to be interconnected to those services must be on the same instance. In other words, each data instance is a separate data VLAN on the same physical topology. When there is any one link failure or any one node failure in the ring, G.8032 protocols are capable of restoring traffic between all remaining nodes in these data instances.

Ethernet R-APS can be configured on any port configured for access mode using dot1q, q-in-q encapsulation enabling support for Ethernet R-APS protected services on the service edge towards the customer site, or within the Ethernet backbone. ELINE and ELAN services can be afforded Ethernet R-APS protection and, although the Ethernet Ring providing the protection uses a ring for protection the services are configured independent of the Ring properties. The intention of this is to cause minimum disruption to the service during Ethernet R-APS failure detection and recovery.

In the 7210 SAS implementation, the Ethernet Ring is built from a VPLS service on each node with VPLS SAPs that provides Ring path with SAPs. As a result, most of the VPLS SAP features are available on Ethernet rings if desired. This results in a fairly feature rich ring service.

The control tag defined under each eth-ring is used for encapsulating and forwarding the CCMs and the G.8032 messages used for the protection function. If a failure of a link or node affects an active Ethernet ring segment, the services will fail to receive the CC messages exchanged on that segment or will receive a fault indication from the Link Layer OAM module.

For fault detection using CCMs, three CC messages plus a configurable hold-off timer must be missed for a fault to be declared on the associated path. The latter mechanism is required to accommodate the existence of additional, 50 ms resiliency mechanism in the optical layer. After it receives the fault indication, the protection module will declare the associated ring link down and the G.8032 state machine will send the appropriate messages to open the RPL and flush the learned addresses.

Flushing is triggered by the G.8032 state machine and the 7210 SAS implementation allows flooding of traffic during the flushing interval to expedite traffic recovery.

The following diagram shows an example G8032 ring. The 0-3 Ring Example on page 60 below illustrates a resilient Ring Service. In the ring example, a QinQ ring (solid line) using VID 500 carries two customer VLANs Dot1q 100 and QinQ 400.1 respectively). The RPL for the G8032 ring is between A and B where B is the RPL owner. Also, illustrated in the figure below is a QinQ service on the (dotted line) ring that uses Dot1q VID 600 for the ring to connect service VLAN 100.50. The two rings have RPLs on different nodes which allow a form of load balancing. The example serves to illustrate that service encapsulations and ring encapsulation can be mixed in various combinations. Also, note that neither of the rings is a closed loop. A ring can restore connectivity when any one node or link fails to all remaining nodes within a small amount of transfer time (signaling time after detection).

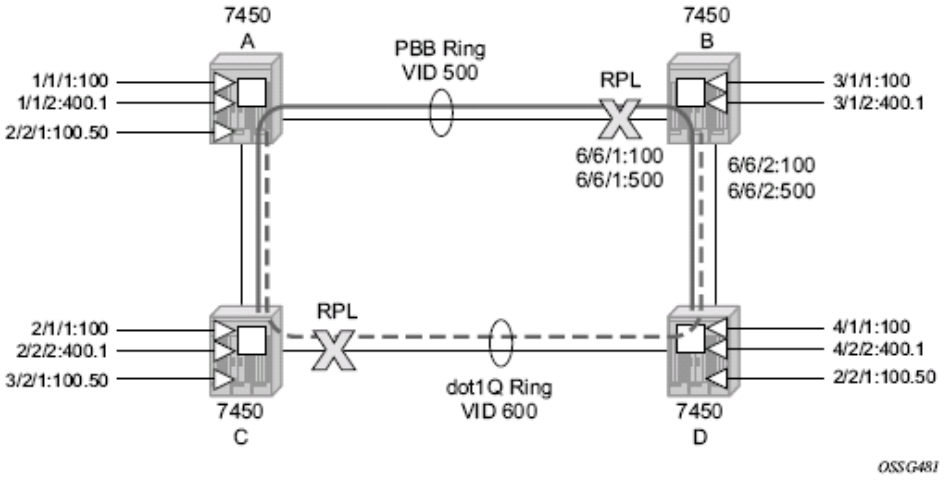

Figure 9: 0-3 Ring Example

The following is a configuration example for G8032 ring for the figure illustrated above:

Step #1 - Configure G8032 ring Paths and Control MEPs used for failure detection on ring ports.

Sample Configuration:

```
configure eth-ring 1
     description "Ring PBB BLUE on Node B"
     revert-time 100
     guard-time 5
     ccm-hold-time down 100 up 200
     rpl-node owner
     path a 1/1/1 raps-tag 100 // CC Tag 100
          description "To A ring link"
          rpl-end
          eth-cfm
               mep 1 domain 1 association 1 direction down  // Control MEP
                    no shutdown
               ccm-enable
               control-mep
               exit
          exit
          no shutdown  // would allow protect switching
                // in absence of the "force" cmd
     exit
     path b 6/6/2 raps-tag 100 //Tag 100
          description "to D Ring Link"
          eth-cfm
               mep 1 domain 1 association 1 direction down
                    no shutdown
                         ccm-enable
                         control-mep
```

```
             exit
         exit
         no shutdown
     exit
no shutdown
exit
```

Step #2 - Configure VPLS service used for G8032 control instances (identified with VID tag 100 and 500) used to exchange R-APS messages for the 2 control instances created on the ring.

```
configure> service
     vpls 10 customer 1 create // Ring APS SAPs
          description "Ring Control VID 100"
          sap 1/1/1:100 eth-ring 1 create  // TAG for the Control Path a
          exit
          sap 6/6/2:100 eth-ring 1 create  // TAG for the Control Path b
          exit
          no shutdown
     exit
configure> service
     vpls 40 customer 1 create //Data Channel on Ring
          description "Ethernet Ring 1 VID 500"
          sap 1/1/1:500 eth-ring 1 create   // TAG for the Data Channel Path a
          exit
          sap 6/6/2:500 eth-ring 1 create   // TAG for the Data Channel Path b
          exit
     exit
```

Step #3 - Configure VPLS data services that will use G8032 for protection.

```
configure> service vpls 1001 // CPE traffic
          sap 3/1/2:400.1 create
          // CPE SAP
          exit

          sap 6/6/1:2000 create
          // uplink SAP
          exit

          no shutdown
     exit

configure> service vpls 1001 // CPE traffic
          sap 2/2/1:100.50 create
          // CPE SAP
          exit

          sap 6/6/1:600 create
          // uplink SAP
          exit

          no shutdown
     exit
```

G8302- Ethernet Ring Sub-Rings

Ethernet Sub-Rings offer a dual redundant way to interconnect rings. The 7210 SAS supports Sub-Rings connected to major rings and a sub-ring connected to a VPLS (LDP based) for access rings support in VPLS networks. Figure 10 illustrates a Major ring and Sub Ring scenario. In this scenario, any link can fail in either ring (ERP1 or ERP2) and each ring is protected. Furthermore, the sub ring (ERP2) relies on the major Ring (ERP1) as part of its protection for the traffic from C and D. The nodes C and D are configured as inter connection nodes.

NOTE: 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C cannot be used as the interconnection nodes. They can be used only as the ring nodes in the sub-ring.

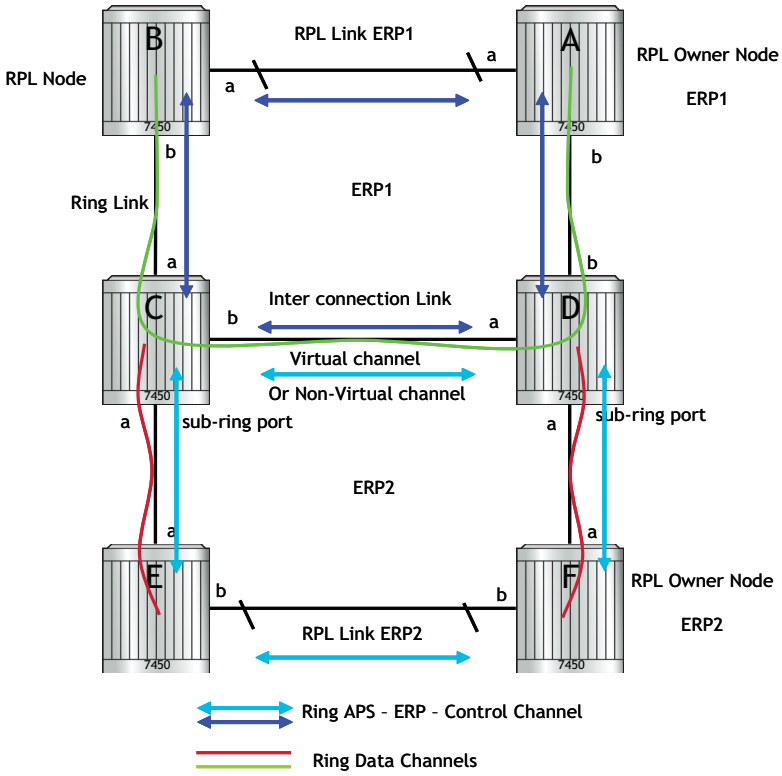

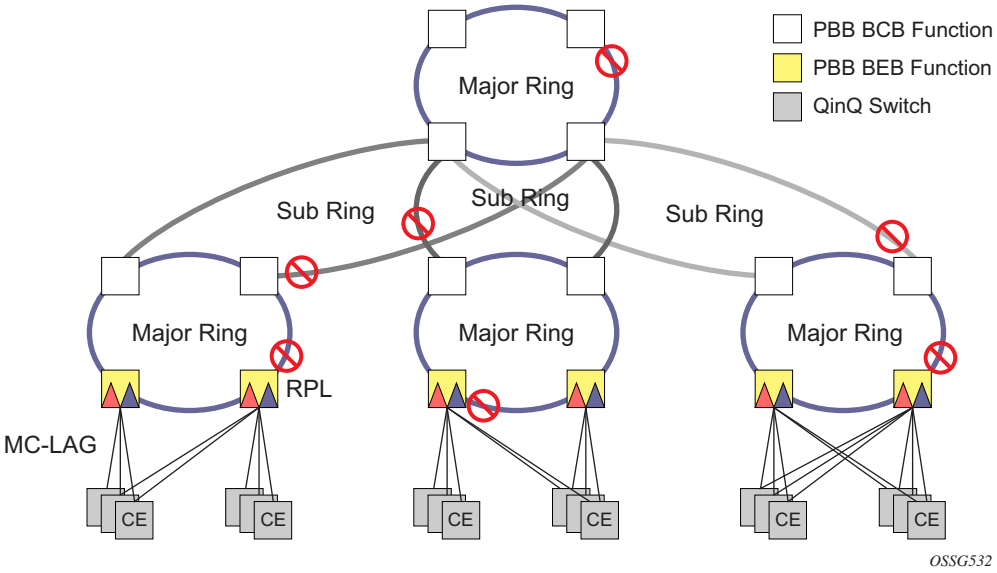

Figure 10: G.8032 Sub-Ring

Sub-Rings and Major Rings run similar state machines for the ring logic, however there are some differences. When Sub-Rings protect a link, the flush messages are propagated to the major ring. (A special configuration allows control of this option on the 7210 SAS.) When major rings change topology, the flush is propagated around the major ring and does not continue to any sub-rings. The reason for this is that Major Rings are completely connected but Sub-Rings are dependent on another ring or network for full connectivity. The topology changes need to be propagated to the other ring or network usually. Sub-Rings offer the same capabilities as major rings in terms of control and data so that all link resource may be utilized.

G8302- Virtual and Non-Virtual Channel

The following illustrates a sample Sub-Ring using virtual-link configuration on Node C, interconnecting node:

```
eth-ring 2
        description "Ethernet Sub Ring on Ring 1"
            interconnect ring-id 1 // Link to Major Ring 1
                propagate-topology-change
            exit
        exit
        path a 1/1/3 raps-tag 100 // Ring control uses VID 100
            eth-cfm
                mep 9 domain 1 association 4
                    ccm-enable
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        no shutdown

    exit

sub-ring non-virtual-link // Not using a virtual link

# Control Channel for the Major Ring ERP1 illustrates that Major ring
# control is still separate from Sub-ring control
  vpls 10 customer 1 create
      description "Control VID 10 for Ring 1 Major Ring"
      stp shutdown
      sap 1/1/1:10 eth-ring 1 create
          stp shutdown
          exit
      sap 1/1/4:10 eth-ring 1 create
          stp shutdown
          exit
      no shutdown
  exit

# Data configuration for the Sub-Ring

  vpls 11 customer 1 create
      description "Data on VID 11 for Ring 1"
      stp shutdown
      sap 1/1/1:11 eth-ring 1 create // VID 11 used for ring
          stp shutdown
      exit
      sap 1/1/4:11 eth-ring 1 create
          stp shutdown
      exit
      sap 1/1/3:11 eth-ring 2 create // Sub-ring data
          stp shutdown
      exit
      sap 3/2/1:1 create
      description "Local Data SAP"
```

```
            stp shutdown
        no shutdown
    exit

# Control Channel for the Sub-Ring using a virtual link. This is
# a data channel as far as Ring 1 configuration. Other Ring 1
# nodes also need this VID to be configured.

    vpls 100 customer 1 create
        description "Control VID 100 for Ring 2 Interconnection"
        split-horizon-group "s1" create //Ring Split horizon Group
        exit
        stp shutdown
        sap 1/1/1:100 split-horizon-group "s1" eth-ring 1 create
            stp shutdown
        exit
        sap 1/1/4:100 split-horizon-group "s1" eth-ring 1 create
            stp shutdown
        exit
        sap 1/1/3:100 eth-ring 2 create
            stp shutdown
        exit
        no shutdown
    exit
```

G8302- Ethernet Ring Sub Ring using non-virtual-link

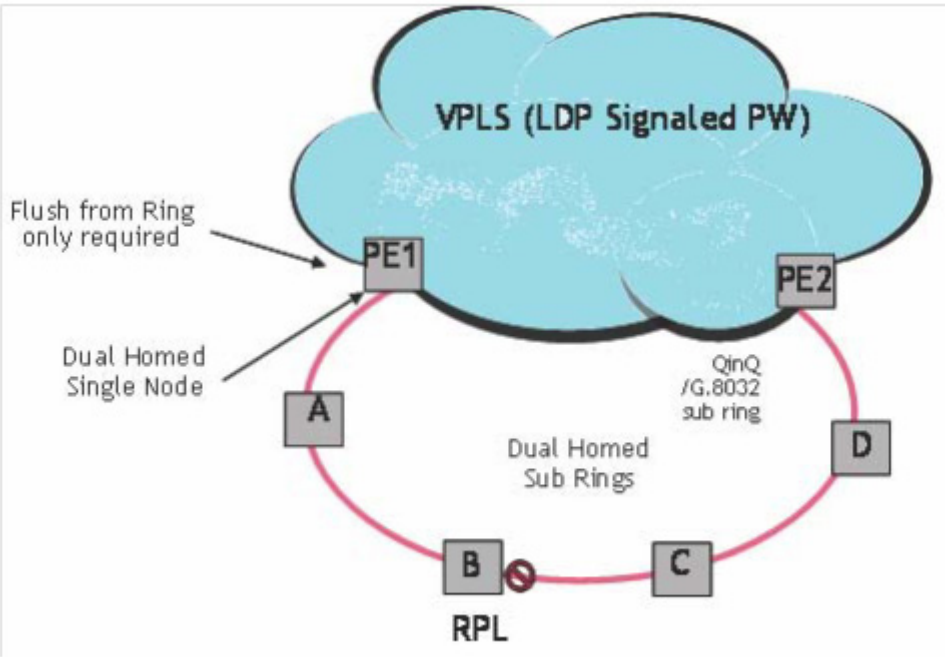

Figure 11: 0-6 Sub-Ring Homed to VPLS

NOTE: In this solution, the 7210 SAS nodes can only be the ring nodes. It cannot be used as the interconnection PE nodes.

The following illustrates a sample Sub-Ring using non-virtual-link configuration on PE1, interconnecting node:

```
eth-ring 1
      description "Ethernet Ring 1"
      guard-time 20
      no revert-time
      rpl-node nbr
      sub-ring non-virtual-link
          interconnect vpls // VPLS is interconnection type
              propagate-topology-change
          exit
      exit
      path a 1/1/3 raps-tag 1.1
```

```
          description "Ethernet Ring : 1 Path on LAG"
          eth-cfm
          mep 8 domain 1 association 8
               ccm-enable
               control-mep
               no shutdown
            exit
        exit
        no shutdown
    exit
    no shutdown
exit
```

All the Sub Ring nodes part of Sub Ring with non-virtual-link should be configured with “sub-ring non-virtual-link” option.

```
eth-ring 1
        sub-ring non-virtual-link
        exit
        path a 1/1/1 raps-tag 1.1
            eth-cfm
                mep 5 domain 1 association 4
                    ccm-enable
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        path b 1/1/2 raps-tag 1.1
            eth-cfm
                mep 6 domain 1 association 3
                    ccm-enable
                    control-mep
                    no shutdown
                exit
            exit
            no shutdown
        exit
        no shutdown
    exit
# Control Channel for Sub-Ring using non-virtual-link on interconnecting node:
vpls 1 customer 1 create
      description "Ring 1 Control termination"
      stp shutdown
      sap 1/1/3:1.1 eth-ring 1 create //path a control
          stp shutdown
      exit
      no shutdown
  exit
# Configuration for the ring data into the VPLS Service

  vpls 5 customer 1 create
      description "VPLS Service at PE1"
      stp
          no shutdown
      exit
      sap 1/1/3:2.2 eth-ring 1 create
```

```
            stp shutdown
        exit
        sap 1/1/5:1 create
        exit
        mesh-sdp 5001:5 create //sample LDP MPLS LSPs
        exit
        mesh-sdp 5005:5 create
        exit
        mesh-sdp 5006:5 create
        exit

        no shutdown
    exit
# Control Channel for Sub-Ring using non-virtual-link on sub-Ring nodes:
vpls 1 customer 1 create
            stp
                shutdown
            exit
            sap 1/1/1:1.1 eth-ring 1 create
                stp
                    shutdown
                exit
            exit
            sap 1/1/2:1.1 eth-ring 1 create
                stp
                    shutdown
                exit
            exit
            no shutdown
          exit
```

The following illustrates a sample Sub-Ring using non-virtual-link configuration homed to a major ring:

```
eth-ring 1
      description "Ethernet Ring 1"
      guard-time 20
      no revert-time
      rpl-node nbr
      sub-ring non-virtual-link
          interconnect ring-id <major ring index>
               propagate-topology-change
          exit
      exit
      path a 1/1/3 raps-tag 1.1
          description "Ethernet Ring : 1 Path on LAG"
          eth-cfm
          mep 8 domain 1 association 8
               ccm-enable
               control-mep
               no shutdown
            exit
        exit
        no shutdown
    exit
    no shutdown
exit
```

G8302- OAM Considerations

Ethernet CFM can be enabled on each individual path under an Ethernet ring. Only down MEPs can be configured on each of them and CCM sessions can be enabled to monitor the liveliness of the path using interval of 100 msec. Different CCM intervals can be supported on the path a and path b in an Ethernet ring. CFM is optional if hardware supports Loss of Signal for example.

UP MEPs on service SAPs which multicast into the service and monitor the active path may be used to monitor services.

NOTE: 7210 SAS-E does not support UP MEPs.

G8302- QoS Considerations

Ethernet ring CC messages transmitted over the SAP queues using the default egress QoS policy will use NC (network class) as a forwarding class. If user traffic is assigned to the NC forwarding class, it will compete for the same bandwidth resources with the Ethernet CCMs. As CCM loss could lead to unnecessary switching of the Ethernet ring, congestion of the queues associated with the NC traffic should be avoided. The operator must configure different QoS Policies to avoid congestion for the CCM forwarding class by controlling the amount of traffic assigned into the corresponding queue.

Note: The operator must configure appropriate ingress QoS policies to ensure that R-APS messages get appropriate QoS treatment and is processed and/or transmitted without delays to enable better failover time.

G8302- Service and Solution Combinations

The Ethernet rings are supported Layer 2 service. The following considerations apply:

- Only ports in access mode and access-uplink mode can be configured as eth-ring paths.
- Dot1q and QinQ ports are supported as eth-ring path members.

G8302- Configuration guidelines

The following are the configuration guidelines for G8032:

- Service level MEPs are not available on all SAPs tied to an eth-ring instance on a port.
- G.8032 instances cannot be configured over a LAG.
- For 7210 SAS-D devices, to improve the service fail-over time due to failures in the ring path, fast flood is enabled by default. On a failure detection in one of the paths of the eth ring, along with MAC flush the system starts to flood the traffic onto the available path.No explicit user configuration is needed for this and it does not affect scaling for filters/ACLs.
- For 7210 SAS-D devices, down MEPs used with services and G.8032 share common hardware resources.
- The 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C device support fast-flood by default to improve service fail-over time. It does not require explicit user configuration and does not share resources with either Filters/ACLS or Service Down MEPs.

Layer 2 Control Processing (L2CP)

Operators providing Epipe and VPLS services need to be able to transparently forward Layer-2 control frames received from the customers. This allows their customers to run these control protocols between the different locations which are part of the L2 VPN service. The 7210 SAS platforms provide user with the following capability:

- An option to tunnel, discard or peer for EFM OAM, LLDP, Dot1x, and LACP.
- BPDU translation and Layer 2 Protocol Tunneling support for xSTP and CISCO control protocols. This is supported only in a VPLS service. For more information, see the "L2PT and BPDU Translation on page 306".

NOTE: The CDP, VTP, DTP, PAgP, and UDLD management protocols are forwarded transparently in an Epipe service.

By default, LACP, LLDP, EFM OAM, and Dot1x Layer-2 control protocol untagged packets are discarded if the protocol is not enabled on the port where these frames are received. User has an option to enable peering by enabling the protocol on the port and configuring the appropriate parameters for the protocol. User also has an option to tunnel these packets using an Epipe or VPLS service.

In a VPLS service, the layer-2 control frames are sent out of all the SAPs configured in the VPLS service. It is recommended to use this feature carefully and only when an VPLS is used to emulate an end-to-end Epipe service (that is, an Epipe configured using a 3-point VPLS Service, with one access SAP and 2 access-uplink SAP/SDPs for redundant connectivity). In other words, if the VPLS service is used for multi-point connectivity, it is not recommended to use this feature. When a layer-2 control frame is forwarded out of dot1q SAP or a QinQ SAP, the SAP tags of the egress SAP are added to the packet.

The following SAPs can be configured for tunneling the untagged L2CP frames (corresponding protocol tunneling needs to be enabled on the port):

- If the port encapsulation is null, user has an option to tunnel these packets by configuring a NULL SAP on a port
- If the port encapsulation is dot1q, user an option to use dot1q explicit null SAP (e.g. 1/1/10:0) or a dot1q default SAP (For example: 1/1/11:*) to tunnel these packets.
- If the port encapsulation is QinQ, user has an option to use 0.* SAP (For example 1/1/10:0.*) to tunnel these packets.

In addition to the protocols listed above, protocols not supported on 7210, For example: GARP, GVRP, ELMI, and others are transparently forwarded in case of a VPLS service. These protocols are transparently forwarded if a NULL SAP, dot1q default SAP, dot1q explicit null SAP or 0.* SAP is configured on the port and received packet is untagged. If the received packet is tagged and matches the tag of any of the SAPs configured on the port, it is forwarded in the context of the

SAP and the service. Else if the received packet is untagged and none of the NULL or dot1q default or dot1q explicit null or 0.* SAP is configured, it is discarded.

If a 7210 receives a tagged L2CP packet on any SAP (includes NULL, dot1q, dot1q range, QinQ, QinQ default), it is forwarded transparently in the service similar to normal service traffic (xSTP processing behavior is different in VPLS service and is listed below).

The xSTP processing behavior in a VPLS service is as follows:

- If xSTP is enabled in the service, and if the tag in the STP BPDU matches the tag of the configured SAP, the received xSTP BPDU is processed by the local xSTP instance on the node for that service when xSTP is enabled on the SAP and discarded when xSTP is disabled on the SAP.
- If the tags do not match, xSTP BPDU packets are transparently forwarded in the service similar to normal service traffic.
- If xSTP is disabled in the service, STP BPDU packets are transparently forwarded in the service similar to normal service traffic

.

Table 7: L2CP for 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K5 and 7210 SAS-12 Platforms

Packet Type	7210 SAS-E	7210 SAS-D	7210 SAS-K5 and K12
LACP	Option to Discard or Peer or Tunnel	Option to Tunnel or Discard or Peer	Option to Tunnel or Discard or Peer
Dot1x	Option to Discard or Peer	Option to Tunnel or Discard or Peer	Option to Tunnel or Discard or Peer
LLDP	Option to Discard or Peer	Option to Tunnel or Discard or Peer (see **Note1**)	Option to Tunnel or Discard or Peer (see **Note1**)
EFM	Option to Tunnel or Discard or Peer	Option to Tunnel or Discard or Peer	Option to Tunnel or Discard or Peer
L2PT	Supported	Supported	Supported
BPDU Tunneling	Supported	Supported	Supported
xSTP	Option to Peer or Tunnel	Option to Peer or Tunnel	Option to Peer or Tunnel

Note1: For more information read the 7210 SAS Interfaces guide to know more about options available for LLDP tunneling.

Service Creation Process Overview on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

NOTE: This section provides overview of service creation process with access-uplink ports. It is applicable to 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

Figure 12 displays the overall process to provision core and subscriber services.

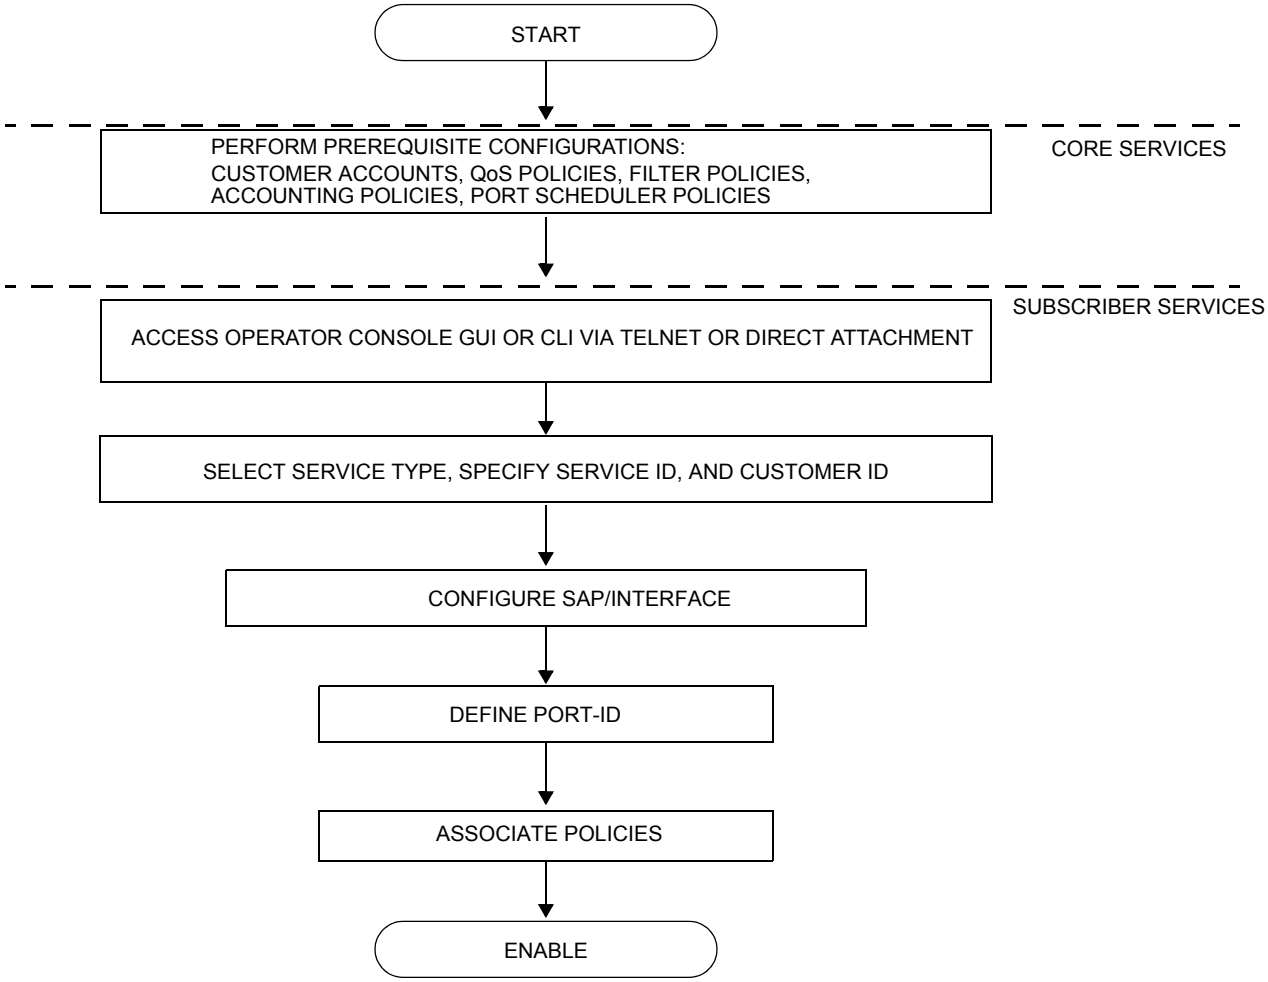

Figure 12: Service Creation and Implementation Flow

Deploying and Provisioning Services

The service model provides a logical and uniform way of constructing connectivity services. The basic steps for deploying and provisioning services can be broken down into three phases.

Phase 1: Core Network Construction

Before the services are provisioned, the following tasks should be completed:

- Build the IP or IP/MPLS core network.
- Configure routing protocols.

Phase 2: Service Administration

Perform preliminary policy configurations to control traffic flow, operator access, and to manage fault conditions and alarm messages, the following tasks should be completed:

- Configure group and user access privileges.
- Build templates for QoS, filter and/or accounting policies needed to support the core services.

Phase 3: Service Provisioning

- Provision customer account information.
- If necessary, build any customer-specific QoS, filter or accounting policies.
- Provision the customer services on the service edge routers by defining SAPs, binding policies to the SAPs.

Configuration Notes

This section describes service configuration caveats.

General

Service provisioning tasks can be logically separated into two main functional areas, core tasks and subscriber tasks and are typically performed prior to provisioning a subscriber service.

Core tasks include the following:

- Create customer accounts
- Create template QoS, filter, scheduler, and accounting policies

Subscriber services tasks include the following:

- Configure interfaces (where required) and SAPs
- Create exclusive QoS and filter policies

Configuring Global Service Entities with CLI

This section provides information to create subscriber (customer) accounts using the command line interface.

Topics include:

- Service Model Entities on page 77
- Configuring Customers on page 79
- ETH-CFM Features on page 107
- Service Management Tasks on page 100

Service Model Entities

The Alcatel-Lucent service model uses logical entities to construct a service. The service model contains four main entities to configure a service.

- Subscribers on page 79
- Services:
  - → Ethernet Pipe (Epipe) Services on page 146
  - → VPLS on page 343
- Service Access Points (SAPs)
  - → Ethernet Pipe (Epipe) Services on page 146
  - → VPLS SAP on page 356

Basic Configuration

The most basic service configuration must have the following:

- A customer ID
- A service type
- A service ID
- A SAP identifying a port and encapsulation value

Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure a customer account

Configuring Customers

The most basic customer account *must* have a customer ID. Optional parameters include:

- Description
- Contact name
- Telephone number

Customer Information

Use the following CLI syntax to create and input customer information:

CLI Syntax:
```
config>service# customer customer-id create
 contact contact-information
 description description-string
 phone phone-number
```

The following displays a basic customer account configuration.

```
A:ALA-12>config>service# info
----------------------------------------
...
       customer 5 create
           description "Nokia Customer"
           contact "Technical Support"
           phone "650 555-5100"
       exit
...
----------------------------------------
A:A:ALA-12>config>service#
```

Service Creation Process Overview on 7210 SAS-K 2F4T6C

Figure 13 displays the overall process to provision core and subscriber services.

NOTE: This section describes service creation with MPLS uplinks using SDPs for interconnecting routers. This is applicable to only 7210 SAS-K 2F4T6C. It is not supported on 7210 SAS-E, 7210 SAS-D, and 7210 SAS-K 2F2T1C.

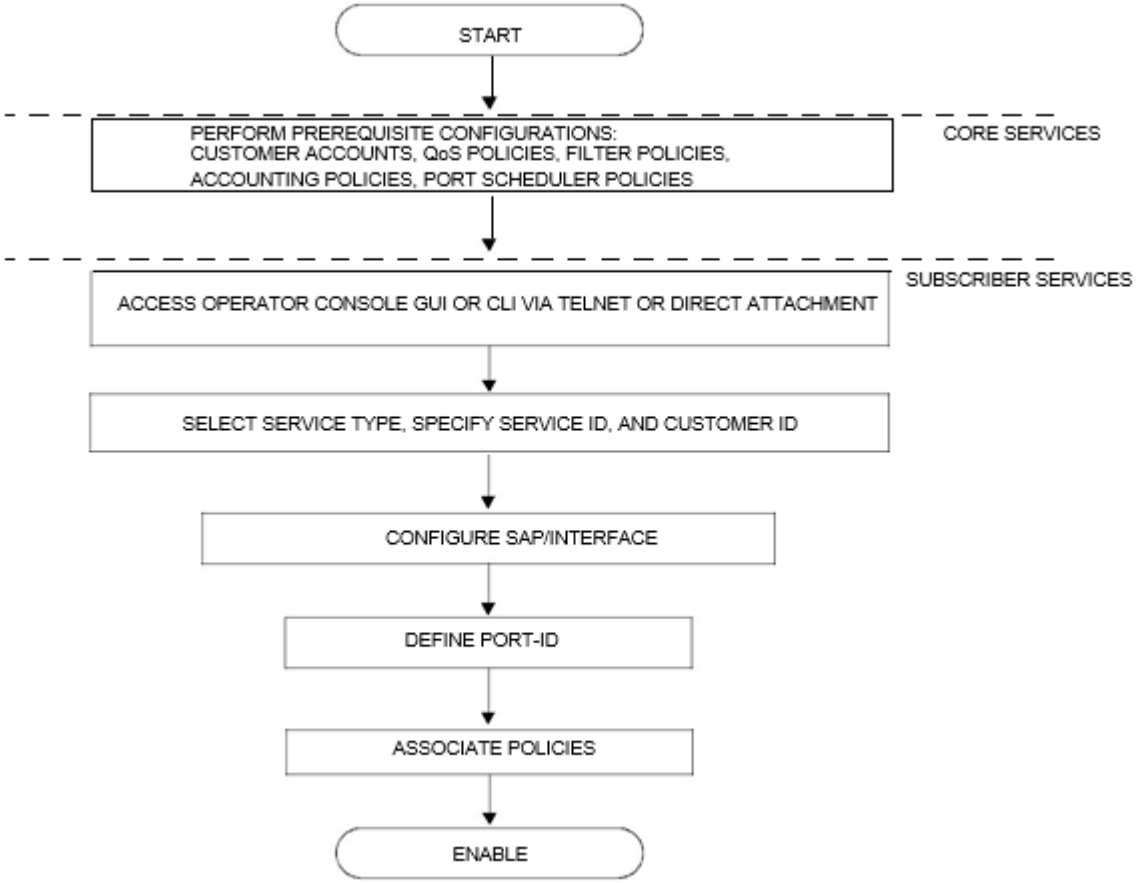

Figure 13: Service Creation and Implementation flow

Deploying and Provisioning Services - 7210 SAS-K 2F4T6C

The service model provides a logical and uniform way of constructing connectivity services. The basic steps for deploying and provisioning services can be broken down into three phases.

Phase 1: Core Network Construction

Before the services are provisioned, the following tasks should be completed:

- Build the IP or IP/MPLS core network.
- Configure routing protocols.
- Configure MPLS LSPs (if MPLS is used).
- Configure Access uplink SAPs (if only L2 uplinks are desired).

Phase 2: Service Administration

Perform preliminary policy configurations to control traffic flow, operator access, and to manage fault conditions and alarm messages, the following tasks should be completed:

- Configure group and user access privileges.
- Build templates for QoS, filter and/or accounting policies needed to support the core services.

Phase 3: Service Provisioning

- Provision customer account information.
- If necessary, build any customer-specific QoS, filter or accounting policies.
- Provision the customer services on the service edge routers by defining SAPs, binding policies to the SAPs.

Configuration Notes

This section describes service configuration caveats.

General

Service provisioning tasks can be logically separated into two main functional areas, core tasks and subscriber tasks and are typically performed prior to provisioning a subscriber service.

Core tasks include the following:

- Create customer accounts
- Create template QoS, filter, scheduler, and accounting policies
- Create SDPs (if using MPLS uplinks)

Subscriber services tasks include the following:

- Create Epipe and VPLS services.
- Create a VPRN service
- Bind SDPs (if using MPLS uplinks)
- Configure interfaces (where required) and SAPs
- Create exclusive QoS and filter policies

Configuring Global Service Entities with CLI

This section provides information to create subscriber (customer) accounts using the command line interface.

Topics include:

Service Model Entities

The Alcatel-Lucent service model uses logical entities to construct a service. The service model contains four main entities to configure a service.

- Subscribers on page 79

- Services:
  - → Ethernet Pipe (Epipe) Services on page 146
  - → VPLS on page 343
  - → IES on page 501
  - → VPLS SAP on page 356
  - → VPRN Service Configuration Commands on page 610

Basic Configuration

The most basic service configuration must have the following:

- A customer ID
- A service type
- A service ID
- A SAP identifying a port and encapsulation value
- An associated SDP for distributed services using MPLS uplinks.

The following example provides an Epipe service configuration displaying the SDP and Epipe service entities. SDP ID 1 was created with the far-end node 10.20.1.2. Epipe ID 101 was created for customer ID 1 which uses the SDP ID 1.

```
*A:Dut-B>config>service# info
----------------------------------------------
        sdp 1 mpls create
            far-end 10.20.1.2
            ldp
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        customer 1 create
            description "Default customer"
        exit
        epipe 101 customer 1 svc-sap-type any create
            spoke-sdp 1:1 create
                no shutdown
            exit
            no shutdown
        exit
----------------------------------------------
```

Common Configuration Tasks

This section provides a brief overview of the tasks performed to configure a customer account and an SDP. SDP configuration is not needed when 7210 SAS devices are configured to use only access-uplink ports.

Configuring Customers

The most basic customer account must have a customer ID. Optional parameters include:

- Description
- Contact name
- Telephone number

Customer Information

Use the following CLI syntax to create and input customer information:

CLI Syntax:

```
config>service# customer customer-id create
                  contact contact-information
                  description description-string
                  phone phone-number
```

The following displays a basic customer account configuration.

```
*A:K-SASK12>config>service>cust# info detail
----------------------------------------------
             description "default Customer"
             contact "Technical Support"
             phone "650 555-5100"
----------------------------------------------
*A:K-SASK12>config>service>cust#
```

Configuring an SDP

The most basic SDP must have the following:

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

SDP Configuration Tasks

This section provides a brief overview of the tasks performed to configure SDPs and provides the CLI commands.

Consider the following SDP characteristics:

- SDPs can be created as MPLS.
- Each distributed service must have an SDP defined for every remote router to provide VLL, VPLS, IES and VPRN services.
- A distributed service must be bound to an SDP. By default, no SDP is associated with a service. Once an SDP is created, services can be associated to that SDP.
- An SDP is not specific or exclusive to any one service or any type of service. An SDP can have more than one service bound to it.
- The SDP IP address must be a 7210 SAS-Series system IP address.
- In order to configure an MPLS SDP, LSPs must be configured first and then the LSP-to-SDP association must be explicitly created.
- In the SDP configuration, automatic ingress and egress labeling (targeted LDP) is enabled by default. Ingress and egress VC labels are signaled over a TLDP connection between two 7210 SAS-Series routers.

Note that if signaling is disabled for an SDP, then services using that SDP must configure ingress and egress vc-labels manually.

To configure a basic SDP, perform the following steps:

1. Specify an originating node.

2. Create an SDP ID.
3. Specify an encapsulation type.
4. Specify a far-end node.

Configuring an SDP

Use the following CLI syntax to create an SDP and select an encapsulation type. Only MPLS encapsulation is supported.

NOTE: When you specify the far-end ip address, you are creating the tunnel. In essence, you are creating the path from Point A to Point B. When you configure a distributed service, you must identify an SDP ID. Use the show service sdp command to display the qualifying SDPs.

When specifying MPLS SDP parameters, you must specify an LSP. If an LSP name is specified, then RSVP is used for dynamic signaling within the LSP.

LSPs are configured in the **config>router>mpls context**. See the 7210 SAS MPLS Guide for configuration and command information.

Use the following CLI syntax to create an MPLS SDP:

CLI Syntax:

```
config>service>sdp sdp-id [mpls] create
```

The following displays an LSP-signaled MPLS SDP configuration.

```
*A:K-SASK12>config>service>sdp$ info detail
----------------------------------------------
exit
        sdp 1 mpls create
            shutdown
            no description
            signaling tldp
            no ldp
            no bgp-tunnel
            no path-mtu
            no adv-mtu-override
            keep-alive
                shutdown
                hello-time 10
                hold-down-time 10
                max-drop-count 3
                timeout 5
                no message-length
            exit
----------------------------------------------
```

```
*A:K-SASK12>config>service>sdp$
```

Configuring a Mixed-LSP SDP

Use the following command to configure an SDP with mixed-LSP mode of operation:

config>service>sdp mpls>mixed-lsp-mode

The primary is backed up by the secondary. Two combinations are possible: primary of RSVP is backed up by LDP and primary of LDP is backed up by 3107 BGP.

The no form of this command disables the mixed-LSP mode of operation. The user first has to remove one of the LSP types from the SDP configuration or the command will fail.

The user can also configure how long the service manager must wait before it reverts the SDP to a higher priority LSP type, when it becomes available by using the following command:

config>service>sdp mpls>mixed-lsp-mode>revert-time revert-time

A special value of the timer dictates that the SDP must never revert to another higher priority LSP type unless the currently active LSP type is down:

config>service>sdp mpls>mixed-lsp-mode>revert-time infinite

The BGP LSP type is allowed. The bgp-tunnel command can be configured under the SDP with the lsp or ldp commands.

Ethernet Connectivity Fault Management (ETH-CFM)

Ethernet Connectivity Fault Management (ETH-CFM) is defined in two similar standards: IEEE 802.1ag and ITU-T Y.1731. They both specify protocols, procedures, and managed objects to support transport fault management, including discovery and verification of the path, detection and isolation of a connectivity fault for each Ethernet service instance. CFM functionalities are supported on 7210 SAS platforms.

The configuration is split into multiple areas. There is the base ETH-CFM configuration which defines the different Management constructs and administrative elements. This is performed in the ETH-CFM context. The individual management points are configure within the specific service contexts in which they are applied.

The 7210 SAS Services Guide provides the basic service applicable material to build the service specific management points, MEPs and MIPs. The different service types support a subset of the features from the complete ETH-CFM suite.

The support available varies across services and on different 7210 platforms.

The troubleshooting tools ETH-LBM/LBR, LTM/ LTR ETH-TST defined by the IEEE 802.1ag specification and the ITU-T Y.1731 recommendation are applicable to all MEPs (MIPs where appropriate). The advanced notification function AIS defined by the ITU-T Y.1731 is supported on Epipe services.

The advanced performance functions, 1DM, DMM/DMR and SLM/SLR are supported on all service MEPs.

For a description of the individual features and functions that are supported, see the OAM and Diagnostics Guide.

For more information about MEP and MIP support on different service endpoints (i.e. SAPs, SDP Bindings, etc.) varies based on the 7210 platform capabilities. Please see the 7210 platform table below to know more about the service endpoints supported and MEP/MIP supported..

Acronym	Call-out
1DM	One way Delay Measurement (Y.1731)
AIS	Alarm Indication Signal
CCM	Continuity check message
CFM	Connectivity fault management
DMM	Delay Measurement Message (Y.1731)
DMR	Delay Measurement Reply (Y.1731)

Acronym	Call-out
LBM	Loopback message
LBR	Loopback reply
LTM	Linktrace message
LTR	Linktrace reply
ME	Maintenance entity
MA	Maintenance association
MA-ID	Maintenance association identifier
MD	Maintenance domain
MEP	Maintenance association end point
MEP-ID	Maintenance association end point identifier
MHF	MIP half function
MIP	Maintenance domain intermediate point
OpCode	Operational Code
RDI	Remote Defect Indication
TST	Ethernet Test (Y.1731)
SLM	Synthetic Loss Message (Y.1731)
SLR	Synthetic Loss Reply (Y.1731)

ETH-CFM capabilities may be deployed in many different Ethernet service architectures. The Ethernet based SAPs and SDP bindings provide the endpoint on which the management points may be created. The basic functions can be used in different services, VPLS and Epipe. The following figure shows two possible example scenarios for ETH-CFM deployment in Ethernet access and aggregation networks.

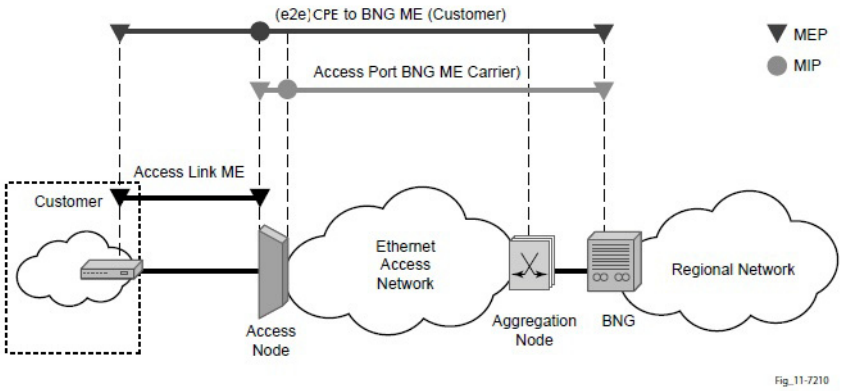

Figure 14: Ethernet OAM Model for Ethernet Access - Business

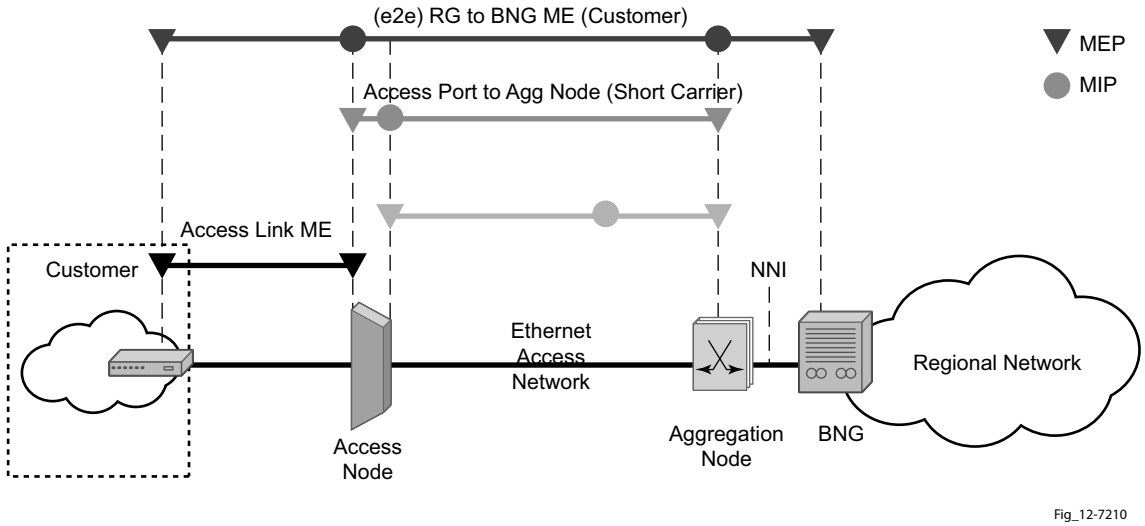

Figure 15: Ethernet OAM Model for Ethernet Access - Wholesale

The following functions are supported:

- On 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C, CFM can be enabled or disabled on a SAP basis.
- On 7210 SAS-K 2F4T6C, CFM can be enabled or disabled on a SAP and SDP Binding basis.

- The eight ETH-CFM levels are suggested to be broken up numerically between customer 7-5, service provider 4-3 and Operator 2-1.
  - → Level 0 typically is meant to monitor direct connections without any MIPs and should be reserved for port-based G8032 MEPs.
- Down MEP and UP MEP with an MEP-ID on a SAP/SDP binding for each MD level can be configured, modified, or deleted. Each MEP is uniquely identified by the MA-ID and MEP-ID tuple.
  - → MEP creation on a SAP is allowed only for Ethernet ports (with null, q-tags, qinq encapsulations).
  - → MEP support in different service and the endpoints configured in the service (SAPs, SDPs, IP interfaces, etc.) varies across service and 7210 platforms. The table below lists the support available for MEP on different 7210 platforms.
  - → On 7210 SAS-D, UP MEPs cannot be created by default on system boot-up. The user needs to explicitly allocate hardware resources for use with UP MEP feature, using the commands that appear under configure> system> resource- profile CLI context. Only after resources have been allocated by the user, UP MEPs can be created. Until resources are not allocated to UP MEP, the software fails all attempts to create an UP MEP.
  - → On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, no explicit resource allocation is required before configuring UP MEPs.
  - → 7210 SAS-E does not support UP MEPs.
- MIP creation on a SAP for each MD level can be enabled and disabled. MIP creation is automatic or manual when it is enabled. When MIP creation is disabled for an MD level, the existing MIP is removed.
  - → 7210 SAS-D and 7210 SAS-E platforms have the notion of ingress and egress MIPs. In other words, on these platforms, MIPs are not bi-directional by default. Ingress MIP responds to OAM messages that is received from the wire. Egress MIP responds to OAM messages that is being sent out to the wire. Ingress and Egress MIP support for SAP, SDP Bindings and services varies and is listed in the table below. For more information on MEP and MIP support, see MEP and MIP Support on page 94.
  - → On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, all MIPs are bi-directional. MIP responds to OAM messages that is received from the wire and also responds to OAM messages that is being sent out to the wire. MIP support for SAP, SDP Bindings and services varies and is listed in the table below. For more information on MEP and MIP support, see MEP and MIP Support on page 94.

Common Actionable Failures

It is important to note that AIS operates independently from the **low-priority-defect** setting. The **low-priority-defect** setting configuration parameter affects only the ETH-CFM fault propagation and alarming outside the scope of AIS. Any fault in the MEP state machine generates AIS when it is configured. Table 8 illustrates the ETH-CC defect condition groups, configured low-priority-defect setting, priority and defect as it applies to fault propagation.

Table 8: ETH-CC defect condition groups

Defect	Low Priority Defect	Description	Causes	Priority
DefNone	n/a	No faults in the association	Normal operations	n/a
DefRDICCM	allDef	Remote Defect Indication	Feedback mechanism to inform unidirectional faults exist. It provides the feedback loop to the node with the unidirectional failure conditions	1
DefMACStatus (default)	macRemErrXcon	MAC Layer	Remote MEP is indicating a remote port or interface not operational.	2
DefRemoteCCM	remErrXon	No communication from remote peer.	MEP is not receiving CCM from a configured peer. The timeout of CCM occurs at 3.5x the local CC interval. As per the specification, this value is not configurable.	3
DefErrorCCM	errXcon	Remote and local configures do not match required parameters.	Caused by different interval timer, domain level issues (lower value arriving at a MEP configured with a higher value), MEP receiving CCM with its MEPID	4
DefXconn	Xcon	Cross Connected Service	The service is receiving CCM packets from a different association. This could indicate that two services have merged or there is a configuration error on one of the SAP or bindings of the service, incorrect association identification.	5

MEP and MIP Support

The following is a general table that indicates the ETH-CFM support for the different services and endpoints. It is not meant to indicate the services supported or the requirements for those services on the individual platforms.

ETH-CFM Support for 7210 SAS-D

Table 9: ETH-CFM Support Matrix for 7210 SAS-D

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress and Egress	Not Supported
VPLS	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress MIP only	Not Supported (see Note1)
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
	SAP	Not Supported	Not Supported	Not Supported	Not Supported

NOTE 1 - On 7210 SAS-D, Primary VLAN support is available with UP MEP and Dot1q range SAPs only. For more information see the section on "Epipe Emulation using Dot1q VLAN range SAP in VPLS with G.8032".

ETH-CFM Support for 7210 SAS-E

Table 10: ETH-CFM Support Matrix for 7210 SAS-E

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP (Access and Access-uplink SAP)	Yes	Not Supported	Not Supported	Not Supported
VPLS	SAP (Access and Access-uplink SAP)	Yes	Not Supported	Ingress MIP only	Not Supported
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
	SAP	Not Supported	Not Supported	Not Supported	Not Supported

ETH-CFM Support for 7210 SAS-K 2F2T1C

Table 11: ETH-CFM Support Matrix for 7210 SAS-K 2F2T1C

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress and Egress	Supported (only for UP MEPs)
VPLS	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress and Egress	Not Supported

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
RVPLS	SAP	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
	SAP	Not Supported	Not Supported	Not Supported	Not Supported

ETH-CFM Support for 7210 SAS-K 2F4T6C

Table 12: ETH-CFM Support Matrix for 7210 SAS-K 2F4T6C

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
Epipe	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress and Egress	Not Supported
	Spoke-SDP	Yes	Yes	Ingress and Egress	Not Supported
VPLS	SAP (Access and Access-uplink SAP)	Yes	Yes	Ingress and Egress	Not Supported
	Spoke-SDP	Yes	Yes	Ingress and Egress	Not Supported
	Mesh-SDP	Yes	Yes	Not Supported	Not Supported

Service	Ethernet Connection Type	Down MEP	UP MEP	MIP	Primary VLAN
RVPLS	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
	IP interface (IES or VPRN)	Not Supported	Not Supported	Not Supported	Not Supported
IES	IES IPv4 interface	Not Supported	Not Supported	Not Supported	Not Supported
	SAP	Not Supported	Not Supported	Not Supported	Not Supported

ETM-CFM - MEP and MIP configuration notes:

- To achieve better scaling on the 7210 SAS-E, it is recommended that the MEPs are configured at particular levels. The recommended levels are 0, 1, 3 and 7.
- Ethernet-Rings are not configurable under all service types. Any service restrictions for MEP direction or MIP support will override the generic capability of the Ethernet-Ring MPs. For more information on Ethernet-Rings, refer to the 7210 SAS Interfaces Guide.

Configuring ETH-CFM Parameters

Configuring ETH-CFM requires commands at two different hierarchy levels of the CLI.

A sample of the global ETH-CFM configuration which defines the domains, associations, linkage o the service id or function, and the globally applicable CCM parameters including the interval and building of the remote MEPs database is shown below.

The following example displays a sample configuration.

```
*A:ALU-7_A>config>eth-cfm# info
----------------------------------------------
        domain 1 name "1" level 1
association 2 name 1345
                bridge-identifier 100
                exit
                ccm-interval 60
                remote-mepid 2
                remote-mepid 3
            exit
        exit
----------------------------------------------
*A:ALU-7_A>config>eth-cfm#
```

Defining the MEP and configuring service specific ETH-CFM parameters is performed within the service on the specific SAP or SDP binding. The example using the service VPLS 100 shows this configuration on the SAP.

```
#*A:ALU-7_A>config>service# info
----------------------------------------------
     vpls 100 customer 1 create
          description "VPLS service 100 - Used for MEP configuration example"
               sap 2/2/1:20 create
                    description "2/2/1:20"
                    eth-cfm
                         mep 1 domain 1 association 1 direction down
                              no shutdown
                         exit
                    exit
               exit
          exit
          no shutdown
          exit
          customer 1 create
               description "Default customer"
          exit
          exit
----------------------------------------------
*A:ALU-7_A>config>service#
```

All the examples shown above were based on IEEE 802.1ag. They are not capable of running Y.1731 functions. To build a Y.1731 context the domain format must be none.

The examples below show the global ETH-CFM configuration and the advanced Y.1731 functions that can be configured. The configuration will reject the configuration of Y.1731 functions within an IEEE 802.1ag context.

```
*A:7210-2# config>eth-cfm# info
----------------------------------------------
        domain 1 format none level 1
            association 1 format icc-based name "1234567890123"
                bridge-identifier 100
                exit
                ccm-interval 1
            exit
        exit


*A:7210-2# config>service# info
----------------------------------------------
        vpls 100 customer 1 create
            stp
                shutdown
            exit
            sap 2/2/1:40 create
                eth-cfm
                    mep 1 domain 1 association 1 direction up
                        ais-enable
                            priority 2
                            interval 60
                        exit
                        eth-test-enable
                            test-pattern all-ones crc-enable
                        exit
                        no shutdown
                    exit
                exit
            exit
            no shutdown
        exit
----------------------------------------------
```

Notes:

- To be able to transmit and receive AIS PDUs, a Y.1731 MEP must have **ais-enable** set.
- To be able to transmit and receive ETH-Test PDUs, a Y.1731 MEP must have **eth-test-enable** set.

Service Management Tasks

This section discusses the following service management tasks:

- Modifying Customer Accounts on page 100
- Deleting Customers on page 101

Modifying Customer Accounts

To access a specific customer account, you must specify the customer ID.
To display a list of customer IDs, use the show service customer command.
Enter the parameter (description, contact, phone) and then enter the new information.

CLI Syntax:
```
config>service# customer customer-id create
    [no] contact contact-information
    [no] description description-string
    [no] phone phone-number
```

Example:
```
config>service#  customer 27 create
config>service>customer$ description "Western Division"
config>service>customer# contact "John Dough"
config>service>customer# no phone "(650) 237-5102"
```

Deleting Customers

The no form of the customer command removes a customer ID and all associated information. All service references to the customer must be shut down and deleted before a customer account can be deleted.

CLI Syntax: `config>service# no customer customer-id`

Example:

```
config>service# epipe 5 customer 27 shutdown
config>service# epipe 9 customer 27 shutdown
config>service# no epipe 5
config>service# no epipe 9
config>service# no customer 27
```

Global Services Command Reference

Command Hierarchies

- Customer Commands on page 103
- Pseudo-wire (PW) Commands for 7210 SAS-K 2F4T6C on page 103
- SDP Commands for 7210 SAS-K 2F4T6C on page 105
- SAP Commands for 7210 SAS-E on page 106
- SAP Commands for 7210 SAS-D on page 106
- SAP Commands for 7210 SAS-K 2F2T1C on page 106
- SAP Commands for 7210 SAS-K 2F4T6C on page 107
- ETH-CFM Configuration Commands on page 108
- Show Commands on page 109

Customer Commands

```
config
    — service
        — [no] customer customer-id [create]
            — contact contact-information
            — no contact
            — description description-string
            — no description
            — [no] phone phone-number
```

Pseudo-wire (PW) Commands for 7210 SAS-K 2F4T6C

```
config
    — service
        — [no] pw-template policy-id [use-provisioned-sdp] [create]
            — accounting-policy acct-policy-id
            — no accounting-policy
            — [no] collect-stats
            — [no] control-word
            — [no] disable-learning
            — [no] disable-aging
            — [no] discard-unknown-source
            — [no] force-vlan-vc-forwarding
```

```
— igmp-snooping
        — [no] disable-router-alert-check
        — import policy-name
        — no import
        — last-member-query-interval 1/10 seconds
        — 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
        — version version
        — no version
— limit-mac-move {blockable | non-blockable}
— no limit-mac-move
— [no] mac-pinning
— max-nbr-mac-addr table-size
— no max-nbr-mac-addr
— split-horizon-group group-name
— no split-horizon-group
        — description description-string
        — no description
— vc-type {ether | vlan}
— vlan-vc-tag 0..4094
— no vlan-vc-tag
```

Pseudo-wire (PW) routing Commands for 7210 SAS-K 2F4T6C

```
config
    — service
        — pw-routing
            — boot-timer secs
            — no boot-timer
            — local-prefix local-prefix [create]
            — no local-prefix local-prefix
                — advertise-bgp route-distinguisher rd [community community]
                — no advertise-bgp route-distinguisher rd [community community]
            — path name [create]
            — no path name
                — hop hop-index ip-address
                — no hop hop-index
                — [no] shutdown
            — retry-count [10..10000]
            — no retry-count
            — retry-timer secs
            — no retry-timer
            — spe-address global-id:prefix
            — no spe-address
            — [no] static-route route-name
```

SDP Commands for 7210 SAS-K 2F4T6C

```
config
    — service
        — sdp sdp-id [mpls] [create]
        — no sdp sdp-id
            — accounting-policy acct-policy-id
            — no accounting-policy
            — collect-stats acct-policy-id
            — no collect-stats
            — [no] adv-mtu-override
            — [no] bgp-tunnel
            — [no] collect-stats
            — description description-string
            — no description
            — far-end ip-address no far-end
            — keep-alive
                — hello-time seconds
                — no hello-time
                — hold-down-time seconds
                — no hold-down-time
                — max-drop-count count
                — no max-drop-count
                — message-length octets
                — no message-length
                — [no] shutdown
                — timeout timeout
                — no timeout
            — [no] ldp
            — metric metric
            — no metric
            —  no mixed-lsp-mode
            — mixed-lsp-mode
                — no revert-time
                — revert-time {revert-time |infinite}
            — [no] lsp lsp-name
            — path-mtu octets
            — no path-mtu
            — [no] shutdown
            — signaling [off | tldp]
```

SAP Commands for 7210 SAS-E

```
config
    — service
        — epipe service-id [customer customer-id] [create] [svc-sap-type {null-star| dot1q| dot1q-pre-serve}] [customer-vid vlan-id]
        — no epipe service-id
            — sap sap-id [create]
            — no sap sap-id
        — ies service-id [customer customer-id] [create]
        — no ies service-id
            — sap sap-id [create]
            — no sap sap-id
        — vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type {null-star|dot1q|dot1q-preserve}] [customer-vid vlan-id]
        — no vpls service-id
            — sap sap-id [create] [eth-ring ring-index]
            — no sap sap-id
```

SAP Commands for 7210 SAS-D

```
config
    — service
        — epipe service-id [customer customer-id] [create] [svc-sap-type {null-star|dot1q-pre-serve|any|dot1q-range}] [customer-vid vlan-id]
        — no epipe service-id
            — sap sap-id [create]
            — no sap sap-id
        — ies service-id [customer <customer-id>] [create]
        — no ies service-id
            — sap sap-id [create]
            — no sap sap-id
        — vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type {null-star | any | dot1q-preserve}] [customer-vid vlan-id]
        — no vpls service-id
            — sap sap-id [create] [eth-ring ring-index]
            — no sap sap-id
```

SAP Commands for 7210 SAS-K 2F2T1C

```
config
    — service
        — epipe service-id [customer customer-id] [create] [svc-sap-type {any| dot1q-range}]
        — no epipe service-id
            — sap sap-id [create]
            — no sap sap-id
        — ies service-id [customer customer-id] [create] [vpn vpn-id]
        — no ies service-id
```

```
— vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {any|dot1q-range}]
  [r-vpls]
— no vpls service-id
      — sap sap-id [create]  [split-horizon-group group-name]
      — no sap sap-id
```

SAP Commands for 7210 SAS-K 2F4T6C

```
config
    — service
        — epipe service-id [customer customer-id] [create] [svc-sap-type {any| dot1q-range}]
        — no epipe service-id
            — sap sap-id [create]
            — no sap sap-id
        — ies service-id [customer customer-id] [create] [vpn vpn-id]
        — no ies service-id
        — vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {any|dot1q-range}]
          [r-vpls]
        — no vpls service-id
            — sap sap-id [create]  [split-horizon-group group-name]
            — no sap sap-id
        — vprn service-id [customer customer-id] [create]
        — no vprn service-id
```

ETH-CFM Configuration Commands

```
config
    — eth-cfm
        — domain md-index [format md-name-format] [name md-name] level level
        — domain md-index
        — no domain md-index
            — association ma-index [format ma-name-format] name ma-name
            — association ma-index
            — no association ma-index
                — [no] bridge-identifier bridge-id
                    — mhf-creation {default | none | explicit| static}
                    — no mhf-creation
                    — mip-ltr-priority priority
                    — no mip-ltr-priority
                    —  vlan vlan-id
                    — no vlan
                — ccm-interval
                — no ccm-interval {10ms| 100ms|1|10|60|600}
                — [no] remote-mepid mep-id
        — slm
            — [no] inactivity-timer timer
        — system
            — enable-dmm-version-interop
            — no enable-dmm-version-interop
```

Show Commands

```
show
    — service
        — customer [customer-id] [site customer-site-name]
        — id service-id
        — sdp [sdp-id | far-end ip-addr] [detail | keep-alive-history]
        — pw-template-using sdp-id pw-port [pw-port-id]
        — pw-template-using [consistent|inconsistent|na] egressifs
        — pw-template-using sdp-id keep-alive-history
        — pw-template-using far-end ip-address keep-alive-history
        — pw-template-using [sdp-id] [detail]
        — pw-template-using far-end ip-address [detail]
        — sdp-using [sdp-id[:vc-id] | far-end ip-address]
        — service-using [epipe] [ies] [vpls][mirror][customer customer-id]
    — eth-ring [status]
    — eth-ring ring-index hierarchy
    — eth-ring ring-index [path {a|b}]
    — eth-cfm
        — association [ma-index] [detail]
        — cfm-stack-table [port [port-id [vlan vlan-id]][level 0..7] [direction down]
        — mep mep-id domain md-index association ma-index [loopback] [linktrace]
        — mep mep-id domain md-index association ma-index remote-mepid mep-id | all-remote-mepids
        — mep mep-id domain md-index association ma-index eth-test-results [remote-peer mac-address]
        — mep mep-id domain md-index association ma-index one-way-delay-test [remote-peer mac-address]
        — mep mep-id domain md-index association ma-index two-way-delay-test [remote-peer mac-address]
        — pw-routing {local-prefix|static-route|paths|all}
        — pw-routing route-table [all-routes]
        — pw-routing route-table summary
    — pw-template
    — pw-template-using
```

Tools Perform Commands for 7210 SAS-K 2F4T6C

```
tools
    — perform
        — service
            — eval-pw-template policy-id [allow-service-impact]
            — id service-id
                — endpoint endpoint-name
                    — force-switchover sdp-id:vc-id
                    — no force-switchover
                    — force-switchover spoke-sdp-fec [1..4294967295]
                — eval-pw-template policy-id [allow-service-impact]
            — eval-expired-fec
                — eval-expired-fec spoke-sdp-fec-id
```

— **eval-expired-fec all**
— **spoke-sdp-fec-release** *global-id[:prefix[:ac-id]]*

Global Service Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>eth-cfm>mep

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Special Cases **Service Admin State —** While the service is shutdown, all customer packets are dropped and counted as discards for billing and debugging purposes.

description

Syntax **description** *description-string*
no description

Context config>service>customer

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Customer Management Commands

customer

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

Context config>service

Description This command creates a customer ID and customer context used to associate information with a particular customer. Services can later be associated with this customer at the service level.

Each *customer-id* must be unique. The *create* keyword must follow each new **customer** *customer-id* entry.

Enter an existing **customer** *customer-id* (without the *create* keyword) to edit the customer's parameters.

Default **customer 1** always exists on the system and cannot be deleted.

The **no** form of this command removes a *customer-id* and all associated information. Before removing a *customer-id*, all references to that customer in all services must be deleted or changed to a different customer ID.

Parameters *customer-id* — Specifies the ID number to be associated with the customer, expressed as an integer.

Values 1 — 2147483647

contact

Syntax **contact** *contact-information*
no contact *contact-information*

Context config>service>customer

Description This command allows you to configure contact information for a customer.

Include any customer-related contact information such as a technician's name or account contract name.

Default No contact information is associated with the *customer-id*.

The **no** form of this command removes the contact information from the customer ID.

Parameters *contact-information* — The customer contact information entered as an ASCII character string up to 80 characters in length. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Any printable, seven bit ASCII characters may be used within the string.

phone

Syntax [**no**] **phone** *string*

Context config>service>customer *customer-id*

Description This command adds telephone number information for a customer ID.

Default none

The **no** form of this command removes the phone number value from the customer ID.

Parameters *string* — The customer phone number entered as an ASCII string string up to 80 characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Any printable, seven bit ASCII characters may be used within the string.

Pseudo-wire Commands for 7210 SAS-K 2F4T6C

pw-routing

Syntax **pw-routing**

Context config>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command enables the context to configure dynamic multi-segment pseudowire (MS-PW) routing. Pseudowire routing must be configured on each node that will be a T-PE or an S-PE.

Default disabled

boot-timer

Syntax **boot-timer secs**
no boot-timer

Context config>service>pw-routing

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command configures a hold-off timer for MS-PW routing advertisements and signaling and is used at boot time.

The no form of this command removes a previously configured timer and restores it to its default.

Default 10

Parameters *timer-value*

The value of the boot timer in seconds.

Values 0 — 600

local-prefix

Syntax **local-prefix local-prefix [create]**
no local-prefixlocal-prefix

Context config>service>pw-routing

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command configures one or more node prefix values to be used for MS-PW routing. At least one prefix must be configured on each node that is an S-PE or a T-PE.

The no form of this command removes a previously configured prefix, and causes the corresponding route to be withdrawn if it has been advertised in BGP.

Default no local-prefix

Parameters *local-prefix*

Specifies a 32 bit prefix for the AII. One or more prefix values, up to a maximum of 16 may be assigned to the 7210 node. The global ID can contain the 2-octet or 4-octet value of the provider's Autonomous System Number (ASN). The presence of a global ID based on the provider's ASN ensures that the AII for spoke-SDPs configured on the node will be globally unique.

Values <global-id>:<ip-addr>|<raw-prefix>

ip-addr a.b.c.d

raw-prefix 1

4294967295

global-id 1

4294967295

advertise-bgp

Syntax **advertise-bgp route-distinguisher rd [community community]**
no advertise-bgp route-distinguisher rd

Context config>service>pw-routing

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command enables a given prefix to be advertised in MP-BGP for dynamic MS-PW routing.

The no form of this command will explicitly withdraw a route if it has been previously advertised.

Default no advertise-bgp

Parameters *rd*

Specifies a 32 bit prefix for the AII. One or more prefix values, up to a maximum of 16 may be assigned to the 7210 node. The global ID can contain the 2-octet or 4-octet value of the provider's Autonomous System Number (ASN). The presence of a global ID based on the provider's ASN ensures that the AII for spoke-SDPs configured on the node will be globally unique.

Values (6 bytes, other 2 Bytes of type will be automatically generated) asn:number1 (RD Type 0): 2bytes ASN and 4 bytes locally administered number ip-address:number2 (RD Type 1): 4bytes IPv4 and 2 bytes locally administered number;

community community

An optional BGP communities attribute associated with the advertisement. To delete a previously advertised community, advertise-bgp route-distinguisher must be run again with the same value for the RD but excluding the community attribute.

Values community {2-byte-as-number:comm-va1}

2-byte-asnumber 1— 65535

comm.-val 0 — 65535

path

Syntax **path name [create]**
no path name

Context config>service>pw-routing

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command configures an explicit path between this 7210 T-PE and a remote 7210 T-PE. For each path, one or more intermediate S-PE hops must be configured. A path can be used by multiple multi-segment pseudowires. Paths are used by a 7210 T-PE to populate the list of Explicit Route TLVs included in the signaling of a dynamic MS-PW.

A path may specify all or only some of the hops along the route to reach a T-PE.

The no form of the command removes a specified explicit path from the configuration.

Default no path

Parameters *path-name*

Specifies a locally-unique case-sensitive alphanumeric name label for the MS-PW path of up to 32 characters in length.

hop

Syntax **hop hop-index ip-address**
no hop hop-index

Context config>service>pw-routing>hop

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command configures each hop on an explicit path that can be used by one or more dynamic MSPWs. It specifies the IP addresses of the hops that the MS-PE should traverse. These IP addresses can correspond to the system IP address of each S-PE, or the IP address on which the T-LDP session to a given S-PE terminates.

The no form of this command deletes hop list entries for the path. All the MS-PWs currently using this path are unaffected. Additionally, all services actively using these MS-PWs are unaffected. The path must be shutdown first in order to delete the hop from the hop list. The `no hop hop-index' command will not result in any action, except for a warning message on the console indicating that the path is administratively up.

Default no hop

Parameters hop-index

Specifies a locally significant numeric identifier for the hop. The hop index is used to order the hops specified. The LSP always traverses from the lowest hop index to the highest. The hop index does not need to be sequential.

Values 1 — 16

ip-address

Specifies the system IP address or terminating IP address for the T-LDP session to the S-PE corresponding to this hop. For a given IP address on a hop, the system will choose the appropriate SDP to use.

retry-count

Syntax **retry-count [10..10000]**
no retry-count

Context config>service>pw-routing

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

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 retry-count

Specifies the maximum number of retries before putting the spoke-sdp into the shutdown state.

Values 10 — 10000

retry-timer

Syntax **retry-timer secs**
no retry-timer

Context config>service>pw-routing

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

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 "AII unreachable".

The no form of this command reverts the timer to its default value.

Default 30

Parameters retry-count

The initial retry-timer value in seconds.

Values 10 – 480

spe-address

Syntax **spe-address global-id:prefix**
no spe-address

Context config>service>pw-routing

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

NOTE: 7210 SAS-K 2F4T6C cannot be configured to be a S-PE node. The command allows the 7210 SAS-K 2F4T6C configured as T-PE node to determine the S-PE node to use for Multi-segment PWs.

This command configures a single S-PE Address for the node to be used for dynamic MS-PWs. This value is used for the PW switching point TLV used in LDP signaling, and is the value used by PW status signaling to indicate the PE that originates a PW status message. Configuration of this parameter is mandatory to enable dynamic MS-PW support on a node.

If the S-PE Address is not configured, spoke-sdps that use dynamic MS-PWs and pw-routing localprefixes cannot be configured on a T-PE. Furthermore, and 7210 node will send a label release for any label mappings received for FEC129 AII type 2.

The S-PE Address cannot be changed unless the dynamic ms-pw configuration is removed.

Furthermore, changing the S-PE Address will also result in all dynamic MS-PWs for which this node is an S-PE being released. It is recommended that the S-PE Address should be configured for the life of an MS-PW configuration after reboot of the 7210.

The no form of this command removes the configured S-PE Address.

Default no spe-address

Parameters *global-id*

Specifies a 4-octet value that is unique to the service provider. For example, the global ID can contain the 2-octet or 4-octet value of the provider's Autonomous System Number (ASN).

Syntax: <*global-id:prefix*>: <*global-id*>:{<prefix>|<ipaddress>}

global-id 1 — 4294967295

prefix 1 — 4294967295

ipaddress a.b.c.d

static-route

Syntax **[no] static-route *route-name***

Context config>service>pw-routing

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command configures a static route to a next hop S-PE or T-PE. Static routes may be configured on either S-PEs or T-PEs.

A default static route is entered as follows:

static-route 0:0:next_hop_ip_addresss

or

static-route 0:0.0.0.0:next_hop_ip_address

The no form of this command removes a previously configured static route.

Default no static-route

Parameters route-name

Specifies the static pseudowire route.

Values route-name <*global-id*>:<*prefix*>:<next-hop-ip_addr>

<global-id>:0

4294967295

prefix a.b.c.d | 0— 4294967295

ip_addr a.b.c.d

pw-template

Syntax **[no] pw-template policy-id [use-provisioned-sdp] [create]**

Context config>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures an SDP template.

Parameters *use-provisioned-sdp*

Specifies whether to use an already provisioned SDP. When specified, the tunnel manager will be consulted for an existing active SDP. Otherwise, the default SDP template will be used to use for instantiation of the SDP.

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.

control-word

Syntax **[no] control-word**

Context config>service>pw-template

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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.

Default no control-word

SDP Commands

sdp

Syntax **sdp sdp-id[mpls] [create]**
no sdp sdp-id

Context config>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command creates or edits a Service Distribution Point (SDP). SDPs must be explicitly configured.

An SDP is a logical mechanism that ties a far-end 7210 SAS to a service without having to specifically define far end SAPs. Each SDP represents a method to reach a 7210 SAS router.

7210 SAS supports only Multi-Protocol Label Switching (MPLS) encapsulation. A 7210 SAS supports both signaled and non-signaled Label Switched Paths (LSPs) through the network. Non-signaled paths are defined at each hop through the network. Signaled paths are communicated by protocol from end to end using Resource Reservation Protocol (RSVP). Paths may be manually defined or a constraint-based routing protocol (such as OSPF-TE or CSPF) can be used to determine the best path with specific constraints. An LDP LSP can also be used for an SDP when the encapsulation is MPLS. The use of an LDP LSP type or an RSVP/Static LSP type are mutually exclusive except when the mixed-lsp option is enabled on the SDP.

SDPs are created and then bound to services. Many services may be bound to a single SDP. The operational and administrative state of the SDP controls the state of the SDP binding to the service.

If *sdp-id* does not exist, a new SDP is created. When creating an SDP, the **mpls** keyword must be specified. SDPs are created in the admin down state (**shutdown**) and the **no shutdown** command must be executed once all relevant parameters are defined and before the SDP can be used.

If *sdp-id* exists, the current CLI context is changed to that SDP for editing and modification. For editing an existing SDP, the **mpls** keyword is specified. If a keyword is specified for an existing *sdp-id*, an error is generated and the context of the CLI will not be changed to the specified *sdp-id*.

The **no** form of this command deletes the specified SDP. Before an SDP can be deleted, it must be administratively down (shutdown) and not bound to any services. If the specified SDP is bound to a service, the **no sdp** command will fail generating an error message specifying the first bound service found during the deletion process. If the specified *sdp-id* does not exist an error will be generated.

Default none

Parameters *sdp-id*

The SDP identifier.

Values 1 - 17407

accounting-policy

Syntax **accounting-policy acct-policy-id**
no accounting-policy

Context config>service>sdp
config>service>pw-template

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command creates the accounting policy context applied to an SDP. An accounting policy must be defined before it can be associated with a SDP. If the policy-id does not exist, an error message is generated.

A maximum of one accounting policy can be associated with a SDP at one time. Accounting policies are configured in the config>log context.

The no form of this command removes the accounting policy association from the SDP, and the accounting policy reverts to the default.

Default Default accounting policy.

Parameters *acct-policy-id*

Enter the accounting policy-id as configured in the config>log>accounting-policy context.

Values 1 - 99

collect-stats

Syntax **[no] collect-stats**

Context config>service>sdp
config>service>pw-template

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables accounting and statistical data collection for either the SDP. 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

disable-learning

Syntax **[no] disable-learning**

Context config>service>pw-template

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables learning of new MAC addresses.

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-source

Syntax **[no] discard-unknown-source**

Context config>service>pw-template

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

When this command is enabled, packets received with an unknown source MAC address will be dropped only if the maximum number of MAC addresses have been reached. 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.

Default no discard-unknown

egress

Syntax **egress**

Context config>service>pw-template

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure spoke SDP binding egress filter parameters.

ingress

Syntax **ingress**

Context config>service>pw-template

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure spoke SDP binding ingress filter parameters.

force-vlan-vc-forwarding

Syntax **[no] force-vlan-vc-forwarding**

Context config>service>pw-template

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command forces vc-vlan-type forwarding in the data path for spoke and mesh SDPs that have either vc-type. This command is not allowed on vlan-vc-type SDPs.

The no version of this command sets default behavior.

Parameters This feature is disabled per default.

split-horizon-group

Syntax **split-horizon-group** *group-name* [**create**]

Context config>service>vpls
config>service>pw-template

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command is used to create a new split horizon group for the VPLS instance. Traffic arriving on a SAP or spoke SDP within this split horizon group will not be copied to other SAPs or spoke SDPs in the same split horizon group.

A split horizon group must be created before SAPs and spoke SDPs can be assigned to the group. The split horizon group is defined within the context of a single VPLS instance. The same group-name can be re-used in different VPLS instances.

he **no** form of the command removes the group name from the configuration.

Parameters *group-name* — Specifies the name of the split horizon group to which the SAP or Spoke-SDP belongs.

create — Mandatory keyword to create a split-horizon group.

vc-label

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

Context config>service>pw-template>ingress

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the ingress VC label.

Parameters *vc-label*

A VC ingress value that indicates a specific connection.

Values 2048 - 18431

limit-mac-move

Syntax **limit-mac-move [blockable | non-blockable]**
no limit-mac-move

Context config>service>pw-template

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command indicates whether or not the mac-move agent will limit the MAC re-learn (move) rate.

Default blockable

Parameters *blockable* — The agent will monitor the MAC re-learn rate, and it will block it when the re-learn rate is exceeded.

non-blockable — When specified, a SAP will not be blocked, and another blockable SAP will be blocked instead.

vc-type

Syntax **vc-type {ether | vlan}**

Context config>service>pw-template

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command overrides the default VC type signaled for the binding to the far end 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 vctype command can still be used to define the dot1q value expected by the far-end provider equipment.

A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

VC types are derived according to IETF draft-martini-l2circuit-trans-mpls.

- The VC type value for Ethernet is 0x0005.
- The VC type value for an Ethernet VLAN is 0x0004.

Parameters *ether* — Defines the VC type as Ethernet. The ethernet and vlan keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. Defining Ethernet is the same as executing no vc-type and restores the default VC type for the spoke SDP binding. (hex 5)

vlan — Defines the VC type as VLAN. The ethernet and vlan keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings.

vlan-vc-tag

Syntax **vlan-vc-tag 0..4094**
no vlan-vc-tag [0..4094]

Context config>service>pw-template

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The no form of this command disables the command

Default no vlan-vc-tag

Parameters 0..4094 — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

adv-mtu-override

Syntax **[no] adv-mtu-override**

Context config>service>sdp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command overrides the advertised VC-type MTU of all spoke-sdps of Layer 2 services using this SDP-ID. When enabled, the router signals a VC MTU equal to the service MTU, which includes the Layer 2 header. It also allows this router to accept an MTU advertized by the far-end PE which value matches either its advertised MTU or its advertised MTU minus the Layer 2 headers.

By default, the router advertizes a VC-MTU equal to the Layer 2 service MTU minus the Layer 2 header and always matches its advertized MTU to that signaled by the far-end PE rotuer, otherwise the spoke-sdp goes operationally down.

When this command is enabled on the SDP, it has no effect on a spoke-sdp of an IES/VPRN spoke interface using this SDP-ID. The router continues to signal a VC MTU equal to the net IP interface MTU, which is min (ip-mtu, sdp operational path mtu - Layer 2 headers). The router also continues to make sure that the advertized MTU values of both PE routers match or the spoke-sdp goes operationally down.

The no form of the command disables the VC-type MTU override and returns to the default behavior.

Default no adv-mtu-override

bgp-tunnel

Syntax **[no] bgp-tunnel**

Context config>service>sdp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command allows the use of BGP route tunnels available in the tunnel table to reach SDP far-end nodes. Use of BGP route tunnels are only available with MPLS-SDP. Only one of the transport methods is allowed per SDP - LDP, RSVP-LSP or BGP-Tunnel (BGP-Tunnel is not supported on multi-mode LSP)

NOTE: The 7210 SAS provides an option to install labels for only those BGP 3107 labelled routes which are in use by services. For more details about this option, see the 7210 SAS Routing Protocols User Guide.

The no form of the command disables resolving BGP route tunnel LSP for SDP far-end.

Default no bgp-tunnel (BGP tunnel route to SDP far-end is disabled)

far-end

Syntax **far-end ip-address node-id node-id**
no far-end

Context config>service>sdp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the system IP address of the far-end destinationrouter for the Service Distribution Point (SDP) that is the termination point for a service.

The far-end IP address must be explicitly configured. The destination IP address must be asystem IP address.

If the SDP uses MPLS encapsulation, the far-end ip-address is used to check LSP names when added to the SDP. If the “to IP address” defined within the LSP configuration does not exactly match the SDP far-end ip-address, the LSP will not be added to the SDP and an error will be generated.

If the SDP uses MPLS encapsulation, the far-end ip-address is used to check LSP names when added to the SDP. If the “to IP address” defined within the LSP configuration does not exactly match the SDP far-end ip-address, the LSP will not be added to the SDP and an error will be generated.

Alternatively, and SDP that uses MPLS can have an MPLS-TP node with an MPLS-TP node-id and (optioanlly) global-id. In this case, the SDP must use an MPLS-TP LSP and the SDP signaling parameter must be set to off.

An SDP cannot be administratively enabled until a far-end ip-address or MPLS-TP node-id is defined. The SDP is operational when it is administratively enabled (no shutdown) and the far-end ipaddress is contained in the IGP routing table as a host route. OSPF ABRs should not summarize host routes between areas. This can cause SDPs to become operationally down. Static host routes (direct and indirect) can be defined in the local dev ice to alleviate this issue.

The no form of this command removes the currently configured destination IP address for the SDP.

The ip-address parameter is not specified and will generate an error if used in the no far-end command. The SDP must be administratively disabled using the config service sdp shutdown command before the no far-end command can be executed. Removing the far end IP address will cause all lsp-name associations with the SDP to be removed.

Default none

Parameters *ip-address* — The system address of the far-end 7210 devices for the SDP in dotted decimal notation.

node-id node-id — The MPLS-TP Node ID of the far-end system for the SDP, either in dotted decimal notaion (a.b.c.d) or an

unsigned 32-bit integer (1 – 4294967295).

metric

Syntax **metric metric**
no metric

Context config>service>sdp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the metric to be used within the tunnel table manager for decision making purposes. When multiple SDPs going to the same destination exist, this value is used as a tie-breaker by tunnel table manager users such as MP-BGP to select the route with the lower value.

Parameters metric — Specifies the SDP metric.

Values 0 — 65535

mixed-lsp-mode

Syntax **[no] mixed-lsp-mode**

Context config>service>sdp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the use by an SDP of the mixed-LSP mode of operation. This command indicates to the service manager that it must allow a primary LSP type and a backup LSP type in the same SDP configuration. For example, the lsp and ldp commands are allowed concurrently in the SDP configuration. The user can configure one or two types of LSPs under the same SDP. Without this command, these commands are mutually exclusive.

The user can configure an RSVP LSP as a primary LSP type with an LDP LSP as a backup type. The user can also configure a BGP RFC 3107 BGP LSP as a backup LSP type.

If the user configures an LDP LSP as a primary LSP type, then the backup LSP type must be an RFC 3107 BGP labeled route.

At any given time, the service manager programs only one type of LSP in the linecard that will activate it to forward service packets according to the following priority order:

1. RSVP LSP type. One RSVP LSP can be configured per SDP.

In the case of the RSVP/LDP SDP, the service manager will program the NHLFE(s) for the active LSP type preferring the RSVP LSP type over the LDP LSP type. If no RSVP LSP is configured or all configured RSVP LSPs go down, the service manager will re-program the linecard with the LDP LSP if available. If not, the SDP goes operationally down.

When a higher priority type LSP becomes available, the service manager reverts back to this LSP at the expiry of the sdp-revert-time timer or the failure of the currently active LSP, whichever comes first. The service manager then re-programs the linecard accordingly. If the infinite value is configured, then the SDP reverts to the highest priority type LSP only if the currently active LSP failed.

Note however, that LDP uses a tunnel down damp timer which is set to three seconds by default. When the LDP LSP fails, the SDP will revert to the RSVP LSP type after the expiry of this timer. For an immediate switchover this timer must be set to zero. Use the configure>router>ldp>tunneldown-damp-time command.

Note: For more information on, "***configure>router>ldp>tunnel-down-damp-time***" command, see 7210 SAS OS MPLS Guide.

If the user changes the value of the sdp-revert-time timer, it will take effect only at the next use of the timer. Any timer which is outstanding at the time of the change will be restarted with the new value. In the case of the LDP/BGP SDP, the service manager will prefer the LDP LSP type over the BGP LSP type. The service manager will re-program the linecard with the BGP LSP if available otherwise it brings down the SDP operationally.

Also Note the following difference in behavior of the LDP/BGP SDP compared to that of an RSVP/LDP SDP. For a given /32 prefix, only a single route will exist in the routing table: the IGP route or the BGP route. Thus, either the LDP FEC or the BGP label route is active at any given time. The impact of this is that the tunnel table needs to be re-programmed each time a route is deactivated and the other is activated. Furthermore, the SDP revert-time cannot be used since there is no situation where both LSP types are active for the same /32 prefix.

The no form of this command disables the mixed-LSP mode of operation. The user first has to remove one of the LSP types from the SDP configuration or the command will fail.

Default no mixed-lsp-mode

revert-time

Syntax **revert-time {revert-time | infinite}**
no revert-time

Context config>service>sdp>mixed-lsp-mode

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the delay period the SDP must wait before it reverts to a higher priority LSP type when one becomes available.

The no form of the command resets the timer to the default value of 0. This means the SDP reverts immediately to a higher priority LSP type when one becomes available.

Default 0

Parameters *revert-time* — Specifies the delay period, in seconds, that the SDP must wait before it reverts to a higher priority LSP type when one becomes available. A value of zero means the SDP reverts immediately to a higher priority LSP type when one becomes available.

Values 0 — 600

infinite — This keyword forces the SDP to never revert to another higher priority LSP type unless the currently active LSP type is down.

ldp

Syntax **[no] ldp**

Context config>service>sdp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables LDP-signaled LSP's on MPLS-encapsulated SDPs.

In MPLS SDP configurations either one LSP can be specified or LDP can be enabled. The SDP ldp and lsp commands are mutually exclusive. If an LSP is specified on an MPLS SDP, then LDP cannot be enabled on the SDP. To enable LDP on the SDP when an LSP is already specified, the LSP must be removed from the configuration using the no lsp lsp-name command.

Alternatively, if LDP is already enabled on an MPLS SDP, then an LSP cannot be specified on the SDP. To specify an LSP on the SDP, the LDP must be disabled. The LSP must have already been created in the config>router>mpls context with a valid far-end IP address. The above rules are relaxed when the mixed-lsp option is enabled on the SDP.

Default no ldp (disabled)

lsp

Syntax **lsp lsp-name**
no lsp lsp-name

Context config>service>sdp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command creates associations between one label switched paths (LSPs) and an Multi-Protocol

Label Switching (MPLS) Service Distribution Point (SDP). This command is implemented only on MPLS-type encapsulated SDPs.

In MPLS SDP configurations either one LSP can be specified.

The LSP must have already been created in the config>router>mpls context. with a valid far-end IP address. RSVP must be enabled.

If no LSP is associated with an MPLS SDP, the SDP cannot enter the operationally up state. The SDP can be administratively enabled (no shutdown) with no LSP associations. The lsp-name may be shutdown, causing the association with the SDP to be operationally down (the LSP will not be used by the SDP).

The no form of this command deletes one LSP associations from an SDP. If the lsp-name does not exist as an association or as a configured LSP, no error is returned. An lsp-name must be removed from all SDP associations before the lsp-name can be deleted from the system. The SDP must be administratively disabled (shutdown) before the last lsp-name association with the SDP is deleted.

Default none

Parameters *lsp-name* — The name of the LSP to associate with the SDP. An LSP name is case sensitive and is limited to 32 ASCII 7-bit printable characters with no spaces. If an exact match of lsp-name does not already exist as a defined LSP, an error message is generated. If the lsp-name does exist and the LSP to IP address matches the SDP **far-end** IP address, the association is created.

signaling

Syntax **signaling {off | tldp}**

Context config>service>sdp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the signaling protocol used to obtain the ingress and egress pseudowire labels in frames transmitted and received on the SDP. When signaling is off then labels are manually configured when the SDP is bound to a service. The signalling value can only be changed while the administrative status of the SDP is down.

The no form of this command is not applicable. To modify the signaling configuration, the SDP must be administratively shut down and then the signaling parameter can be modified and re-enabled.

Default tldp

Parameters *off* — Ingress and egress signal auto-labeling is not enabled. If this parameter is selected, then each service using the specified SDP must manually configure VPN labels. This configuration is independent of the SDP's transport type, MPLS (RSVP or LDP).

tldp — Ingress and egress pseudowire signaling using T-LDP is enabled.

path-mtu

Syntax **path-mtu bytes**
no path-mtu

Context config>service>sdp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the Maximum Transmission Unit (MTU) in bytes that the Service Distribution Point (SDP) can transmit to the far-end device router without packet dropping or IP fragmentation overriding the SDP-type default path-mtu.

The default SDP-type path-mtu can be overridden on a per SDP basis. Dynamic maintenance protocols on the SDP like RSVP may override this setting.

If the physical mtu on an egress interface indicates the next hop on an SDP path cannot support the current path-mtu, the operational path-mtu on that SDP will be modified to a value that can be transmitted without fragmentation.

The no form of this command removes any path-mtu defined on the SDP and the SDP will use the system default for the SDP type.

Default The default **path-mtu** defined on the system for the type of SDP is used.

SDP Keepalive Commands

keep-alive

Syntax **keepalive**

Context config>service>sdp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Context for configuring SDP connectivity monitoring keepalive messages for the SDP ID. SDP-ID keepalive messages use SDP Echo Request and Reply messages to monitor SDP connectivity. The operating state of the SDP is affected by the keepalive state on the SDP-ID. SDP Echo Request messages are only sent when the SDP-ID is completely configured and administratively up. If the SDP-ID is administratively down, keepalives for that SDP-ID are disabled.

SDP Echo Requests (when sent for keepalive messages) are always sent with the originator-sdp-id.

All SDP-ID keepalive SDP Echo Replies are sent using generic IP OAM encapsulation. When a keepalive response is received that indicates an error condition, the SDP ID will immediately be brought operationally down. Once a response is received that indicates the error has cleared and the hold-down-time interval has expired, the SDP ID will be eligible to be put into the operationally up state. If no other condition prevents the operational change, the SDP ID will enter the operational state.

A set of event counters track the number of keepalive requests sent, the size of the message sent, nonerror replies received and error replies received. A keepalive state value is kept indicating the last

response event. A keepalive state timestamp value is kept indicating the time of the last event. With each keepalive event change, a log message is generated indicating the event type and the timestamp value.

The table below describes keepalive interpretation of SDP echo reply response conditions and the

effect on the SDP ID operational status.

hello-time

Syntax **hello-time seconds**
no hello-time

Context config>service>sdp>keep-alive

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Configures the time period between SDP keepalive messages on the SDP-ID for the SDP connectivity

monitoring messages.

The no form of this command reverts the hello-time seconds value to the default setting.

Default *hello-time 10* — 10 seconds between keepalive messages

seconds — The time period in seconds between SDP keepalive messages, expressed as a decimal integer.

Values 1 — 3600

hold-down-time

Syntax **hold-down-time seconds**
no hold-down-time

Context config>service>sdp>keep-alive

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Configures the minimum time period the SDP will remain in the operationally down state in response to SDP keepalive monitoring.

This parameter can be used to prevent the SDP operational state from "flapping" by rapidly transitioning between the operationally up and operationally down states based on keepalive messages.

When an SDP keepalive response is received that indicates an error condition or the max-drop-count keepalive messages receive no reply, the sdp-id will immediately be brought operationally down. If a keepalive response is received that indicates the error has cleared, the sdp-id will be eligible to be put into the operationally up state only after the hold-down-time interval has expired.

The no form of this command reverts the hold-down-time seconds value to the default setting.

Default *hold-down-time 10* — The SDP is operationally down for 10 seconds after an SDP keepalive error.

Parameters *seconds* — The time in seconds, expressed as a decimal integer, the sdp-id will remain in the operationally down state before it is eligible to enter the operationally up state. A value of 0 indicates that no hold-down-time will be enforced for sdp-id.

Values 0 — 3600

max-drop-count

Syntax **max-drop-count** *count*
no max-drop-count

Context config>service>sdp>keep-alive

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the number of consecutive SDP keepalive failed request attempts or remote replies that can be missed after which the SDP is operationally downed. If the max-dropcount consecutive keepalive request messages cannot be sent or no replies are received, the SDP-ID will be brought operationally down by the keepalive SDP monitoring.

The no form of this command reverts the max-drop-count count value to the default settings.

Default max-drop-count 3

Parameters *count* — The number of consecutive SDP keepalive requests that are failed to be sent or replies missed, expressed as a decimal integer.

Values 1 — 5

message-length

Syntax **message-length** *octets*
no message-length

Context config>service>sdp>keep-alive

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the SDP monitoring keepalive request message length transmitted. The no form of this command reverts the message-length octets value to the default setting.

Default 0 — The message length should be equal to the SDP's operating path MTU as configured in the **pathmtu** command. If the default size is overridden, the actual size used will be the smaller of the operational SDP-ID Path MTU and the size specified.

octets — The size of the keepalive request messages in octets, expressed as a decimal integer. The size keyword overrides the default keepalive message size.

Values 40 — 9198

timeout

Syntax **timeout** *timeout*
no timeout

Context config>service>sdp>keep-alive

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the time interval that the SDP waits before tearing down the session.

Default 5

Parameters *timeout* — The timeout time, in seconds.

Values 1 — 10

ETH-CFM Configuration Commands

eth-cfm

Syntax **eth-cfm**

Context config

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables the context to configure ETH CFM parameters.

domain

Syntax **domain** *md-index* **[format** *md-name-format***] [name** *md-name***] level** *level*
domain *md-index*
no domain *md-index*

Context config>eth-cfm

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures Connectivity Fault Management domain parameters.

The no form of the command removes the MD index parameters from the configuration.

Parameters *md-index* — Specifies the Maintenance Domain (MD) index value.

Values 1 — 4294967295

format {dns | mac | none | string} — Specifies a value that represents the type (format).

Values dns: Specifies the DNS name format.

mac: X:X:X:X:X:X-u

X: [0..FF]h

u: [0..65535]d

none: Specifies a Y.1731 domain format and the only format allowed to execute Y.1731 specific functions.

string Specifies an ASCII string.

Default string

name *md-name* — Specifies a generic Maintenance Domain (MD) name

Values 1 — 43 characters

level *level* — Specifies the integer identifying the maintenance domain level (MD Level). Higher numbers correspond to higher maintenance domains, those with the greatest physical reach, with the highest values for customers' CFM packets. Lower numbers correspond to lower maintenance domains, those with more limited physical reach, with the lowest values for single bridges or physical links.

Values 0 — 7

association

Syntax **association** *ma-index* **[format** *ma-name-format***] name** *ma-name*
association *ma-index*
no association *ma-index*

Context config>eth-cfm>domain

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the Maintenance Association (MA) for the domain.

Parameters *ma-index* — Specifies the MA index value.

Values 1 — 4294967295

format {icc-based | *integer* | *string* | *vid* | *vpn-id}* — Specifies a value that represents the type (format).

Values icc-based: Only applicable to a Y.1731 context where the domain format is configured as none. Allows for exactly a 13 character name.

integer: 0 — 65535 (integer value 0 means the MA is not attached to a VID.)

string: raw ascii

vid: 0 — 4095

vpn-id: RFC-2685, Virtual Private Networks Identifier

xxx:xxxx, where x is a value between 00 and FF.

for example 00164D:AABBCCDD

Default integer

name *ma-name* — Specifies the part of the maintenance association identifier which is unique within the maintenance domain name.

Values 1 — 45 characters

bridge-identifier

Syntax **[no] bridge-identifier bridge-id**

Context config>eth-cfm>domain>association

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the cross-reference required to link the CFM function with the service context. The link is created when the bridge-id, service-id, and vlan-id (for a primary VLAN) match.

Under the association context, this command is used to specify various MEP and MIP creation parameters. The VLAN parameter is not tied to the bridge-identifier statement, but rather is an object under the bridge-identifier context.

The no form of this command is only available under the association context. Negating the line will remove the bridge-identifier and the link between the ETH-CFM configuration and the matching service-id.

Parameters bridge-id — Specifies the bridge ID for the domain association.Note that there is no verification that a service has been created with a matching service ID.

Values 1 — 2147483647

mhf-creation

Syntax **mhf-creation {default | none | explicit| static}**
no mhf-creation

Context config>eth-cfm>domain>association>bridge-identifier

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command determines whether to allow MIP creation for the MA. Use of the none, default and explicit parameters are only allowed for MHFs (MIPs) not associated with a configured Primary VLAN.

NOTE: On 7210 platforms, we specifically call out support for Ingress MIP and Egress MIP. Ingress MIP responds to OAM messages that is received from the wire. Egress MIP responds to OAM messages that is being sent out to the wire.

Please refer to the tables, for MEP and MIP support available for different services on different platforms.

Default none

Parameters *default* — Specifies that MHFs can be created for this VID only on bridge ports through which this VID can pass without the requirement for a MEP at some lower MA level.

none — Specifies that no MHFs can be created for this VID.

explicit — Specifies that MHFs can be created for this VID only on bridge ports through which this VID can pass, and only if a MEP is created at some lower MA level. There must be at least one lower level MEP provisioned on the same SAP.

mip-ltr-priority

Syntax **mip-ltr-priority priority**
no mip-ltr-priority

Context config>eth-cfm>domain>association>bridge-identifier

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command allows the operator to set the priority of the Linktrace Response Message (ETH-LTR) from a MIP for this association. If this command is not specified a LTR priority of 7 will be used.

Default no mip-ltr-priority

Parameters *priority* — Specifies the priority of the Linktrace Response Message (ETH-LTR) from a MIP for this association.

Values 0 — 7

vlan

Syntax **vlan** *vlan-id*
no vlan

Context config>eth-cfm>domain>association>bridge-identifier

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the bridge-identifier primary VLAN ID. Note that it is informational only, and no verification is done to ensure MEPs on this association are on the configured VLAN. Note: Also see the command description for bridge-identifier.

Default no vlan

Parameters *vlan-id* — Specifies a VLAN ID monitored by MA.

Values 0 — 4094

ccm-interval

Syntax **ccm-interval {10ms | 100ms | 1 | 10 | 60 | 600}**
no ccm-interval

Context config>eth-cfm>domain>association

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the CCM transmission interval for all MEPs in the association.

NOTES:

- 100ms timer value is supported only for Down MEPs on 7210 SAS-D. The minimum timer for UP MEPs on 7210 SAS-D is 1 second.
- On 7210 SAS-E, the minimum timer value supported for Down MEPs (including G8032 Down MEPs) is 1 second. In other words, 100 ms timers are not supported.
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, the minimum timer value supported for both Down Service MEPs and UP Service MEPs is 1 second. 10ms and 100ms CCM timers are supported only for G.8032 Down MEPs.

The no form of the command reverts the value to the default.

Default 10 seconds

Parameters *interval* — Specifies the interval between CCM transmissions to be used by all MEPs in the MA.

remote-mepid

Syntax **[no] remote-mepid mep-id**

Context config>eth-cfm>domain>association

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the remote maintenance association end point (MEP) identifier.

Parameters *mep-id* — Maintenance association end point identifier of a remote MEP whose information from the MEP database is to be returned.

Values 1 — 8191

slm

Syntax **slm**

Context config>eth-cfm

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This is the container that provides the global configuration parameters for ITU-T Synthetic Loss Measurement (ETH-SL).

inactivity-timer

Syntax **inactivity-timer** *timer*
[no] inactivity-timer

Context config>eth-cfm>slm

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

The time the responder keeps a test active. The time between packets exceed this values within a test the responder marks the previous test as complete. The timer treats any new packets from a peer with the same test-id, source-mac and MEP-ID as a new test responding with the sequence number one.

Default 100 seconds

Parameters *timer* — Specifies the amount of time in seconds.

Values 10 100

system

Syntax **system**

Context config>eth-cfm

Description **Platforms Supported**: 7210 SAS-D.

This command configures Connectivity Fault Management General System parameters.

enable-dmm-version-interop

Syntax **enable-dmm-version-interop**
no enable-dmm-version-interop

Context config>eth-cfm>system

Description **Platforms Supported**: 7210 SAS-D.

This command enables processing of DMM version 1 messages that are received by the node, as specified by ITU-T Y.1731 standards for interoperability for nodes that support either version 0 or version 1 implementation. 7210 SAS nodes support processing as recommended for DMM version 0 messages. In other words, without the use of this command, 7210 SAS nodes only process DMM version 0 messages and do not respond to DMM version 1 messages.

When this command is enabled, the 7210 SAS processes all received DMM PDU messages according to version 0 rules. DMM reply messages are sent with version field values that are identical to that of the received DMM PDU. For example, if a DMM PDU with a version value of 1 is received, the DMM reply message is sent with a version field value of 1.

On the 7210 SAS-D, when this command is disabled, timestamping for DMM messages is applied in hardware for both receive and transmit directions. When this command is enabled, timestamping for DMM messages is applied in hardware for the receive direction only, and timestamping for the transmit direction is applied in software by the CPU.

Default no enable-dmm-version-interop

tools

Syntax **tools**

Context root

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command enables the context to enable useful tools for debugging purposes.

Default none

Parameters *dump* — Enables dump tools for the various protocols.

perform — Enables tools to perform specific tasks.

perform

Syntax **perform**

Context tools

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command enables the context to enable tools to perform specific tasks.

Default none

service

Syntax **services**

Context tools>perform

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command enables the context to configure tools for services.

id

Syntax **id service-id**

Context tools>perform>service

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command enables the context to configure tools for a specific service.

Parameters *service-id* — Specify an existing service ID.

Values 1 — 2147483647

endpoint

Syntax **endpoint** *endpoint-name*

Context tools>perform>service>id

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command enables the context to configure tools for a specific VLL service endpoint.

Parameters *endpoint-name* — Specify an existing VLL service endpoint name.

force-switchover

Syntax **force-switchover** *sdp-id:vc-id*
no force-switchover
force-switchover spoke-sdp-*fec [1..4294967295]*

Context tools>perform>service>id>endpoint

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command forces a switch of the active spoke SDP for the specified service.

Parameters *sdp-id:vc-id* — Specify an existing spoke SDP for the service.

spoke-sdp-fec spoke-sdp-fec-id — The spoke-sdp-fec-id for a FEC129 AII Type 2 spoke-sdp. This parameter is mutually exclusive with sdp:vc-id used for a FEC 128 spoke-sdp.

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 : 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#
```

eval-pw-template

Syntax **eval-pw-template**

Context tools>perform>service

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command re-evaluates the pseudowire template policy.

Parameters policy-id — Specifies the pseudowire template policy.

eval-expired-fec

Syntax **eval-expired-fec** *spoke-sdp-fec-id*
eval-expired-fec *all*

Context tools>perform>service>pw-routing

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command will reset the retry counter and retry timer for the specified spoke-SDP and attempt to establish the spoke-SDP again.

spoke-sdp-fec-release

Syntax **spoke-sdp-fec-release** *global-id[:prefix[:ac-id]]*

Context tools>perform>service

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command is used to clear the MS-PW bindings associated with particular SAII or TAII on an SPE.

VLL Services

In This Chapter

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

Topics in this section include:

- Ethernet Pipe (Epipe) Services on page 146
- Epipe Oper State decoupling on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C on page 149
- VLL Service Considerations on page 198

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 147
  - → on page 148
  - → QoS Policies on page 199
  - → Filter Policies on page 199
  - → MAC Resources on page 199
- Configuring a VLL Service with CLI on page 201

Epipe Service Overview

An Epipe service is a Layer 2 point-to-point service where the customer data is encapsulated and transported across a service provider's network. An Epipe service is completely transparent to the subscriber's data and protocols. The 7210 SAS 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.

The following support is available on different platforms:

- The 7210 SAS-E, support only local Epipe services. The 7210 SAS-E provide Null and Dot1q SAPs to provide a point-point Layer 2 local services.
- The 7210 SAS-D, support only local Epipe services. The 7210 SAS-D provide Null, QinQ SAPs and Dot1q SAPs to provide a point-point Layer 2 local services.
- The 7210 SAS-K 2F2T1C, support only local Epipe services. The 7210 SAS-K 2F2T1C provide Null, QinQ SAPs and Dot1q SAPs to provide a point-point Layer 2 local services.
- The 7210 SAS-K 2F4T6C supports both local and distributed Epipe services. The 7210 SAS-K12 provide Null, QinQ SAPs and Dot1q SAPs to provide a point-point Layer 2 local and distributed service. It supports the use of MPLS tunnels for distributed services.

Each SAP configuration includes a specific port on which service traffic enters the 7210 SAS router from the customer side (also called the access side). Each port is configured with an encapsulation type. If a port is configured with an IEEE 802.1Q (referred to as Dot1q) encapsulation, then a unique encapsulation value (ID) must be specified.

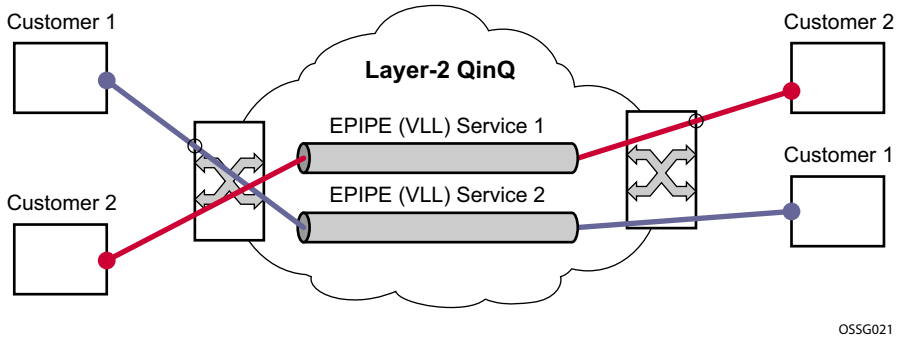

Figure 16: Local Epipe/VLL Service using Access or Access-uplink ports on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

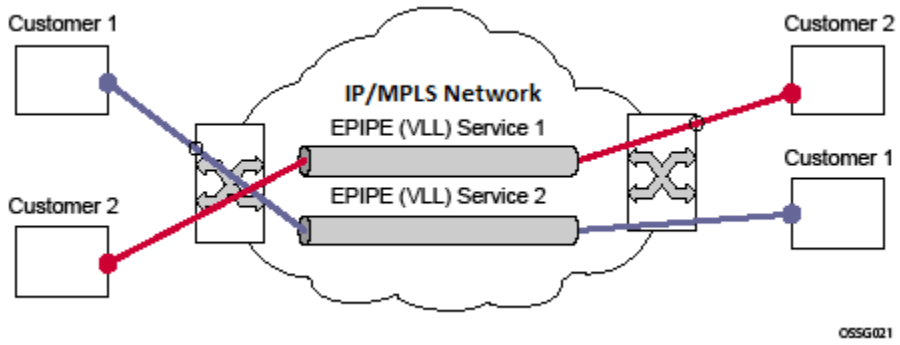

Figure 17: Distributed Epipe/ VLL service using Network ports on 7210 SAS-K 2F4T6C

Epipe Oper State decoupling on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

An Epipe service transitions to an operation state, ‘Down’ when only a single entity SAP or Binding is active and the operation state of the mate is down or displays an equivalent state. The default behavior does not allow operators to validate the connectivity and measure performance metrics. With this feature an option is provided to allow operators to validate the connectivity and measure performance metrics of an Epipe service prior to the customer hand-off. The operator can also maintain performance and continuity measurement across their network regardless of the connectivity between the terminating node and the customer. If the SAP between the operator and the customer enters a Oper Down state, the Epipe remains Operationally UP, so the results can continue to be collected uninterrupted. The operator receives applicable port or SAP alerts/alarms. This option is available only for the customer facing SAP failures. If a network facing SAP or Spoke-SDP fails the operational state of the Epipe service is set to 'Down'. In other words, there is no option to hold the service in an UP state, if a network component fails.

The following functionality is supported:

- Configuration under SAP is required to change the default behavior of the Epipe service in response to the SAP failure.
- The user can create a SAP on a LAG where the LAG has no port members. In this case, the operator configures the “*ignore-oper-state*” on the SAP and the service remains operational. However, as there are no ports existing in the LAG member group, there is no extraction function that can be created. This feature protects against an established working configuration with full forwarding capabilities from failing to collect PM data. The user should shutdown their equipment and place the Epipe SAP in an operationally down state.
- The SAP connecting the provider equipment to the customer is configured to hold the Epipe service status UP when the customer facing SAP enters any failed state. Only one SAP per Epipe is allowed to be configured.
- Any failure of the network entity (network SAP or SDP-Binding) still cause the Epipe service to transition to OPER=DOWN.
- As the service remains operationally up, all bindings should remain operationally up and should be able to receive and transmit data. The PW status represents the failed SAP in the LDP status message, but this does not prevent the data from using the PW as a transport, in or out. This is the same as LDP status messaging.
- The SAP failure continues to trigger normal reactions, except the operational state of the service
- ETH-CFM PM measurement tools (DMM/SLM) can be used with the UP MEP on the failed SAP to collect performance metric. Additionally, CFM troubleshooting tools & connectivity (LBM, LTM, AIS, CCM) can be used and will function normally.

- ETH-CFM CCM processing and fault propagation does not change. Even when a SAP fails with the hold service UP configuration, CCM sets the Interface Status TLV to "*Down*".
- VPLS services remain operationally UP until the final entity in the service enters a failed operational state. There are no changes to VPLS services and the change is specific to Epipe.

Dynamic Multi-Segment Pseudowire Routing

Dynamic Multi-Segment Pseudowire Routing - Overview

Note: 7210 SAS-K 2F4T6C nodes can only be configured as a T-PE node and not as a S-PE node.

The following sections describe the end-to-end solution with BGP PW-routing, assuming appropriate platforms are used for various functions.

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 18 illustrates the operation of dynamic MS-PWs.

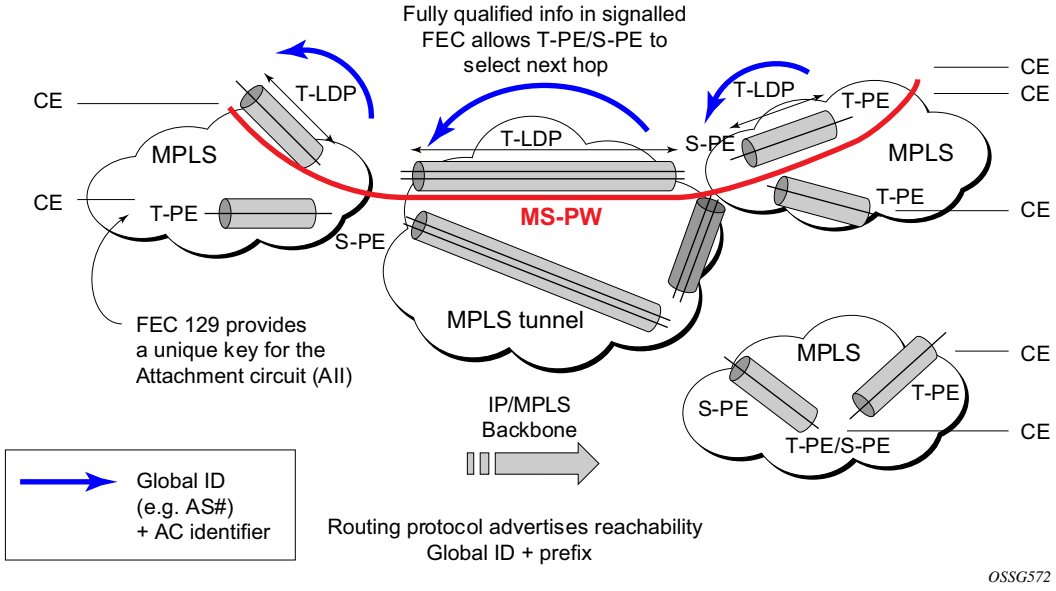

Figure 18: Dynamic MS-PW Overview

The FEC 129 AII Type 2 structure depicted in Figure 19 is used to identify each individual pseudowire endpoint:

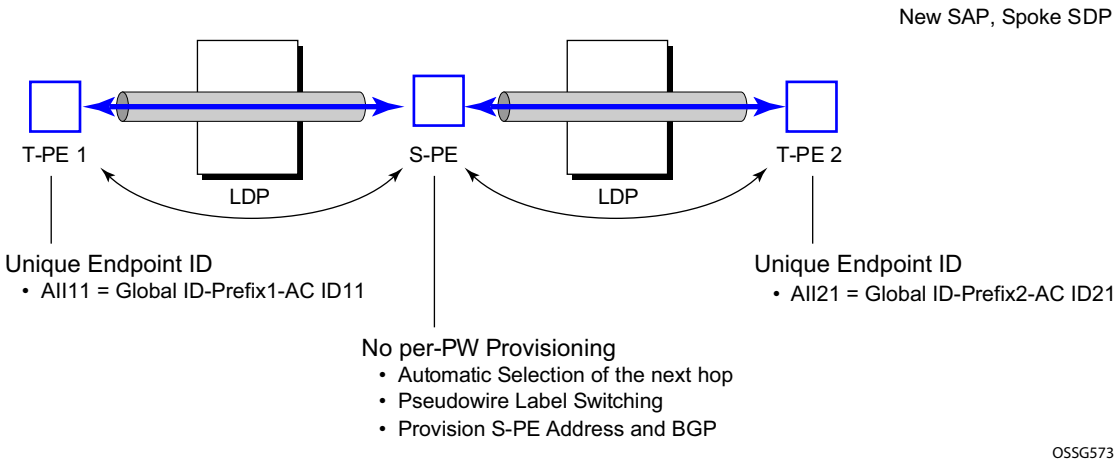

Figure 19: Figure 2 MS-PW Addressing using FEC129 AII Type 2

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 20:

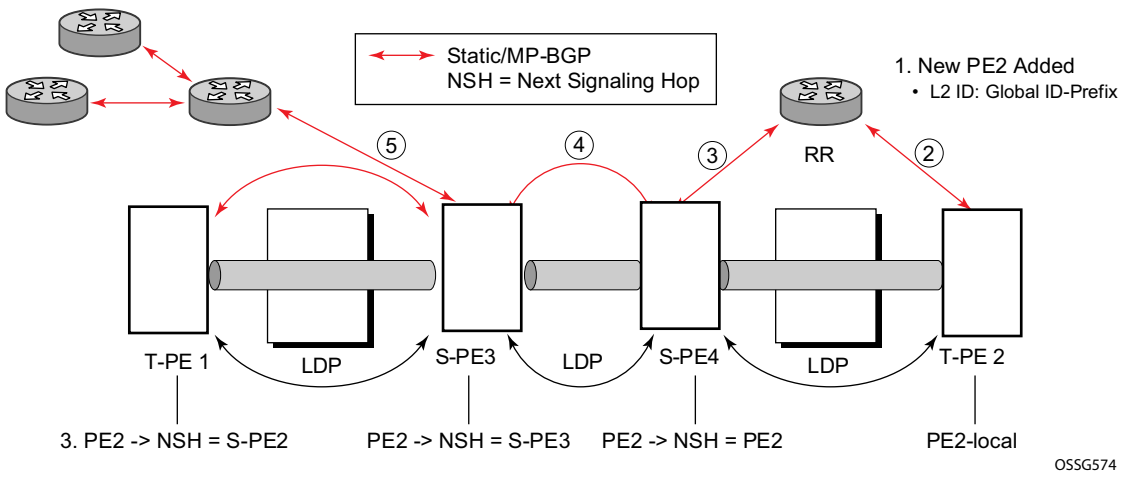

Figure 20: Advertisement of PE Addresses by PW Routing

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 21.

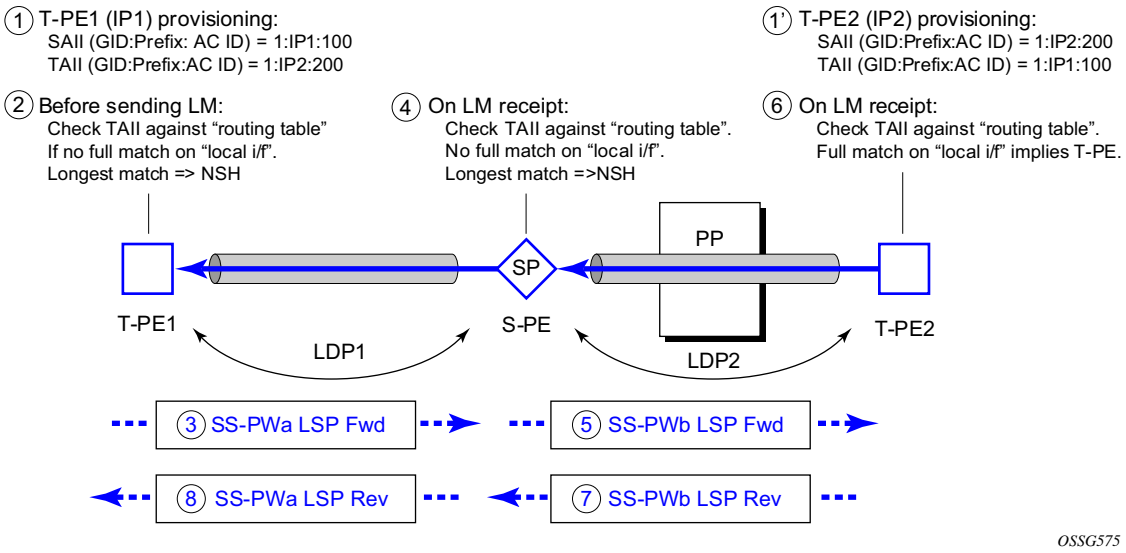

Figure 21: Signaling of Dynamic MS-PWs using T-LDP

In step 1 and 1' the T-PEs are configured with the local and remote endpoint information, Source AII (SAII), Target AII (TAII). On the 7210, the AIIs are locally configured for each spoke SDP, according to the model shown in Figure 22. The 7210 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.

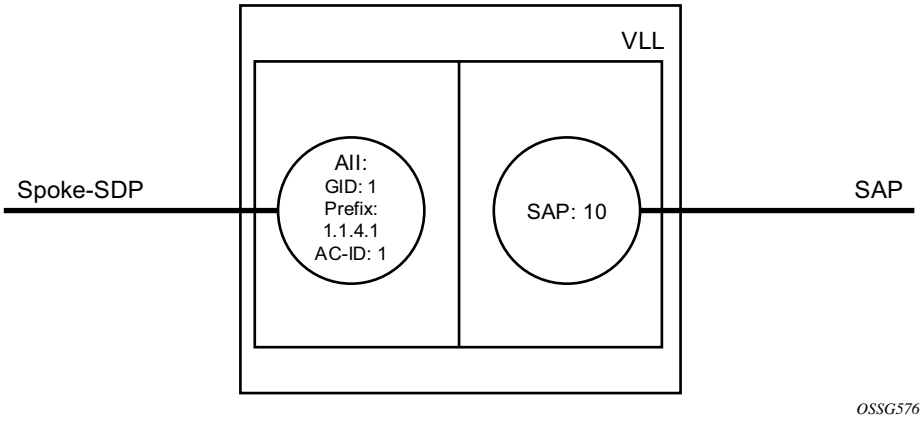

Figure 22: Mapping of AII to SAP

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 23.

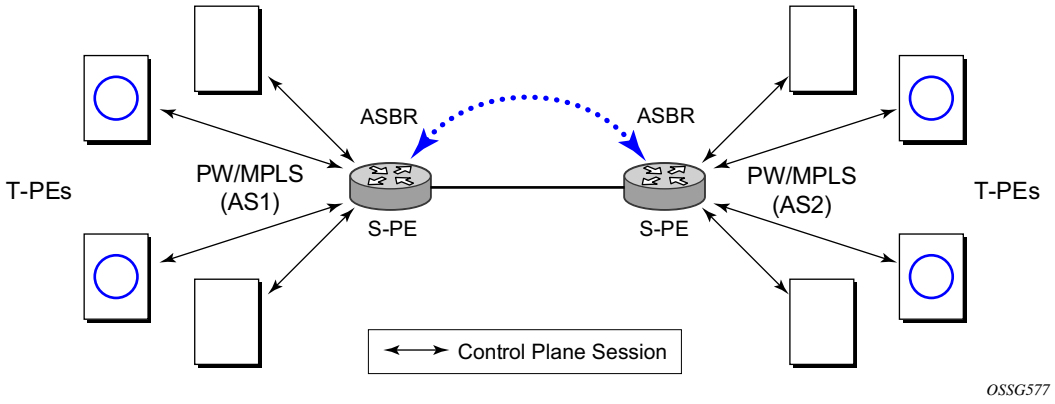

Figure 23: VLL Using Dynamic MS-PWs, Inter-AS Scenario

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

Multi-Segment Pseudowire Routing - Pseudowire Routing

Note: 7210 SAS-K 2F4T6C nodes can only be configured to a T-PE node and not a S-PE node.

The following sections describe the end-to-end solution with BGP PW-routing, assuming appropriate platforms are used for various functions.

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 7210 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 7210 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.

Multi-Segment Pseudowire Routing - 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.

Multi-Segment Pseudowire Routing - 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.

Multi-Segment Pseudowire 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.

Multi-Segment Pseudowire Routing - 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 7210 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.

Multi-Segment Pseudowire Routing - 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.

Multi-Segment Pseudowire Routing - 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 TAII.

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.

Multi-Segment Pseudowire Routing - 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.

Multi-Segment Pseudowire Routing - Pseudowire Redundancy

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

Figure 24 illustrates the use of pseudowire redundancy.

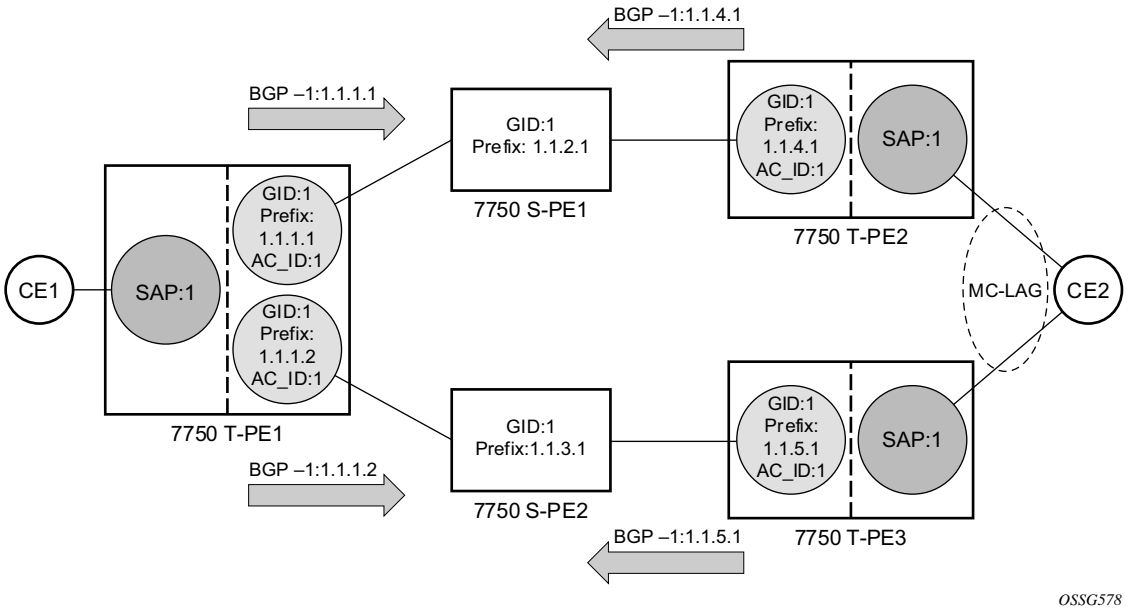

Figure 24: Pseudowire Redundancy

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

- Each MS-PW in the redundant set must have a unique SAII/TAII set and is 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.

Multi-Segment Pseudowire Routing - 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.

Multi-Segment Pseudowire Routing - VCCV-Ping on Dynamic MS-PWs

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

If the SAII:TAII is entered by the user in the vccv-ping command, then those values are 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 SAII:TAII for a remote T-PE of that MS-PW. Note that is SAII:TAII is entered in addition to the *spoke-sdp-fec-id*, then the system will verify the entered values against the values stored in the context for that *spoke-sdp-fec-id*.

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

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

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

Multi-Segment Pseudowire Routing - VCCV-Trace on Dynamic MS-PWs

The 7210 SAS 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.

Multi-Segment Pseudowire Routing - 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 7210 nodes. The network consists of two 7210 T-PEs and two 7210 playing the role of S-PEs, as shown in the following figure. Each 7210 peers with its neighbor using LDP and BGP.

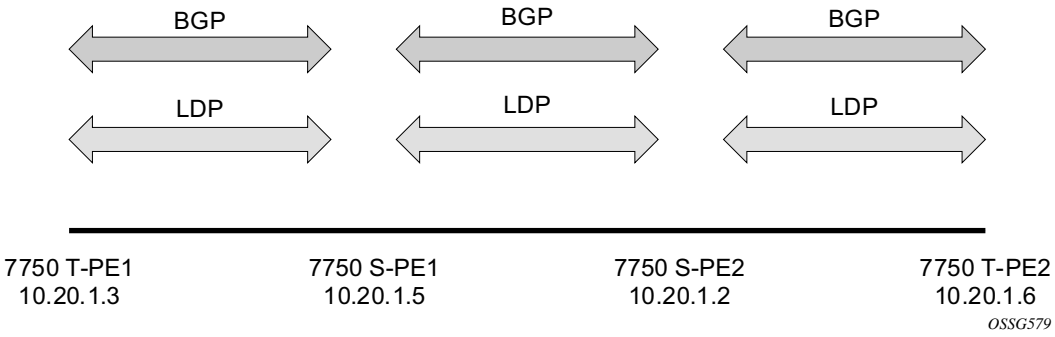

Figure 25: Dynamic MS-PW Example

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. 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.

T-PE-1

```
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
```

T-PE-2

```
    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
```

PW Redundancy - Master-Slave Operation

NOTE: 7210 SAS devices support only standby-signaling-master option. 7210 does not support the CLI command standby-signaling-slave. In the discussion below, reference to standby-signaling-slave command is only used to describe the solution. 7210 device can be used only where standby-signaling-master is used in the example below.

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.

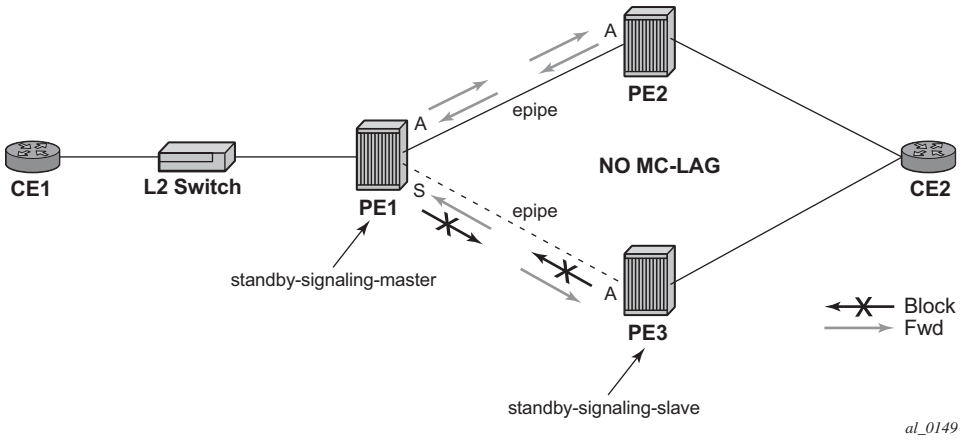

Figure 26: Master-Slave Pseudowire Redundancy

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 PW Redundancy - Operation of Master-Slave Pseudowire Redundancy with Existing Scenarios on page 171.

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.

PW Redundancy - Local Rules at Slave VLL PE

It must not be possible to configure 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.

PW Redundancy - Operation of Master-Slave Pseudowire Redundancy with Existing Scenarios

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

PW Redundancy - VLL Resilience

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

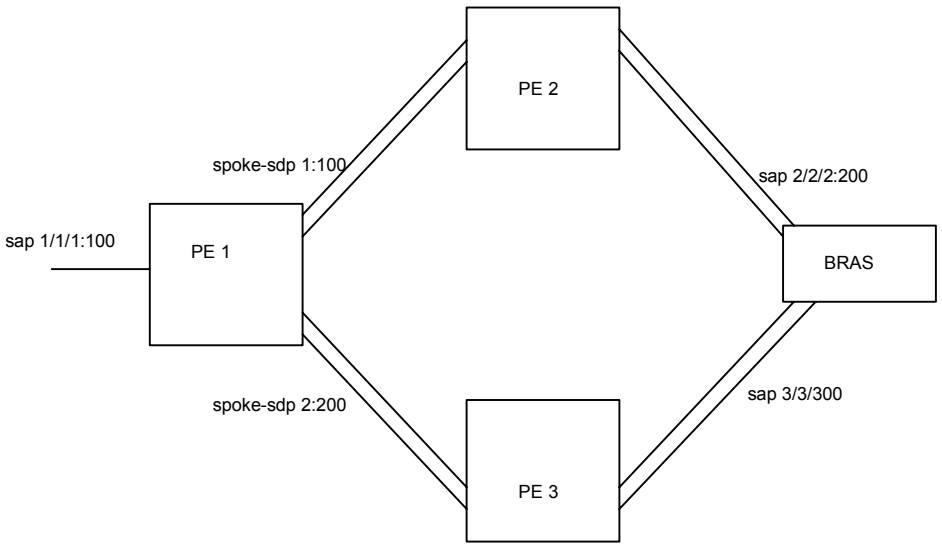

Figure 27: VLL Resilience

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

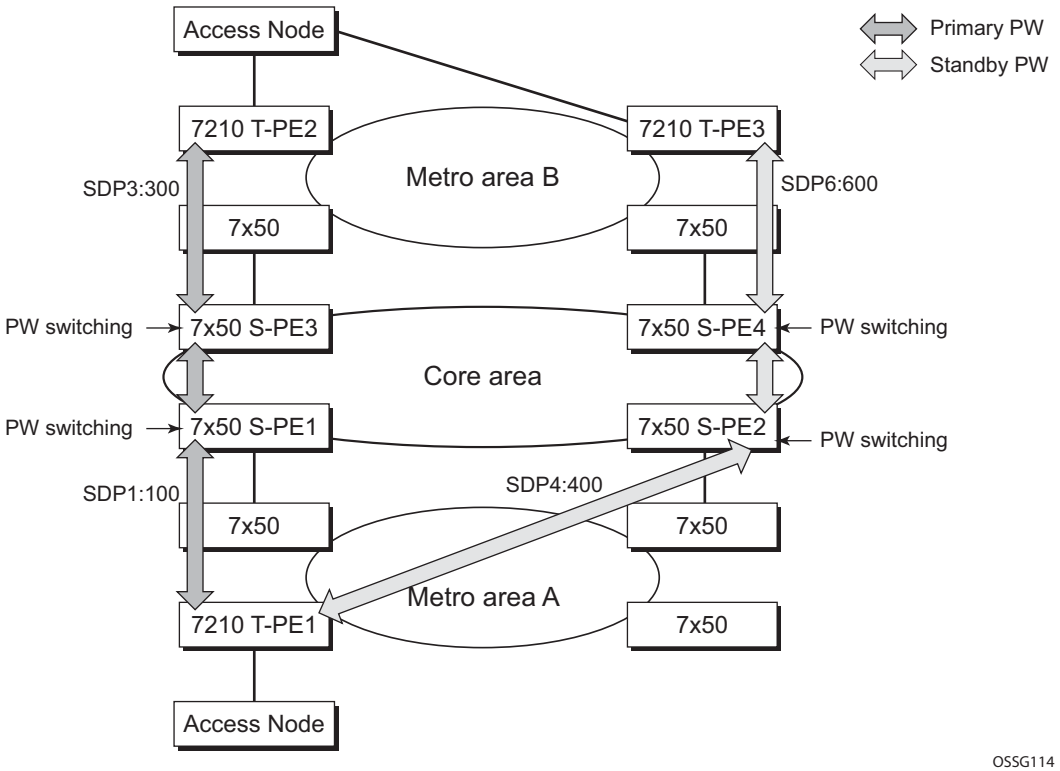

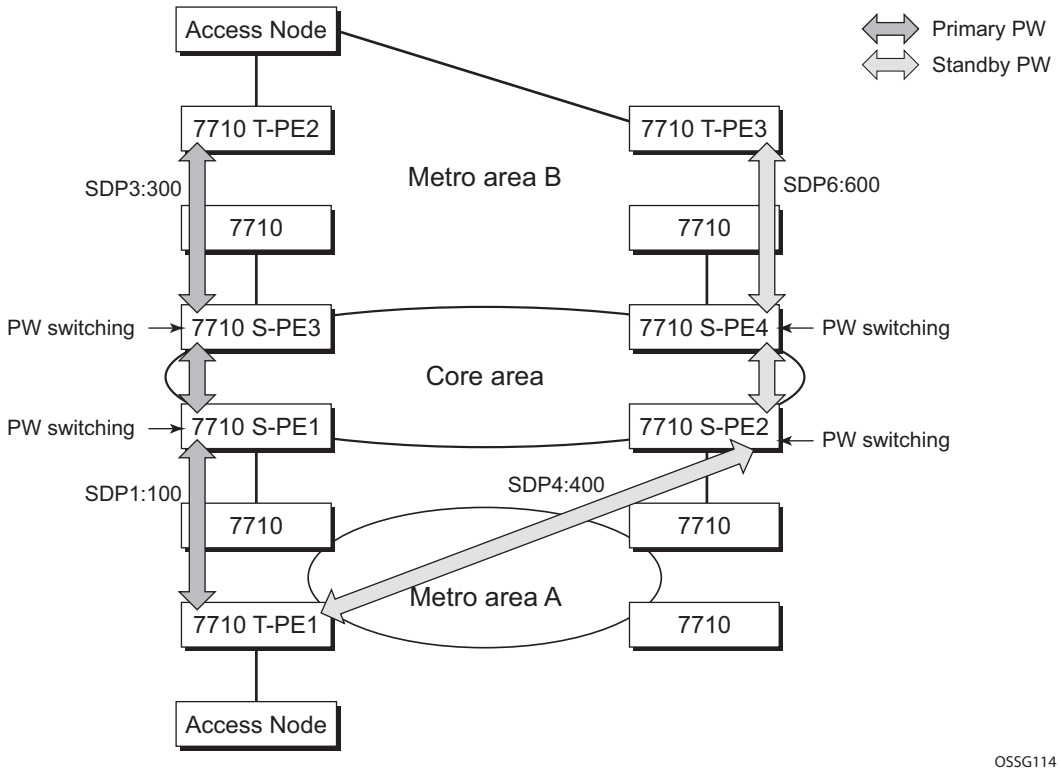

Figure 28: VLL Resilience with Pseudowire Redundancy and Switching

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 339, 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.

PW Redundancy - VLL Resilience for a Switched PW Path

Note: 7210 SAS-K 2F4T6C nodes can only be configured to a T-PE node and not a S-PE node.

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

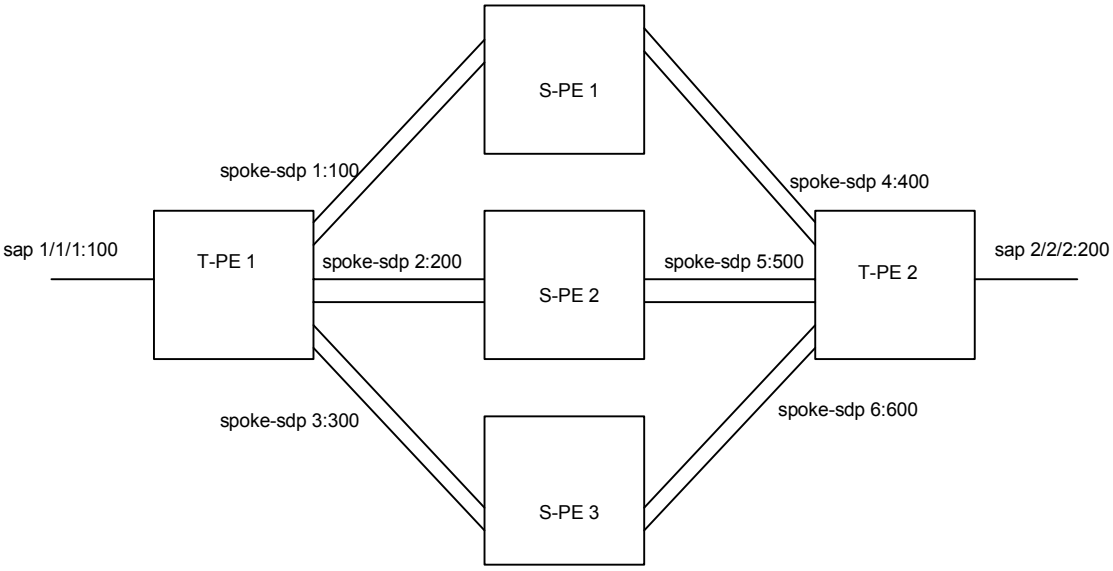

Figure 29: VLL Resilience with Pseudowire Switching

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
          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
```

Access Node Resilience Using MC-LAG and Pseudowire Redundancy

Figure 30 shows the use of both Multi-Chassis Link Aggregation (MC-LAG) in the access network and pseudowire redundancy in the core network to provide a resilient end-to-end VLL service to the customers.

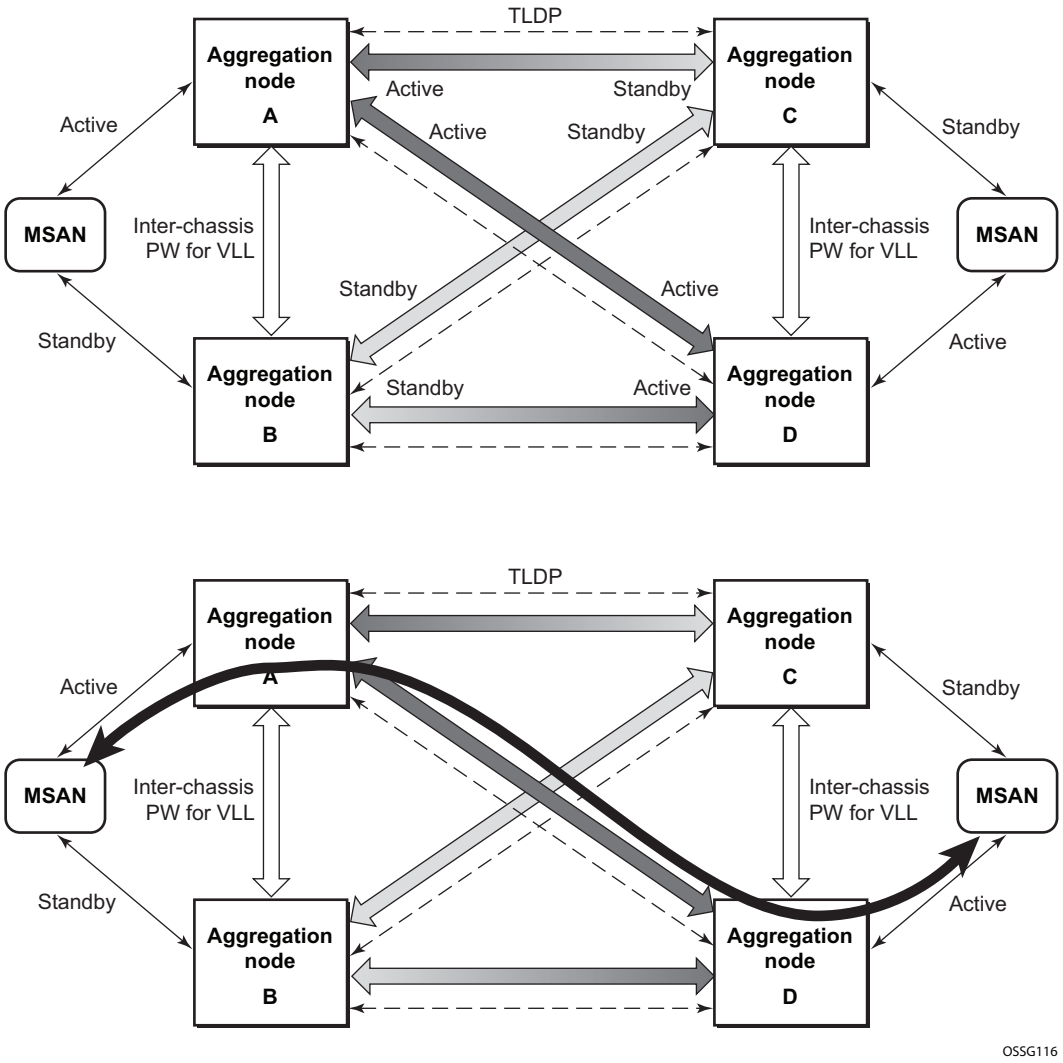

Figure 30: Access Node Resilience

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 7210 SAS 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.

7210 SAS 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, a

7210 SAS 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 7210 SAS standby node forwards these packets to the active node bye 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. A 7210 SAS 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.

VLAN Range for SAPs (or Dot1q range SAPs) in an Epipe Service

7210 SAS VLAN ranges (or Dot1q range SAPs) provide a mechanism to group a range of VLAN IDs as a single service entity. This allows the operator to provide the service treatment (forwarding, ACL, QoS, Accounting, and others) to the group of VLAN IDs as a whole.

NOTE: Grouping a range of VLAN IDs to a SAP is supported only for Virtual Leased Lines (VLL) Ethernet services.

Processing behavior for Dot1q range SAPs in access-uplink mode on 7210 SAS-D and 7210 SAS-K 2F2T1C

NOTE: Dot1q range SAPs are supported only on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C and 7210 SAS-D. It is not supported on 7210 SAS-E.

The access SAPs that specifies VLAN range values using connection-profile (also known as, dot1q range SAPs) is allowed in Epipe service and in VPLS service. For more information on functionality supported, see VLAN Range SAPs feature Support and Restrictions on page 179. The system allows only one range SAP in an Epipe service. It fails any attempt to configure more than one range SAP in an Epipe service. Range SAP can be configured only on access ports. The other endpoint in the Epipe service has to be a “Q.* SAP” in access-uplink mode. The processing and forwarding behavior for packets received on range SAPs are listed below:

- No VLAN tags are removed/stripped on ingress of access dot1q SAP configured to use VLAN ranges. A single tag (Q1) is added to the frame when it is forwarded out of the Q1.* access-uplink SAP.
- When a packet is received on the access-uplink Q1.* SAP, the outermost tag is removed and the packet is forwarded out of the access dot1q range SAP. The system does not check if the inner VLAN tag matches the VLANs IDs (both range and individual values specified in the “connection-profile”) of the dot1q access SAPs configured in the service.
- The dot1q range sap can be supported in a service with svc-sap-type set to ‘dot1q-range’.

VLAN Range SAPs feature Support and Restrictions

- The access SAPs that specifies VLAN range values (using connection-profile) is allowed only in E-Pipe service. The system allows only one range SAP in an Epipe service. It will fail any attempt to configure more than one range SAP in an Epipe service. Range SAP can be configured only on access ports.
- On 7210 SAS-D, the dot1q range sap can be configured only in a service with svc-sap-type set to 'dot1q-range'.
- On 7210 SAS-K 2F2T1C, the dot1q range sap can be configured only in a service with svc-sap-type set to 'dot1q-range'.
- The access SAPs using VLAN range values are allowed only for port or a LAG configured for Dot1q encapsulation.
- A connection profile is used to specify either range of VLAN IDs or individual VLANs to be grouped together in a single SAP.
- On 7210 SAS-D, Dot1q default SAP and dot1q range SAP is not allowed to be configured on the same access port. In other words, they are mutually exclusive. This restriction does not apply to 7210 SAS-K 2F2T1C.
- Multiple “connection-profile” can be used per port or Lag if the VLAN value specified by each of them does not overlap. The number of VLAN ranges available per port/LAG is limited. The available number must be shared among all the SAPs on the port/LAG.
- “Connection-profile”, associated with a SAP cannot be modified. To modify a connection profile, it must be removed from all SAPs using it.
- ACL support - Filter policies are supported on SAP ingress. IP criteria and MAC criteria based filter policy is available for use with access SAPs. Note: For more information on ACL on range SAPs, see the 7210 SAS Router Configuration Guide. Note: For more information on ACL on range SAPs, see the 7210 SAS Router Configuration Guide.
- On 7210 SAS-K 2F2T1C, the outermost VLAN tag and Inner VLAN tag matches are supported for both ingress/egress ACL and ingress QoS classification.

- On 7210 SAS-D Access SAP egress filter are not supported for SAPs configured in a dot1q-range service.This restriction does not apply to 7210 SAS-K 2F2T1C.
- On 7210 SAS-D Access-uplink SAP egress filter are not supported for SAPs configured in a dot1q-range service. This restriction does not apply to 7210 SAS-K 2F2T1C.
- QoS - On 7210 SAS-D, Ingress classification, metering with hierarchical metering support for SAP ingress.

  → SAP ingress classification criteria is available for use with VLAN range SAPs is similar to that available for other SAPs supported in an Epipe service. Dot1p based ingress classification uses the Dot1p bits in the outermost VLAN tag for matching.
- QoS on 7210 SAS-K 2F2T1C- Ingress classification, queuing with hierarchical shaping support for SAP ingress is available.

  → SAP ingress classification criteria is available for use with VLAN range SAPs is similar to that available for other SAPs supported in an Epipe service. Dot1p based ingress classification uses the Dot1p bits in the outermost VLAN tag for matching.
- The amount of hardware resources (such as CAM entries used for matching in QoS classification and ACL match, meters used in SAP ingress policy, and others.) consumed by a single range SAP is equivalent to the amount of resources consumed by a single SAP that specifies a single VLAN ID for service identification. In other words, the hardware can match a range of VLAN values and hence uses 'X' resources for a SAP using a VLAN range instead of X * n, where 'n' is the number of VLANs specified in the range and X is the amount of QoS or ACL resources needed.
- Ingress accounting support is like the support available for other SAPs in an Epipe service. Count of packets or octets received from individual VLANs configured in the connection profile is not available. No support for Egress SAP statistics and accounting is available.
- Mirroring is supported.

Processing behavior for Dot1q range SAPs on SAS-K 2F4T6C

The access SAPs that specifies VLAN range values (using connection-profile) is allowed only in an Epipe service. The system allows only one range SAP in an Epipe service. It will fail any attempt to configure more than one range SAP in an Epipe service. Range SAP can be configured only on access ports. The other endpoint in the Epipe service has to be a Q.* access SAP or a Q.* access uplink SAP or a spoke-sdp (PW) in network mode. The Spoke-SDP processing and forwarding behavior for packets received on range SAPs are listed below:

- No VLAN tags are removed/stripped on ingress of the access dot1q SAPs using VLAN range connection profile.

- When the other endpoint in the service is configured to be an Q1.* access SAP or a Q1.* access-uplink SAP, 7210 adds another tag to the packet and forwards it out of that SAP. If the other endpoint in the service is configured to be a spoke-SDP whose vc-type is set to vc-ether, 7210 SAS adds the appropriate MPLS PW and LSP encapsulations and forwards it out of the SDP.
- In the reverse direction, when the other endpoint is a Q1.* access SAP or a Q1.* access-uplink SAP and a packet is received on it, 7210 SAS removes the outermost VLAN tag and forwards the packet out of the access dot1q SAP using VLAN ranges. When the other endpoint is a spoke-sdp (whose vc-type is set to vc-ether), 7210 SAS removes the MPLS PW and LSP encapsulation and forwards the packet out of the access dot1q SAP using VLAN ranges. The system does not check if the VLAN in the packet matches the VLAN IDs of the dot1q access SAPs configured in the service. In other words, when forwarding the traffic out of the range SAPs, the outermost VLAN tag value is not matched against the range configuration.

VLAN Range SAPs feature Support and Restrictions on SAS-K 2F4T6C

- The access SAPs that specifies VLAN range values (using connection-profile) is allowed only in E-Pipe service. The system allows only one range SAP in an Epipe service. It will fail any attempt to configure more than one range SAP in an Epipe service. Range SAP can be configured only on access ports.
- On 7210 SAS-K 2F4T6C, the dot1q range sap can be configured only in a service with svc-sap-type set to 'dot1q-range'.
- The access SAPs using VLAN range values are allowed only for Dot1q encapsulation port.
- A connection profile is used to specify either range of VLAN IDs or individual VLANs to be grouped together in a single SAP.
- On 7210 SAS-K 2F4T6C, Dot1q default SAP and dot1q range SAP can co-exist on the same access port.
- Multiple “connection-profile” can be used per port or Lag as long as the VLAN value specified by each of them does not overlap. The number of VLAN ranges match resources available per port/LAG is limited. The available number must be shared among all the dot1q range SAPs configured on the port/LAG.
- “Connection-profile”, associated with a SAP cannot be modified. To modify a connection profile, it must be removed from all SAPs using it.
- ACL support - Filter policies are supported on SAP ingress and SAP egress. Only MAC criteria based filter policy is available for use with access SAPs. Note: For more information on ACL on range SAPs, see The 7210 SAS Router Configuration Guide.

Note: For more information on ACL on range SAPs, see the 7210 SAS Router Configuration Guide.

- On 7210 SAS-K 2F4T6C, the outermost VLAN tag and Inner VLAN tag matches are supported for both ingress/egress ACL and ingress QoS classification.
- SDP egress and ingress filter are not supported.
- QoS - On 7210 SAS-K 2F4T6C, Ingress classification, queuing with hierarchical shaping support for SAP ingress.
  - → SAP ingress classification criteria is available for use with VLAN range SAPs is like that available for other SAPs supported in an Epipe service. Dot1p based ingress classification uses the Dot1p bits in the outermost VLAN tag for matching. On access egress, dot1p received from the SDP (on a network port) from another access port is preserved.
- The amount of hardware resources (such as CAM entries used for matching in QoS classification and ACL match, meters used in SAP ingress policy, and others.) consumed by a single range SAP is equivalent to the amount of resources consumed by a single SAP that specifies a single VLAN ID for service identification. In other words, the hardware can match a range of VLAN values and hence uses 'X' resources for a SAP using a VLAN range instead of X * n, where 'n' is the number of VLANs specified in the range and X is the amount of QoS or ACL resources needed.
- Ingress accounting support is like the support available for other SAPs in an Epipe service. Count of packets or octets received from individual VLANs configured in the connection profile is not available. No support for Egress SAP statistics and accounting is available.
- Mirroring is supported.
- Service resiliency mechanisms such as Epipe PW redundancy is supported.

Epipe Pseudowire switching in 7210 SAS-K 2F4T6C

NOTE: 7210 SAS-K 2F4T6C can be configured only as T-PE nodes. It cannot be configured as S-PE nodes. In the text below reference to S-PE nodes refers to other products in the Nokia IP products family that supports S-PE functionality.

The pseudo-wire 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 pseudo-wire switching node instead of what is normally the remote

PE. In addition, the user configures a VLL service on the pseudo-wire switching node using the two SDPs.

The pseudo-wire switching node (i.e. S-PE node) acts in a passive role with respect to signaling of the pseudo-wires. It waits until one or both PEs (i.e. T-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 pseudo-wire switching point TLV is inserted by the switching pseudo-wire 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 pseudo-wire especially if multiple pseudo-wire switching points exist between the two PEs.
- It helps in loop detection of the T-LDP signaling messages where a switching point would receive back a label mapping message it had already relayed.
- The switching point TLV is inserted in pseudo-wire status notification messages when they are sent end-to-end or from a pseudo-wire switching node towards a destination PE.

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

In the following example, the user configures a regular Epipe VLL service PE1 and PE2 (acting as T-PE nodes). These services consist each of a SAP and a spoke SDP. However, the target destination of the SDP is not the remote PE but the pseudo-wire switching node. In addition, the user configures an Epipe VLL service on the pseudo-wire switching node using the two SDPs.

```
|7210 SAS-K 2F4T6C PE1 [T-PE node] (Epipe)|---sdp 2:10---|7450 ESS, 7750 SR, and 7950 XRS
PW SW [S-PE node] (Epipe)|---sdp 7:15---|7210 SAS-K12 PE2 [T-PE node] (Epipe)|
```

Configuration examples can be found in “Configuring Two VLL Paths Terminating on T-PE2”.

Pseudo-wire Switching with Protection

Pseudo-wire 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. This illustrates the use of pseudo-wire redundancy to provide a scalable and resilient VLL service across multiple IGP areas in a provider network.

In the network in Figure 31 on page 184, PE nodes act as masters and pseudo-wire switching nodes act as slaves for the purpose of pseudo-wire 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 pseudo-wire 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 pseudo-wire 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 SPE2 will effect the spoke SDP cross-connect only when both directions of the pseudo-wire have been signaled and matched.

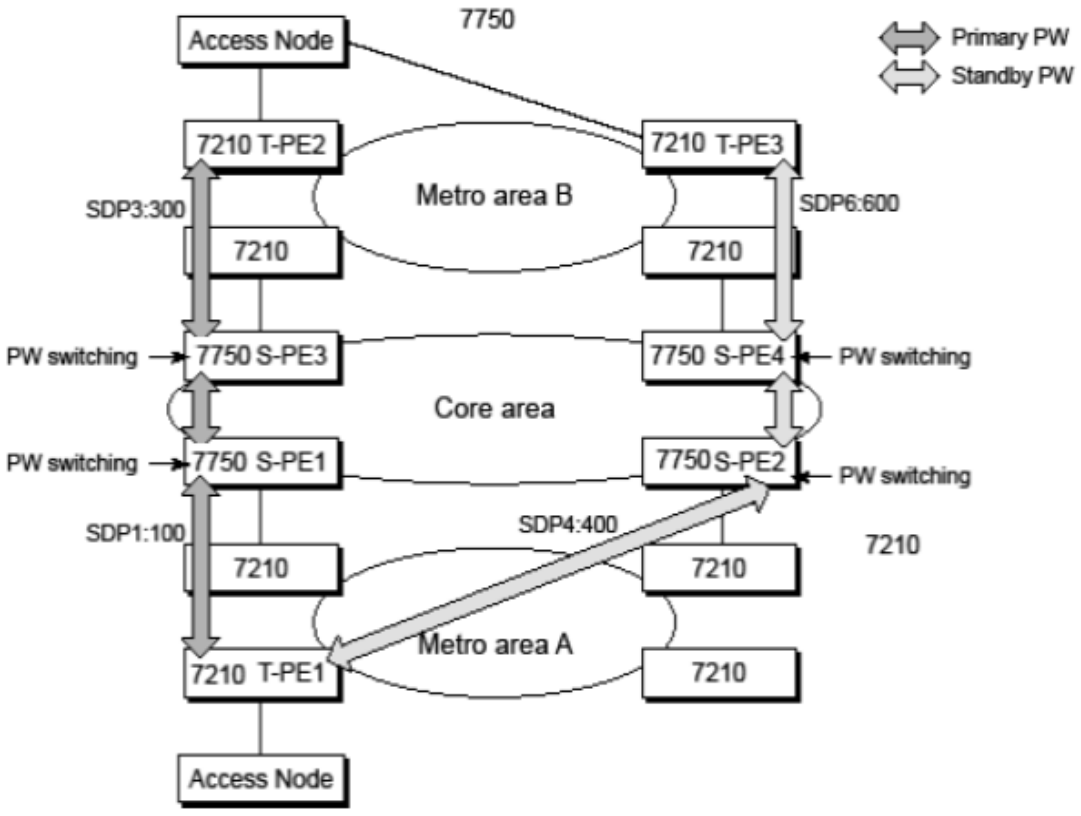

Figure 31: VLL Resilience with Pseudowire Redundancy and Switching

The pseudo-wire switching TLV is useful in a few situations. First, it allows for troubleshooting of the path of the pseudo-wire especially if multiple pseudo-wire 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 pseudo-wire status messages when they are sent from a pseudo-wire switching node towards a destination PE.

Pseudo-wire status messages can be generated by the T-PE nodes and/or the S-PE nodes. Pseudo-wire 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 TPE node and the S-PE should process and pass the message without changes except for the VCID value in the FEC TLV.

Pseudo-wire Switching Behavior

In the network in Figure 31, PE nodes act as masters and pseudo-wire switching nodes act as slaves for the purpose of pseudo-wire 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 pseudo-wire 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. TPE1 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 pseudo-wire have been signaled and matched.

Pseudo-wire status notification messages can be generated by the T-PE nodes and/ or the S-PE nodes. Pseudo-wire 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 pseudo-wire switching TLV, with its system address added to it, to the FEC in the pseudo-wire 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 pseudo-wire 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.

Pseudo-wire Switching TLV

The format of the pseudo-wire switching TLV is as follows:

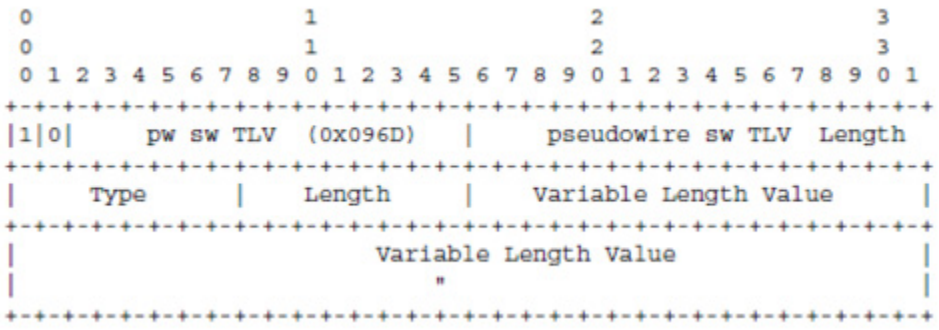

Figure 32: Pseudo-wire Switching TLV

PW sw TLV Length - Specifies the total length of all the following pseudo-wire 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.

Pseudo-wire Switching Point Sub-TLVs

Below are details specific to pseudo-wire switching point sub-TLVs:

- Pseudo-wire ID of last pseudo-wire segment traversed - This sub-TLV type contains a pseudo-wire ID in the format of the pseudo-wire ID
- Pseudo-wire switching point description string - An optional description string of text up to 80 characters long.
- IP address of pseudo-wire switching point.

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

- MH VCCV capability indication

Pseudo-wire Redundancy

Pseudo-wire redundancy provides the ability to protect a pseudo-wire with a pre-provisioned pseudo-wire and to switch traffic over to the secondary standby pseudowire in case of a SAP and/or network failure condition. Normally, pseudo-wires 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 pseudo-wire is also protected. However, there are a couple of applications in which SDP redundancy does not protect the end-to-end pseudo-wire 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 pseudo-wire path is switched in the middle of the network and the pseudowire switching node fails.

Pseudo-wire 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 pseudo-wire or VPLS switchover, which enables carriers to deliver “always on” services across their IP/MPLS networks.

VLL service resiliency with Pseudowire on 7210 SAS-K 2F4T6C

VLL Resilience with Two Destination PE Nodes

Figure 33 illustrates the application of pseudo-wire redundancy to provide Ethernet VLL service resilience for broadband service subscribers accessing the broadband service on the service provider BRAS.

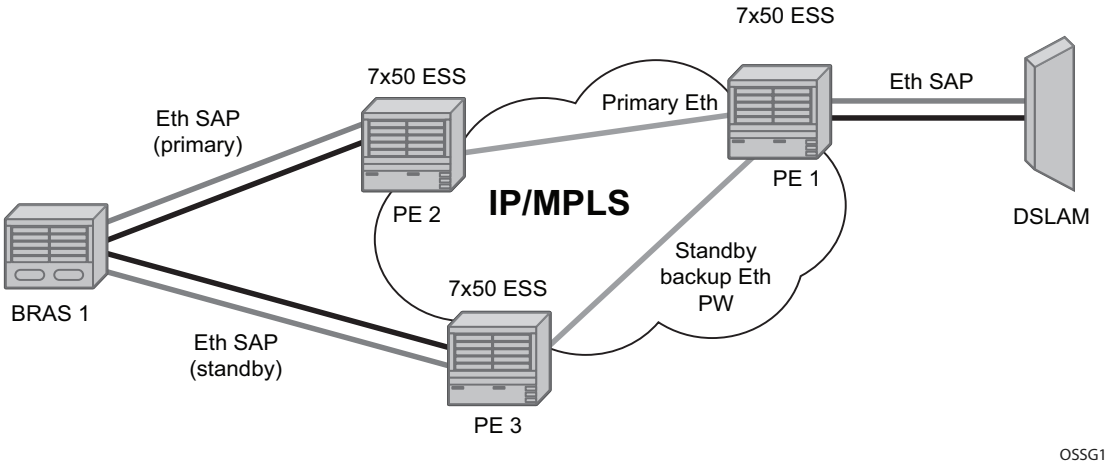

Figure 33: VLL Resilience

If the Ethernet SAP on PE2 fails, PE2 notifies PE1 of the failure by either withdrawing the primary pseudo-wire label it advertised or by sending a pseudo-wire 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 pseudo-wire while transmitting over the backup pseudo-wire. However, in other applications such as those described in Access Node Resilience Using MC-LAG and Pseudo-wire Redundancy, it will be important to minimize service outage to end users.

When the SAP at PE2 is restored, PE2 updates the new status of the SAP by sending a new label mapping message for the same pseudo-wire FEC or by sending pseudo-wire 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 pseudo-wire.

The behavior of the pseudo-wire 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 pseudo-wire:

1. LDP peer (remote PE) node withdrew the pseudo-wire label.
2. T-LDP peer signaled a FEC status indicating a pseudo-wire 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 pseudo-wires need not be the same. In other words, the user can protect a RSVP-TE based spoke SDP with a LDP one. This provides the ability to route the path of the two pseudo-wires over different areas of the network. All VLL service types, for example, Apipe, Epipe, Fpipe, and Ipipe are supported on the 7750 SR.

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

Nokia'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.

On the 7210 SAS, this application is supported with Epipe Service.

VLL Resilience - Master-Slave Operation

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

Figure 34 illustrates the operation of master-slave pseudo-wire 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 pseudo-wire 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.

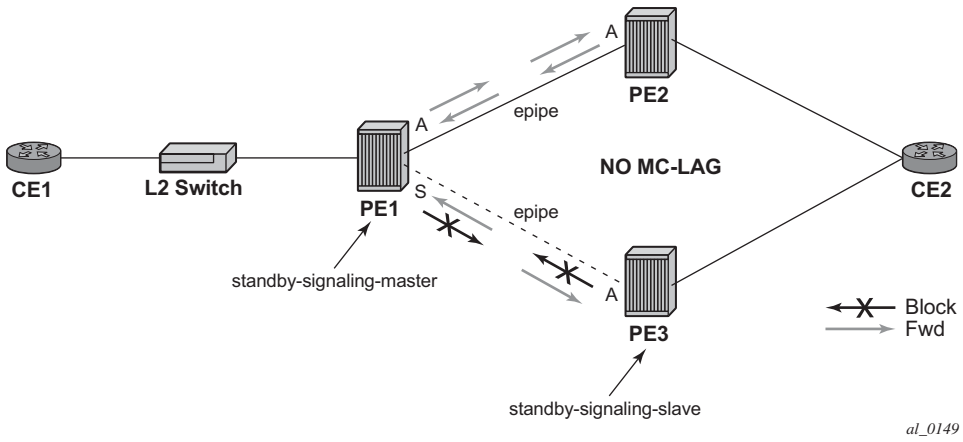

Figure 34: Master-Slave Pseudo-wire Redundancy

Master-slave pseudo-wire redundancy adds the ability for the remote peer to react to the pseudo-wire 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 pseudo-wire status bits:

- Pseudo-wire 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 pseudo-wire 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 pseudo-wire 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 Pseudo-wire Redundancy with Existing Scenarios.

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 pseudo-wire 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.

VLL Resilience - Operation of Master-Slave Pseudo-wire Redundancy with Existing Scenarios

This section discusses how master-slave pseudo-wire redundancy could operate.

Figure 35 shows a VLL resilience path example. An sample configuration follows.

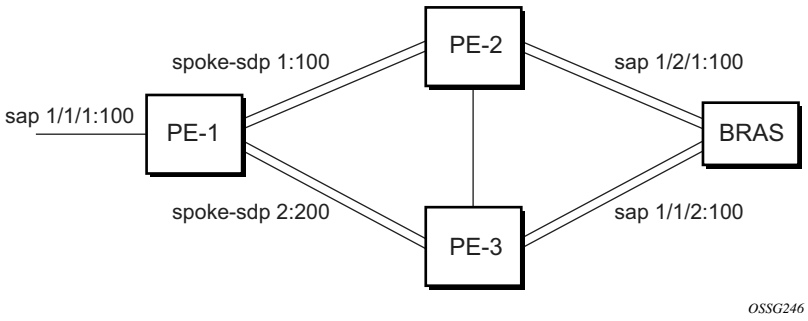

Figure 35: VLL Resilience

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 Pseudo-wire Path

Figure 36 displays a VLL resilience for a switched pseudo-wire path example. A sample configuration follows.

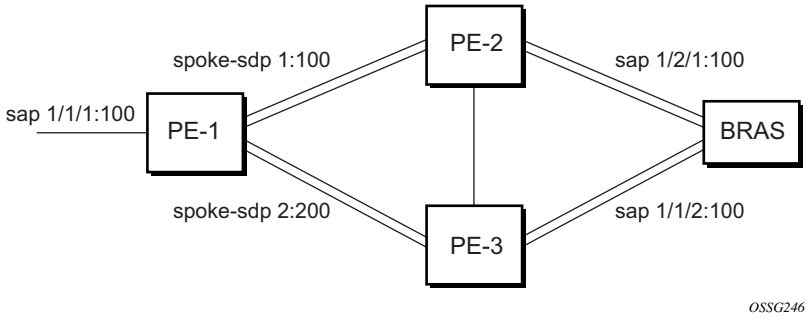

Figure 36: VLL Resilience with Pseudo-wire Switching

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
      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
```

VLL Resilience - Pseudo-wire Redundancy Service Models

This section describes the various pseudo-wire 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, Pseudo-wire Redundancy - Redundant VLL Service Model on page 194.

Pseudo-wire Redundancy - Redundant VLL Service Model

In order to implement pseudo-wire redundancy, a VLL service accommodates more than a single object on the SAP side and on the spoke SDP side. Figure 37 illustrates the model for a redundant VLL service based on the concept of endpoints.

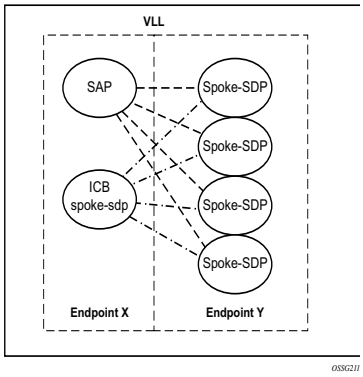

Figure 37: Redundant VLL Endpoint Objects

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 37.

Note that Figure 37 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 pseudo-wire and only switches to a secondary pseudo-wire when it is down the VLL service switches the path to the primary pseudo-wire 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 pseudo-wire to indicate the order in which a secondary pseudo-wire is activated.
- Inter-Chassis Backup (ICB) spoke SDP — Special pseudo-wire used for MC-LAG and pseudo-wire 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.

Pseudo-wire Redundancy - T-LDP Status Notification Handling Rules

Referring to See “Redundant VLL Service Model” on page 221. 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 See “Redundant VLL Service Model” on page 221. can be used on endpoints “X” and “Y”. The following sections refer to the specific combination objects in Figure 37 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 pseudo-wire 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 pseudo-wire 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 pseudo-wire 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 pseudo-wire 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 pseudo-wire 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.

MPLS Entropy Label and Hash Label

MPLS entropy label (RFC 6790) and the Flow Aware Transport label (known as the hash label) (RFC 6391). These labels allow LSR nodes in a network to load-balance labeled packets in a much more granular fashion than allowed by simply hashing on the standard label stack. See the MPLS Guide for further information.

7210 SAS-K 2F4T6C does not support MPLS entropy label and hash label.

VLL Service Considerations

This section describes various of the general 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, MPLS.

SAP Encapsulations

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 Epipe service:

- Ethernet null
- Ethernet dot1q
- QinQ

Note that while different encapsulation types can be used, encapsulation mismatch 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 potentially be double tagged when it is transmitted out of the Dot1q SAP.

QoS Policies

On 7210 SAS-D and 7210 SAS-E, when applied to 7210 SAS device Epipe services, service ingress QoS policies only create the unicast meters defined in the policy. The multi-point meters are not created on the service. Egress QoS policies function as with other services where the class-based queues per port are available.

On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, when applied to 7210 SAS device Epipe services, service ingress QoS policies only create the unicast queues defined in the policy. The multi-point queues are not created on the service. Service egress QoS policies function as with other services where the class-based queues are created as defined in the policy.

NOTES:

- Both Layer 2 or Layer 3 criteria can be used in the QoS policies for traffic classification in a service. For details, refer the 7210 SAS QoS user guide.
- For 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C2, multicast queues are also created for service ingress QoS policy but only unicast queues are used for traffic.

Filter Policies

7210 SAS Epipe services can have a single filter policy associated on both ingress and egress. Both MAC and IP filter policies can be used on Epipe services.

MAC Resources

Epipe services are point-to-point layer 2 VPNs capable of carrying any Ethernet payloads. Although an Epipe is a Layer 2 service, the 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:

- Configuring a VLL Service with CLI on page 201
  - → Basic Configurations on page 202
  - → Common Configuration Tasks on page 202
    - – Common Configuration Tasks on page 202
    - – Using Spoke SDP Control Words on page 239
    - – Pseudowire Configuration Notes on page 240
    - – Configuring VLL Resilience on page 242
    - – Configuring VLL Resilience for a Switched Pseudowire Path on page 243
    - – Service Management Tasks on page 218

  Epipe:
  - → Modifying Epipe Service Parameters on page 219
  - → Disabling an Epipe Service on page 219
  - → Re-Enabling an Epipe Service on page 220
  - → Deleting an Epipe Service on page 220

Basic Configurations

- Creating an Epipe Service for 7210 SAS-E on page 203

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 ingress QoS policies (configured in the **config>qos** context).
  - → Optional - select accounting policy (configured in the **config>log** context).
- Enable the service.

Creating an Epipe Service for 7210 SAS-E

Use the following CLI syntax to create an Epipe service.

CLI Syntax: `config>service# epipe service-id [customer customer-id] [create] [svc-sap-type {null-star | dot1q | dot1q-preserve}] [customer-vid vlan-id]`

```
          description description-string
          no shutdown
```

The following displays an Epipe configuration example:

```
A:ALA-1>config>service# info
----------------------------------------------
...
        epipe 500 customer 1 svc-sap-type null-star create
            description "Local Epipe Service with NULL  SVC_SAP_TYPE"
            no shutdown
        exit
----------------------------------------------
A:ALA-1>config>service#
```

Creating an Epipe Service for 7210 SAS-D

Use the following CLI syntax to create an Epipe service.

CLI Syntax: `config>service# epipe service-id [customer customer-id] [create] [svc-sap-type {null-star | any | dot1q-preserve}] [customer-vid vlan-id]`

The following displays an Epipe configuration example:

```
A:ALA-1>config>service# info
----------------------------------------------
...
        epipe 15 customer 40 svc-sap-type any create
            description "Local Epipe Service with ANY SVC_SAP_TYPE"
            no shutdown
        exit
----------------------------------------------
A:ALA-1>config>service#
```

Creating an Epipe Service with range SAPs

The following displays an example of connection-profile used to configure a range of SAPs and an Epipe configuration using the connection profile:

```
*A:7210SAS>config>connprof# info
----------------------------------------------
        ethernet
            ranges 0 2804-2805 2810-2811 2813 2832-2839
        exit
----------------------------------------------

*A:7210SAS>config>service>epipe# info
----------------------------------------------
            description "Default epipe description for service id 292"
            sap 1/1/4:292.* create
                description "Default sap description for service id 292"
                exit
            exit
            sap 1/1/9:cp-292 create
                description "Default sap description for service id 292"
                exit
            exit
            no shutdown
----------------------------------------------
```

Creating an Epipe Service for 7210 SAS-K 2F2T1C

Use the following CLI syntax to create an Epipe service.

CLI Syntax: `config>service# epipe service-id [customer customer-id] [create] [svc-sap-type {null-star |dot1q|dot1q-preserve|any|dot1q-range|qinq-inner-tag-preserve}]`

The following displays an Epipe configuration example:

```
*A:SAH01-051>config>service>epipe$ info detail
----------------------------------------------
            shutdown
            no description
            service-mtu 1514
            eth-cfm
            exit
            pbb
            exit
----------------------------------------------
*A:SAH01-051>config>service>epipe$

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

Creating an Epipe Service for 7210 SAS-K 2F2T1C with range SAPs

The following displays an example of connection-profile used to configure a range of SAPs and an Epipe configuration using the connection profile:

```
*A:SAH01-051>config>connprof$ info detail
----------------------------------------------
        no description
        ethernet
            no ranges
        exit
----------------------------------------------
*A:SAH01-051>config>connprof$


*A:SAH01-051>config>service>epipe$ info detail
----------------------------------------------
            shutdown
            no description
            service-mtu 1514
            eth-cfm
            exit
            pbb
            exit
----------------------------------------------
*A:SAH01-051>config>service>epipe$

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

Configuring Epipe SAP Parameters

A default QoS policy is applied to each ingress SAP. and a default access egress QoS policy is applied on the port where SAP is egressing. The access egress QoS policy is common to all SAPs on that port. 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 207

CLI Syntax:

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

Local Epipe SAPs

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

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

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 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 1/1/3 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# no shutdown
config>service>epipe>sap# exit

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

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

    exit
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

    exit
exit
```

Distributed Epipe Service

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 service SAPs. On egress, QoS policies are associated with a port. Existing filter policies can be associated with service SAPs on ingress and egress.

Meters (defined in sap-ingress policies) can be applied on ingress. It is associated with SAPs. Scheduler Policies can be applied on egress. It is associated with a port.

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

For SDP configuration information, see Configuring an SDP on page 87. For SDP binding information, see Configuring SDP Bindings (7210 SAS-K 2F4T6C) on page 367.

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# 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# filter ip 1020
          config>service>epipe>sap>ingress# exit
          config>service>epipe>sap# egress
          config>service>epipe>sap>egress# filter ip 6
          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
                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
                exit
            exit
...
----------------------------------------------
A:ALA-2>config>service#
```

Configuring SDP Bindings

Figure 38 displays an example of a distributed Epipe service configuration between two routers, identifying the service and customer IDs, and the unidirectional 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.

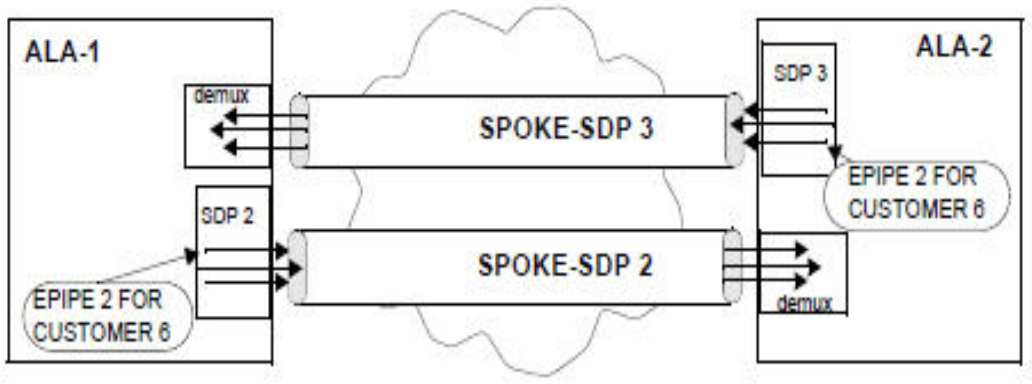

Figure 38: SDPs — Unidirectional Tunnels

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 Service on page 208).

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


ALA-2>config>service# epipe 5500
      config>service>epipe# spoke-sdp 2:456
      config>service>epipe>spoke-sdp# egress
      config>service>epipe>spoke-sdp>egress# vc-label 6600
```

```
      config>service>epipe>spoke-sdp>egress# exit
      config>service>epipe>spoke-sdp# ingress
      config>service>epipe>spoke-sdp>ingress# vc-label 5500
      config>service>epipe>spoke-sdp>ingress# exit
      config>service>epipe>spoke-sdp# no shutdown
```

This example displays the SDP binding for the Epipe service between ALA-1 and ALA-2:

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


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

Using Spoke SDP Control Words

The control word command provides the option to add a control word as part of the packet encapsulation for PW types for which the control word is optional. On 7210 an option is provided to enable it for Ethernet PW (Epipe). The control word might be needed because when ECMP is enabled on the network, packets of a given PW 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 PW (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#
```

Configuring Ingress and Egress SAP Parameters

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

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

This example displays SAP ingress and egress parameters.

```
ALA-1>config>service# epipe 5500
     config>service>epipe# sap 1/1/3:21
     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#
```

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

```
A:ALA-1>config>service#
-----------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to east coast"
            sap 1/1/3:21 create
                ingress
                    qos 555
                    filter ip 1
                exit
            exit
            no shutdown
        exit
-----------------------------------------------
A:ALA-1>config>service#
```

Configuring VLL Resilience

Figure 39 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.

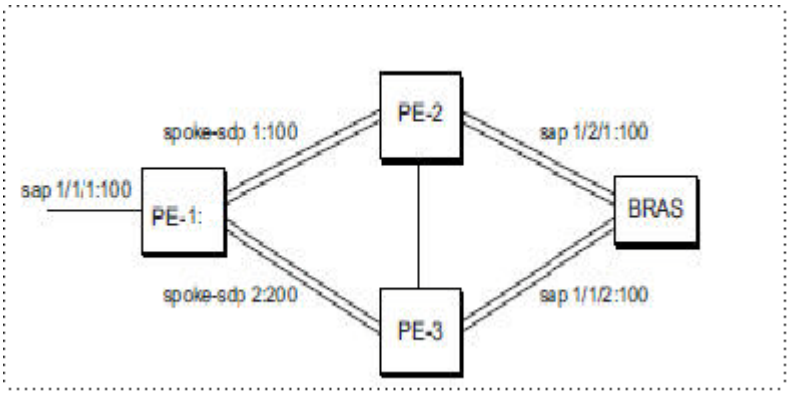

Figure 39: VLL Resilience

PE1:

The following displays an example for the configuration on PE1(a 7210 SAS-K 2F4T6C node).

```
*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 Pseudo-wire Path

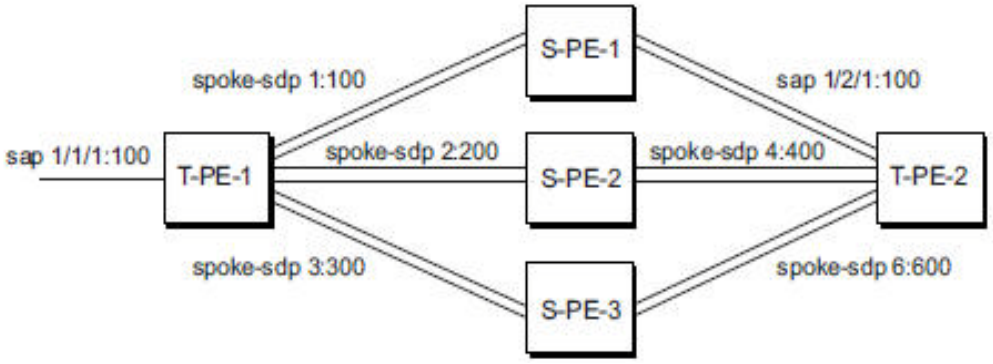

Figure 40: VLL Resilience with Pseudo-wire Switching

T-PE1

The following displays an example for the configuration on TPE1 (which is a 7210 SAS-K 2F4T6C node).

```
*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 Default QinQ SAPs for Epipe Transit Traffic in a Ring Scenario in access-uplink mode

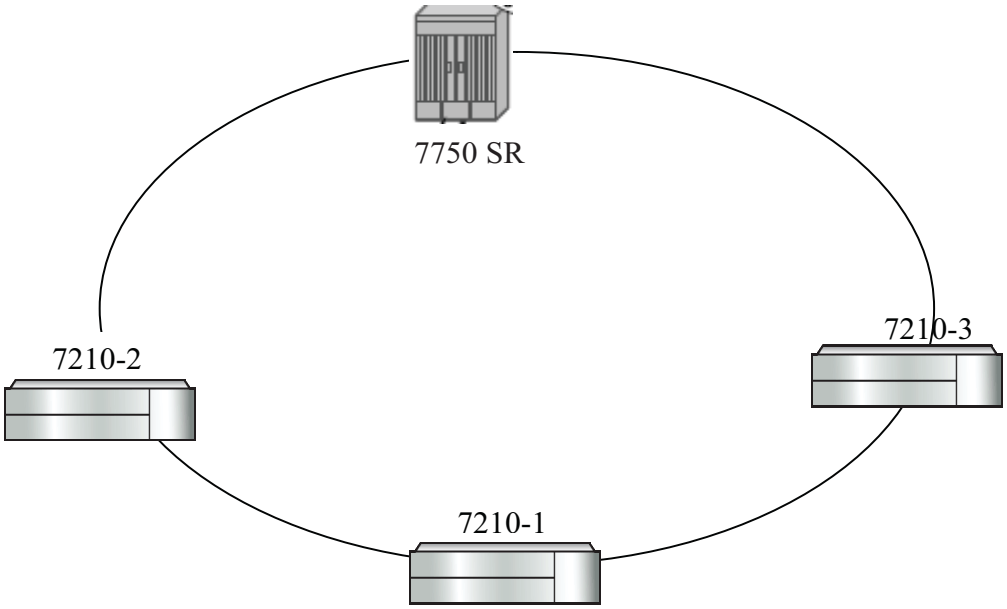

Figure 41: Default QinQ SAP for Transit Traffic in a Ring Scenario

In the Figure 41, 7210-1 is used to deliver some services to customers connected to the device and additionally it needs to pass through transit from other nodes on the ring (example – traffic from 7210-2 to 7210-3 OR from 7210-2 to 7750 –SR onto the core network).

Without Default QinQ SAPs, user would need to configure a service on 7210-1, with access-uplink SAPs for each service originating on some other node in the ring. With support for Default QinQ SAPs, all traffic which does not need to be delivered to any customer service configured on 7210-1 can be switched using the EPIPE service. The example shown below provides the sample configuration commands in this scenario:

```
ALA-1>config>service# epipe 8 customer 1 svc-sap-type null-star create
            sap 1/1/5:*.* create
                statistics
                    ingress
                    exit
                exit
            exit
            sap 1/1/6:*.* create
                statistics
                    ingress
                    exit
                exit
            exit
            no shutdown
        exit
```

Service Management Tasks

This section discusses the following Epipe service management tasks:

- Modifying Epipe Service Parameters on page 219
- Disabling an Epipe Service on page 219
- Re-Enabling an Epipe Service on page 220
- Deleting an Epipe Service on page 220

Modifying Epipe Service Parameters

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

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

The following output displays the SAP configuration:

```
ALA-1>config>service# info
----------------------------------------------
     epipe 2 customer 6 vpn 2 create
             description "Distributed Epipe service to east coast"
             sap 1/1/3:21 create
               accounting-policy 14
             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.
2. Delete the SAP.
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
```

Example:
```
config>service# epipe 2
config>service>epipe# sap 1/1/3:21
config>service>epipe>sap# shutdown
config>service>epipe>sap# exit
config>service>epipe# no sap 1/1/3:21
config>service>epipe# epipe 2
config>service>epipe# shutdown
config>service>epipe# exit
config>service# no epipe 2
```

VLL Services Command Reference

Epipe Service Configuration Commands

- Epipe Service Configuration Commands on page 221
  - Epipe Global Commands for 7210 SAS-E on page 221
  - Epipe Global Commands for 7210 SAS-D on page 222
  - Epipe Global Commands for 7210 SAS-K 2F2T1C on page 222
  - Epipe Global Commands for 7210 SAS-K 2F4T6C on page 222
  - Epipe SAP Configuration Commands for 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C on page 224
  - Epipe SAP Meter Override Commands on page 225
  - Epipe SAP Statistics commands for 7210 SAS-E and 7210 SAS-D on page 225
  - Epipe SAP Configuration- QoS and Filter command for 7210 SAS-D and SAS-E on page 227
  - Epipe SAP Configuration- QoS and Filter command for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C on page 228
- Show Commands on page 229
- Clear Commands on page 229

Epipe Global Commands for 7210 SAS-E

config
— **service**
— **epipe** *service-id* **[customer** *customer-id***] [create] [svc-sap-type {null-star|dot1q|dot1q-pre-serve}] [customer-vid** *vlan-id***]**
— **no epipe** *service-id*
— **description** *description-string*
— **no description**
— **sap** *sap-id* [**create**]
— **no sap** *sap-id*
— **service-name** *service-name*
— **no service-name**
— [**no**] **shutdown**

Epipe Global Commands for 7210 SAS-D

```
config
    — service
        — epipe service-id [customer customer-id] [create] [svc-sap-type {null-star|dot1q-pre-
          serve|any|dot1q-range}] [customer-vid vlan-id]
        — no epipe service-id
            — description description-string
            — no description
            — [no] endpoint endpoint-name [create]
            — sap sap-id [create]
            — no sap sap-id
            — service-name service-name
            — no service-name
            — [no] shutdown
```

Epipe Global Commands for 7210 SAS-K 2F2T1C

```
config
    — service
        — epipe service-id [customer customer-id] [create]
        — epipe service-id [customer customer-id] [create] [svc-sap-type {null-star|dot1q|dot1q-pre-
          serve|any|dot1q-range|qinq-inner-tag-preserve}]
        — epipe service-id [customer customer-id] [create] [svc-sap-type {null-star|dot1q|dot1q-pre-
          serve|any|dot1q-range|qinq-inner-tag-preserve}]
        — no epipe service-id
            — description description-string
            — no description
            — sap sap-id [create]
            — no sap sap-id
            — service-mtu octets
            — no service-mtu
            — service-name service-name
            — no service-name
            — [no] shutdown
```

Epipe Global Commands for 7210 SAS-K 2F4T6C

```
config
    — service
        — epipe service-id [customer customer-id] [create]
        — epipe service-id [customer customer-id] [create] [svc-sap-type {null-star|dot1q|dot1q-pre-
          serve|any|dot1q-range|qinq-inner-tag-preserve}]
        — epipe service-id [customer customer-id] [create] [svc-sap-type {null-star|dot1q|dot1q-pre-
          serve|any|dot1q-range|qinq-inner-tag-preserve}]
        — no epipe service-id
            — description description-string
            — no description
            — [no] endpoint endpoint-name [create]
                — active-hold-delay active-endpoint-delay
                — no active-hold-delay
                — revert-time [revert-time | infinite]
                — no revert-time
```

— **standby-signaling-master**
— [**no**] **standby-signaling-master**
— **sap** *sap-id* [**create**]
— **no sap** *sap-id*
— **service-mtu** *octets*
— **no service-mtu**
— **service-name** *service-name*
— **no service-name**
— [**no**] **shutdown**
— **[no] spoke-sdp**

```
config
    — service
        — epipe service-id
        — no epipe service-id
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy acct-policy-id
                — [no] collect-stats
                — description description-string
                — no description
                — eth-cfm
                    — [no] mep mep-id domain md-index association ma-index
                      [direction {up | down}] primary-vlan-enable
                        — [no] ais-enable
                            — [no] client-meg-level [[level [level ...]]
                            — [no] interval {1 | 60}
                            — [no] priority priority-value
                            — no send-ais-on-port-down
                            — send-ais-on-port-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
                — mac-swap-enable
                — no mac-swap-enable
                — ethernet
                    — [no] ais-enable
                — [no] ignore-oper-down
                — [no] shutdown
```

Epipe SAP Meter Override Commands

```
config
    — service
        — epipe service-id
        — no epipe service-id
            — no sap sap-id
                — ingress
                    — meter-override
                        — meter meter-id [create]
                        — no meter meter-id
                        — adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
                        — cbs size-in-kbytes (supported on 7210 SAS-E)
                        — no cbs
                        — cbs size [kbits|bytes|kbytes] (supported on 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C)
                        — no cbs
                        — mbs size-in-kbits (supported on 7210 SAS-E)
                        — no mbs
                        — mbs size [kbits|bytes|kbytes] (supported on 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C)
                        — no mbs
                        — mode mode
                        — no mode
                        — rate cir cir-rate [pir pir-rate]
                        — no rate
```

Epipe SAP Statistics commands for 7210 SAS-E and 7210 SAS-D

```
config
    — service
        — epipe service-id
        — no epipe service-id
            — no sap sap-id
                — statistics
                    — egress
                        — [no] forwarded-count (supported only on 7210 SAS-D)
                        — [no] packets-forwarded-count (supported only on 7210 SAS-E)
                    — ingress
                        — counter-mode {packet | octet} {in-out-profile-count|forward-drop-count}(supported only on 7210 SAS-E)
                        — [no] shutdown
```

— **counter-mode** **{in-out-profile-count|forward-drop-count}**(supported only on 7210 SAS-D)
— **[no]** **received-count** (supported only on 7210 SAS-D)

Epipe Spoke SDP Configuration Commands for 7210 SAS-K 2F4T6C

config
— **service]**
— **epipe** ***service-id***
— **no 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**
— **no spoke-sdp** *sdp-id*[:*vc-id*]
— **accounting-policy** *acct-policy-id*
— **no accounting-policy**
— [**no**] **collect-stats**
— [**no**] **control-word**
— [**no**] **description**
— [**no**] **egress**
— [**no**] **vc-label** *egress-vc-label*
— **eth-cfm**
— [**no**] **ais-enable**
— [**no**] **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {up | down}]
— [**no**] **ais-enable**
— [**no**] **client-meg-level** [[**level** [level ...]]
— [**no**] **interval** {**1** | **60**}
— [**no**] **priority** *priority-value*
— [**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**
— **mip** **[mac** *mac address***]**
— **mip default-mac**
— **no mip**
— [**no**] **force-vlan-vc-forwarding**
— [**no**] **ingress**
— [**no**] **vc-label** *egress-vc-label*
— **precedence** [*precedence-value*| **primary**]
— **no precedence**
— **no pw-status-signaling**
— **pw-status-signaling**
— [**no**] **shutdown**

— **vlan-vc-tag** *0..4094*
— **no vlan-vc-tag** [*0..4094*]

Epipe SAP Configuration- QoS and Filter command for 7210 SAS-D and SAS-E

config
— **service**
— **epipe** *service-id* [**customer** *customer-id*] [**create**] [**svc-sap-type {null-star|dot1q-pre-serve|any|dot1q-range}**] [**customer-vid** *vlan-id*] (for 7210 SAS-D)
— **epipe** *service-id* [**customer** *customer-id*] [**create**] [**svc-sap-type {null-star|dot1q-pre-serve|dot1q}**] [**customer-vid** *vlan-id*] (for 7210 SAS-E)
— **no epipe** *service-id*
— **no sap** *sap-id* **[create]**
— **egress**
— **aggregate-meter-rate** *rate-in-kbps* [**burst** *burst-in-kbits*] [**enable-stats**]
— **no aggregate-meter-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*]
— **ingress**
— **aggregate-meter-rate** *rate-in-kbps* [**burst** *burst-in-kbits*]
— **no aggregate-meter-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*]
— **qos** *policy-id*
— **no qos**

Epipe SAP Configuration- QoS and Filter command for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

```
config
    — service
        — epipe service-id [customer customer-id] [create] [svc-sap-type {any| dot1q-range}]
        — no epipe service-id
            — no sap sap-id [create]
                — egress
                    — agg-shaper-rate agg-rate
                    — no agg-shaper-rate
                    — dot1p-inner dot1p-inner
                    — no dot1p-inner
                    — no dot1p-outer
                    — dot1p-outer dot1p-outer
                    — filter [ip ip-filter-id]
                    — filter [ipv6 ipv6 -filter-id]
                    — filter [mac mac-filter-id] (app
                    — no filter [ip ip-filter-id] [ ipv6 ipv6 -filter-id] [mac mac-filter-id]
                    — qos policy-id
                    — no qos
                — ingress
                    — agg-shaper-rate agg-rate
                    — no agg-shaper-rate
                    — aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
                    — no aggregate-meter-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]
                    — qos policy-id
                    — no qos
```

Connection Profile Commands

```
config
    — connection-profile conn-prof-id [create]
    — no connection-profile conn-prof-id
        — description description-string
        — no description
        — ethernet
            — no ranges
        — ranges vlan ranges [vlan ranges...(upto 32 max)]
```

Show Commands

```
show
    — service
        — id service-id
            — all
            — base
            — sap
            — stp [sap-id] [detail]]
        — sap-using [sap sap-id]
        — sap-using [ingress | egress] filter filter-id
        — sap-using [ingress ] qos-policy qos-policy-id
        — service-using [epipe] [vpls] [mirror] [cpipe]  [i-vpls] [m-vpls] [sdp sdp-id] [customer customer-id]
show
    — connection-profile [conn-prof-id] [associations]
```

Clear Commands

```
clear
    — service
        — id service-id
        — statistics
            — id service-id
            — counters
            — sap sap-id {all | counters | stp| l2pt}
```

VLL Service Configuration Commands

- Generic Commands on page 232
- VLL Global Commands on page 234
- VLL SAP Commands on page 239
- Connection Profile Commands on page 247
- Service Filter and QoS Policy Commands on page 249

Generic Commands

shutdown

Syntax [**no**] **shutdown**
config>service>epipe
config>service>epipe>sap
config>service>epipe>sap>eth-cfm>mep

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

description

Syntax **description** *description-string*
no description

Context config>service>epipe
config>service>epipe>sap
config>service>epipe>spoke-sdp
config>connection-profile

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command creates a text description stored in the configuration file for a configuration context. The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

eth-cfm

Syntax **eth-cfm**

Context config>service>vll>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command enables the context to configure ETH-CFM parameters.

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {up | down}] **primary-vlan-enable**
no mep *mep-id* **domain** *md-index* **association** *ma-index*

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

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command provisions the maintenance endpoint (MEP).

The no form of the command reverts to the default values.

Note: For more information on ETH-CFM support for different services listed below:

- Table 9, "ETH-CFM Support Matrix for 7210 SAS-D," on page 94
- Table 10, "ETH-CFM Support Matrix for 7210 SAS-E," on page 95
- Table 11, "ETH-CFM Support Matrix for 7210 SAS-K 2F2T1C," on page 95
- Table 12, "ETH-CFM Support Matrix for 7210 SAS-K 2F4T6C," on page 96

Parameters *mep-id* — Specifies the maintenance association end point identifier.

Values 1 — 8191

md-index — Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 4294967295

direction up| down — Indicates the direction in which the maintenance association (MEP) faces on the bridge port.

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

up — Sends ETH-CFM messages towards the MAC relay entity. This value cannot be used on 7210 SAS-E devices.

primary-vlan-enable — **Platforms supported:** 7210 SAS-K 2F2T1C.
Provides a method for linking the with the primary VLAN configured under the bridge-identifier for the MA. This must be configured as part of the creation step and can only be changed by deleting the MEP and recreating it. Primary VLANs are only supported under Ethernet SAPs.

VLL Global Commands

epipe

Syntax

epipe *service-id* [customer *customer-id*] **[create] [svc-sap-type {null-star|dot1q|dot1q-preserve}] [customer-vid** *vlan-id*]
(for 7210 SAS-E)
epipe *service-id* [**customer** *customer-id*] [create] [**svc-sap-type** {null-star|dot1q-preserve|any|dot1q-range}] [**customer-vid** *vlan-id*] **(for 7210 SAS-D)**
epipe *service-id* [**customer** *customer-id*] [create] [**svc-sap-type** {**any|dot1q-range**}] **(for 7210 SAS-K 2F2T1C**)
epipe *service-id* [**customer** *customer-id*] [create] [**svc-sap-type {null-star|dot1q|dot1q-preserve|any|dot1q-range|qinq-inner-tag-preserve}**] **(for 7210 SAS-K 2F4T6C**)
no epipe *service-id*

Context config>service

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

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). In a local service, the SAPs may be defined in one 7210 SAS node and in distributed service the SAPs may be defined on two different 7210 SAS nodes.

NOTE:

- 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K5 platforms support only local SAP to SAP service.
- 7210 SAS-K12 platforms support both local and distributed service.

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.

NOTE: The table lists the allowed SAPs for a particular value of svc-sap-type on different 7210 platforms.

- Table 4, “SAP and Service Combinations for 7210 SAS-E,” on page 46,
- Table 5, “SAP and Service Combinations for 7210 SAS-D,” on page 47,
- Table 6, “SAP and Service Combinations for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C,” on page 49

The **no** form of this command deletes the epipe service instance with the specified *service-id*. The service cannot be deleted until the service has been shutdown.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values *service-id*: 1 — 2147483648
svc-name: 64 characters maximum

customer *customer-id* **—** Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 — 2147483647

svc-sap-type — This parameter specifies the type of access SAPs and access-uplink SAPs allowed in the service.

null-star - Specifies that the allowed SAP in the service, which can be null SAPs, Dot1q default, Q.* SAP, 0.* SAP or Default QinQ SAP (also known as *.* SAP). Supported on 7210 SAS-D and 7210 SAS-E.

dot1q - Specifies that the allowed SAP in the service are Dot1q SAPs and dot1q explicit null SAPs. This is supported only on 7210 SAS-E.

dot1q-preserve - Specifies that the allowed SAP in the service are Dot1q. The Dot1q ID is not stripped after packets matches the SAP. Supported on 7210 SAS-D and 7210 SAS-E.
dot1q-range - Specifies that the access SAP in the service can use VLAN ranges as the SAP tags. The VLAN ranges are configured using the configure> connection-profile CLI command. On ingress of the access dot1q SAP using VLAN ranges, the outermost tag is not removed before forwarding. This is supported only for 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

any - Allows any SAP type. This is supported only on 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Default "any" - 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C
"null-star" - 7210 SAS-E

customer-vid *vlan-id* **—** Defines the dot1q VLAN ID to be specified while creating the local Dot1q SAP for **svc-sap-type dot1q-preserve**. Applicable only for Access-uplink mode.

Values 1 — 4094

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

endpoint

Syntax [**no**] **endpoint** *endpoint-name*
config>service>epipe

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures a service endpoint.

Parameters *endpoint-name* — Specifies an endpoint name.

active-hold-delay

Syntax **active-hold-delay** *active-hold-delay*
no active-hold-delay

Context config>service>epipe>endpoint

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command specifies that the node will delay sending the change in the T-LDP status bits for the VLL endpoint when the MC-LAG transitions the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby** or when any object in the endpoint. For example, SAP, ICB, or regular spoke SDP, transitions from up to down operational state.

By default, when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby**, the node sends immediately new T-LDP status bits indicating the new value of "standby" over the spoke SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

There is no delay applied to the VLL endpoint status bit advertisement when the MC-LAG transitions the LAG subgroup which hosts the SAP from "standby" to "active" or when any object in the endpoint transitions to an operationally up state.

Default 0 — A value of zero means that when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby**, the node sends immediately new T-LDP status bits indicating the new value of **standby** over the spoke SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

1 — A value of one 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** [*revert-time* | **infinite**]

no revert-time

Context config>service>epipe>endpoint

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures the time to wait before reverting back to the primary spoke SDP defined on this service endpoint, after having failed over to a backup spoke SDP.

Parameters *revert-time* — Specify the time, in seconds, to wait before reverting to the primary SDP.

Values 0 — 600

infinite — Causes the endpoint to be non-revertive.

standby-signaling-master

Syntax [**no**] **standby-signaling-master**

Context config>service>vll>endpoint

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

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 no standby-signaling-master

service-mtu

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

Context config>service>epipe

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

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 (i.e., 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.

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.

The **no** form of this command returns the default **service-mtu** for the indicated service type to the default value.

Default epipe: 1514

The following table displays MTU values for specific VC types.

VC-Type	Example Service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
VPLS	1514	1500
VPLS (with preserved dot1q)	1518	1504
VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (QinQ with preserved bottom Qtag)	1518	1504

The size of the MTU in octets, expressed as a decimal integer.

Values 1 — 9194

service-name

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

Context config>service>epipe

Platforms Supported: 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

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 o 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).

VLL SAP Commands

sap

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

Context config>service>epipe

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

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 7210 device. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

On 7210 SAS-E and 7210 SAS-D, in a single physical port only one SAP can belong to one service. Multiple SAPs can be defined over a physical port but each of these SAPs should belong to different service. This restriction does not apply to 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

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. Additionally, in access-uplink mode, SAPs can be defined also on access-uplink port. Access-uplink SAPs are network facing SAPs representing Dot1q or QinQ tunnels used to transport traffic towards the service nodes.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded.

The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The following encapsulations are supported:

- Ethernet access SAPs support null, dot1q, and qinq (QinQ SAP on access ports is not supported on 7210 SAS-E).
- Ethernet access-uplink SAPs support only QinQ encapsulation.

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 (IES), the IP interface must be shutdown before the SAP on that interface may be removed.

Default No SAPs are defined.

Special Cases 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).

sap-id — Specifies the physical port identifier portion of the SAP. See Common CLI Command Descriptions on page 945 for command syntax.

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

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy [<acct-policy-id>]

Context config>service>epipe>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command creates the accounting policy context that can be applied to a SAP.

An accounting policy must be defined before it can be associated with a SAP. If the *policy-id* does not exist, an error message is generated.

A maximum of one accounting policy can be associated with a SAP at one time. Accounting policies are configured in the **config>log** context.

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

Default Default accounting policy.

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1-99

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>epipe>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command enables accounting and statistical data collection for either the SAP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default no collect-stats

ethernet

Syntax **ethernet**

Context config>service>epipe>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Use this command to configure Ethernet properties in this SAP.

ais-enable

Syntax [**no**] **ais-enable**

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

Description This command enables the generation and the reception of AIS messages.

client-meg-level

Syntax **client-meg-level** [[*level* [*level* ...]]
no client-meg-level

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

Description This command configures the client maintenance entity group (MEG) level(s) to use for AIS message generation. Up to 7 levels can be provisioned with the restriction that the client MEG level must be higher than the local MEG level.

Parameters *level* — Specifies the client MEG level.

Values 1 — 7

Default 1

interval

Syntax **interval** {**1** | **60**}
no interval

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

Description This command specifies the transmission interval of AIS messages in seconds.

Parameters **1** | **60** — The transmission interval of AIS messages in seconds.

Default 1

priority

Syntax **priority** *priority-value*
no priority

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

Description This command specifies the priority of AIS messages originated by the node.

Parameters *priority-value* — 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>epipe>sap>eth-cfm>mep

Description This command enables the generation of CCM messages.

The **no** form of the command disables the generation of CCM messages.

ccm-ltm-priority

Syntax **ccm-ltm-priority** *priority*
no ccm-ltm-priority

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

Description This command specifies the priority value for CCMs and LTMs transmitted by the MEP.

The **no** form of the command removes the priority value from the configuration.

Default The highest priority on the bridge-port.

Parameters *priority* — Specifies the priority of CCM and LTM messages.

Values 0 — 7

llf

Syntax [**no**] **llf**

Context config>service>epipe>sap>ethernet

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command enables Link Loss Forwarding (LLF) on an Ethernet port or an ATM port. This feature provides an end-to-end OAM fault notification for Ethernet VLL service. It brings down the Ethernet

port (Ethernet LLF) towards the attached CE when there is a local fault on the Pseudowire or service, or a remote fault on the SAP or pseudowire, signaled with label withdrawal or T-LDP status bits. It ceases when the fault disappears.

The Ethernet port must be configured for null encapsulation.

ignore-oper-down

Syntax **[no] ignore-oper-down**

Context config>service>epipe>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command enables the user to configure the optional command for a specific SAP to ignore the transition of the operational state to down when a SAP fails. Only a single SAP in an ePipe may use this option.

Default no ignore-oper-down

send-ais-on-port-down

Syntax **[no] send-ais-on-port-down**

Context config>service>epipe>sap>eth-cfm>mep>ais-enable
config>service>vpls>sap>eth-cfm>mep>ais-enable

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

By default, the system generates an ETH-AIS message, if enabled, when CCM messages are not received within the configured time period. This command allows the user to specify that ETH-AIS should be generated for client MEPs immediately when port down event is detected on the port where the server MEP (and the associated SAP) resides. On a subsequent port up event, the AIS messages continue to be sent until valid CCMs are received. If there are no remote-meps configured for the MEP then on a subsequent port up event, the AIS messages are not sent.

The no form of the command reverts to default behavior.

Default no send-ais-on-port-down

eth-test-enable

Syntax [**no**] **eth-test-enable**

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

Description For ETH-test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands:

oam eth-cfm eth-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [data-length *data-length*]

A check is done for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP will indicate the problem.

test-pattern

Syntax **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
no test-pattern

Context config>service>epipe>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.

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** *errors*
no bit-error-threshold

Context config>service>epipe>sap>eth-cfm>mep>eth-test-enable

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command is used to specify the threshold value of bit errors.

one-way-delay-threshold

Syntax **one-way-delay-threshold** *seconds*

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

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command enables/disables eth-test functionality on MEP.

Parameters *seconds* — Specifies the one way delay threshold in seconds.

Values 0-600

Default 3

mac-swap-enable

Syntax [**no**] mac-swap-enable

Context config>service>epipe>sap
config>service>vpls>sap

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

When this command is enabled, the source and destination MAC address for all the packets being sent out of the SAP is swapped and then the packets are looped back to ingress of the SAP. SAP loopback is typically used with Y.1564 testhead functionality to loopback packets received on the end which is acting as a reflector looping back packets received from the testhead generator.

This command is available for testing VLL services and VPLS services only. When enabled, no packets egress the service SAP and packets received on ingress are not processed.

NOTES:

- Before enabling this command, turn off all layer 2 and IP control protocols (such as LACP, EFM, 802.1x and so on) on the device and its peer to prevent errors such as protocol flaps due to time-out and so on.
- When SAP loopback with mac-swap is enabled, for broadcast and multicast packets (that is, packets received with multicast or broadcast destination address), the source MAC address is used as the destination MAC address and the system MAC address is the source MAC address.
- The user can enable MAC swap on Multiple SAPs in the same service. The user can also configure static MAC on SAPs and traffic to prevent running into the MAC Move, re-learning issues, and unpredictable behavior of the device.

Default no mac-swap-enable

mip

Syntax **mip [mac mac-address]**
mip default-mac
no mip

Context config>service>epipe>sap>eth-cfm
config>service>epipe>spoke-sdp>eth-cfm - only 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command allows Maintenance Intermediate Points (MIPs) to be created if mhf-creation for the MA is configured using the default option.

NOTE: This command is not supported in 7210 SAS-E devices.

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

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MIP. The MAC must be unicast. Using the all zeros address is equivalent to the no form of this command.

default-mac — Using the no command deletes the MIP. If the operator wants to change the mac back to the default mac without having to delete the MIP and reconfiguring this command is useful.

Default no mip

fault-propagation-enable

Syntax **fault-propagation-enable {use-if-tlv | suspend-ccm}**
no fault-propagation-enable

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

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

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.

Connection Profile Commands

connection-profile

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

Context config

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command creates a profile for the user to configure the list VLAN values to be assigned to an Dot1q SAP in an Epipe service.

A connection profile can only be applied to a Dot1q SAP which is part of an Epipe Service.

The no form of this command deletes the profile from the configuration.

Default none

Parameters *conn-prof-id* — Specifies the profile number.

Values 1 — 1000 - 7210 SAS-D
1 — 8000 - 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

ethernet

Syntax ethernet

Context config>connprof

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Provides the context to configure the VLAN ranges values.

Default none

ranges

Syntax no ranges
ranges vlan-ranges *[vlan-ranges...(upto 32 max)]*

Context config>connprof>ethernet

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Specifies the list of VLAN ranges or individual VLAN ID to be used for mapping the given VLANs to the Epipe SAP.

The system validates that the values specified are valid VLAN ID in the range 0-4094 (VLAN ID 4095 is reserved). Ranges are specified in the format ‘a-b ’, the expression (a < b) should be true. Up

to about 32 individual VLAN values or VLAN ranges can be specified. A maximum of up to 8 VLAN ranges are allowed per connection profile.

Default none

Parameters *vlan-ranges* — Specifies the list of VLAN ranges or individual VLAN ID to be used for mapping the given VLANs to the Epipe SAP.

A list of space separated values specified as either a-b or individual VLAN IDs. Both the VLAN IDs and the value used for 'a' and 'b' must be in the range of 0-4094. Additionally, value 'a' must be less than value 'b'.
For example:

ranges	100-200 5 6 4000-4020
ranges	4 5 6 10 11 12
ranges	250-350 500-600 1000-1023

Service Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>epipe>sap

Description This command enables the context to configure egress SAP parameters.

force-vlan-vc-forwarding

Syntax [**no**] **force-vlan-vc-forwarding**

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

Description **Platform Supported:** 7210 SAS-K 2F4T6C.

This command forces vc-vlan-type forwarding in the data path for spoke which have either vc-type. This comand 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>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

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.

agg-shaper-rate

Syntax **agg-shaper-rate** *agg-rate*
no agg-shaper-rate

Context config>service>epipe>sap>ingress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command allows user to specify the aggregate rate for the access SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues.
The no form of the command disables the use of SAP aggregate rate shaper. In other words, the SAP can use up the maximum bandwidth available.

Default no agg-shaper-rate

Parameters ***agg-rate*** **—** Specifies the rate in kilobits per second.

Values 50 — 3000000 | max

Default max

agg-shaper-rate

Syntax **agg-shaper-rate** *agg-rate*
no agg-shaper-rate

Context config>service>epipe>sap>egress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.
This command allows user to specify the aggregate rate for the access SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use egress queues.
The no form of the command disables the use of SAP aggregate rate shaper. In other words, the SAP can use up the maximum bandwidth available.

Default no agg-shaper-rate

Parameters ***agg-rate*** **—** Specifies the rate in kilobits per second.

Values 50 — 1000000 | max

Default max

aggregate-meter-rate

Syntax **aggregate-meter-rate** *rate-in-kbps* [**burst** *burst-in-kbits*] [**enable-stats**]
no aggregate-meter-rate

Context config>service>epipe>sap>egress
config>service>ies>if>sap>egress
config>service>vpls>sap>egress

Description **Platforms Supported:** 7210 SAS-D.

This command allows the user to configure the access SAP egress aggregate policer. The rate (PIR) of the SAP egress aggregate policer must be specified by the user. The user can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the traffic sent out of the SAP and determines the final disposition of the packet, which is either forwarded or dropped.

The user can optionally associate a set of two counters to count total forwarded packets and octets and total dropped packets and octets. When use of this counter is enabled, the amount of resources required increases by twice the amount of resources taken up when the counter is not used. If the **enable-stats** keyword is specified during the creation of the meter, the counter is allocated by software (if available). To free up the counter and relinquish its use, the user can use the **no aggregate-meter-rate** command and then recreate the meter using the **aggregate-meter rate** command.

If egress frame-based accounting is used, the SAP egress aggregate meter rate accounts for the Ethernet frame overhead. The system accounts for 12 bytes of IFG and 8 bytes of start delimiter. Frame-based accounting does not affect the count of octets maintained by the counter (if in use).

Notes:

- Before enabling this command for a SAP, resources must be allocated to this feature from the egress-internal-tcam resource pool using the command **config>system>resource-profile> egress-internal-tcam>egress-sap-aggregate-meter.** Refer to the 7210 SAS Basic System Guide for more information.
- The egress aggregate meter is not FC-aware. The forward and drop decisions are taken based on the order in which the packets are sent out of the SAP by the egress port scheduler.

The **no** form of the command removes the egress aggregate policer from use.

Default no aggregate-meter-rate

Parameters *rate-in-kbps* — Specifies the rate in kilobits per second.

Values 1 — 4000000 | max

Default max

burst-in-kbits — Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 —16384

Values 512

enable-stats — Specifies if the counter to count forwarded and dropped packets and octets is allocated. If this keyword is used while configuring the meter, the counter is allocated.

aggregate-meter-rate

Syntax **aggregate-meter-rate** *rate-in-kbps* [**burst** *burst-in-kbits*]
no aggregate-meter-rate

Context config>service>vpls>sap>ingress
config>service>epipe>sap>ingress
config>service>ies>if>sap>ingress
config>service>vprn>if>sap>ingress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C; IES and VPRN contexts apply only to the 7210 SAS-K 2F4T6C.

This command allows the user to configure the access SAP ingress aggregate policer. The rate of the SAP ingress aggregate policer must be specified by the user. The user can optionally specify the burst

size for the SAP aggregate policer. The aggregate policer monitors the ingress traffic on different FCs using policers to rate-limit the flow and determines the final disposition of the packet. The packet is either forwarded to an identified profile or dropped.

Note:

- The sum of CIR of the individual FCs configured under the SAP cannot exceed the PIR rate configured for the SAP. Though the 7210 SAS software does not block this configuration, it is not recommended for use.
- The queued traffic flows are not limited by the aggregate-meter. In others words, only metered flows can use aggregate meter. Queues flows can use only aggregate shaper.

The table below provides information about the final disposition of the packet based on the operating rate of the per FC policer and the per SAP aggregate policer:

Per FC meter Operating Rate	Per FC Assigned Color	SAP aggregate meter Operating Rate	Final Packet Color
Within CIR	Green	Within PIR	Green or In-profile
Within CIR	Green	Above PIR	Red and Dropped
Above CIR, Within PIR	Yellow	Within PIR	Yellow or Out-of-Profile
Above CIR, Within PIR	Yellow	Above PIR	Red or Dropped
Above PIR	Red	Within PIR	Red or Dropped
Above PIR	Red	Above PIR	Red or Dropped

Table 13: Final Disposition of the packet based on per FC and per SAP policer or meter.

The SAP ingress meter counters increment the packet or octet counts based on the final disposition of the packet.

The **no** form of the command removes the aggregate policer from use.

Default no aggregate-meter-rate

Parameters *rate-in-kbps* — Specifies the rate in kilobits per second.

Values 1 — 4000000 | max

Default max

burst-in-kbits — Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 —16384

Default 512

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

Description **Platforms Supported:** 7210 SAS-E, 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

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.

NOTE: For filter support available on different 7210 platforms, see the 7210 SAS Router Configuration User Guide.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system.

Special Cases **Epipe —** Both MAC and IP filters are supported on an Epipe service SAP.

Parameters **ip** *ip-filter-id* **—** Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 — 65535

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

dot1p-inner

Syntax **no dot1p-inner**
dot1p-inner

Context config>service>epipe>sap>egress
config>service>vpls>sap>egress

Description **Platforms supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command allows the user to define the Dot1p marking values to be used per SAP on egress for the inner tag when the SAP encapsulation is QinQ (that is, Q1.Q2 SAP)r. The command takes effect only if remarking is enabled in the remark policy associated with this SAP (under the egress context). It overrides the marking values defined in the remark policy associated with this SAP, if any.

Table 14: Table for Dot1p values extracted from the packet on SAP (ingress) when Dot1p-inner and Dot1p-outer CLI command are configured:

Ingress SAP Type	Dot1p value extracted from received packet when Use-rcvd-inner-dot1p	Dot1p values extracted from received packet when Use-rcvd-outer-dot1p
Null SAP	None	None
Dot1q SAP	Dot1p from the outermost VLAN tag	Dot1p from the outermost VLAN tag
Dot1q Default SAP (that is, * SAP)	None	None
Dot1q Explicit NULL SAP (that is,:0 SAP)	Dot1p from the outermost VLAN tag (if priority tagged packet), else none.	Dot1p from the outermost VLAN tag (if priority tagged packet), else none.
Dot1q range SAP	Dot1p from the outermost VLAN tag	Dot1p from the outermost VLAN tag
Q1.Q2 SAP	Dot1p of the inner tag	Dot1p from the outermost VLAN tag
Q1.* SAP	Dot1p from the outermost VLAN tag	Dot1p from the outermost VLAN tag
0.* SAP	Dot1p from the outermost VLAN tag	Dot1p from the outermost VLAN tag

Q1.0 SAP	Dot1p of the inner priority tag, if available, else from outermost VLAN tag	Dot1p from the outermost VLAN tag
QinQ Default SAP (that is,. *.* SAP)	None	None

Table 15: Table for Dot1p values marked in the packet on SAP egress when Dot1p inner and Dot1p-outer CLI command are configured:

Egress SAP Type	Dot1p-inner - Dot1p value marked in the packet sent out of this SAP when Use-rcvd-inner-dot1p	Dot1p-inner Dot1p value marked in the packet sent out of this SAP when Use-rcvd-outer-dot1p	Dot1p-outer Dot1p value marked in the packet sent out of this SAP when Use-rcvd-inner-dot1p	Dot1p-outer Dot1p value marked in the packet sent out of this SAP when Use-rcvd-outer-dot1p
Null SAP	NA	NA	NA	NA
Dot1q SAP	NA	NA	Dot1p bits in the outermost tag	Dot1p bits in the outermost tag
Dot1q Default SAP (that is, * SAP)	NA	NA	NA	NA
Dot1q Explicit NULL SAP (that is, :0 SAP)	NA	NA	NA	NA
Dot1q range SAP	NA	NA	NA	NA
Q1.Q2 SAP	Dot1p bits from the inner tag	Dot1p bits from the outermost tag	Dot1p bits from the inner tag	Dot1p bits from the outermost tag
Q1.* SAP	NA	NA	Dot1p bits from the inner tag	Dot1p bits from the outermost tag
0.* SAP	NA	NA	NA	NA
Q1.0 SAP	NA	NA	Dot1p bits from the inner tag	Dot1p bits from the outermost tag
QinQ Default SAP (that is, *.* SAP)	NA	NA	NA	NA

NOTE: **NA** – egress encapsulation is not done, neither remark policy nor 'use-rcvd' command will be applicable at that level.

If the no form of the command is executed, software will use the values defined in the remark policy associated with this SAP, if any. If no remark policy is associated with SAP egress, then the default values are used.

Default no dot1p-inner

Parameters *use-rcvd-inner-dot1p* — For more information, see :

- Table 14, "Table for Dot1p values extracted from the packet on SAP (ingress) when Dot1p-inner and Dot1p-outer CLI command are configured:," on page 254
- Table 15, "Table for Dot1p values marked in the packet on SAP egress when Dot1p inner and Dot1p-outer CLI command are configured:," on page 255

use-rcvd-outer-dot1p — For more information, see

- Table 14, "Table for Dot1p values extracted from the packet on SAP (ingress) when Dot1p-inner and Dot1p-outer CLI command are configured:," on page 254
- Table 15, "Table for Dot1p values marked in the packet on SAP egress when Dot1p inner and Dot1p-outer CLI command are configured:," on page 255

dot1p-outer

Syntax **no dot1p-outer**
dot1p-outer

Context config>service>epipe>sap>egress
config>service>vpls>sap>egress
config>service>ies>if>sap>egress
config>service>vprn>if>sap>egress

Description **Platforms supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C; IES and VPRN contexts apply only to the 7210 SAS-K 2F4T6C.

This command allows the user to define the Dot1p marking values to be used per SAP on egress for the outer tag when the SAP encapsulation is QinQ or Dot1q. The command takes effect only if remarking is enabled in the remark policy associated with this SAP (under the egress context). It overrides the marking values defined in the remark policy associated with this SAP, if any.

For more information, see:

- Table 14, "Table for Dot1p values extracted from the packet on SAP (ingress) when Dot1p-inner and Dot1p-outer CLI command are configured:," on page 254
- Table 15, "Table for Dot1p values marked in the packet on SAP egress when Dot1p inner and Dot1p-outer CLI command are configured:," on page 255

If the no form of the command is executed, software will use the values defined in the remark policy associated with this SAP, if any. If no remark policy is associated with SAP egress, then the default values are used.

Default no dot1p-outer

Parameters *use-rcvd-inner-dot1p* — For more information, see:

- Table 14, "Table for Dot1p values extracted from the packet on SAP (ingress) when Dot1p-inner and Dot1p-outer CLI command are configured:," on page 254
- Table 15, "Table for Dot1p values marked in the packet on SAP egress when Dot1p inner and Dot1p-outer CLI command are configured:," on page 255

user-rcvd-outer-dot1p — For more information, see:

- Table 14, "Table for Dot1p values extracted from the packet on SAP (ingress) when Dot1p-inner and Dot1p-outer CLI command are configured:," on page 254
- Table 15, "Table for Dot1p values marked in the packet on SAP egress when Dot1p inner and Dot1p-outer CLI command are configured:," on page 255

meter-override

Syntax **[no] meter-override**

Context config>service>epipe>sap>ingress
config>service>vpls>sap>ingress
config>service>ies>interface>sap>ingress
config>service>vprn>interface>sap>ingress - supported only on 7210 SAS-K 2F4T6C

Description **Platforms supported:** 7210 SAS-E, 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command, within the SAP ingress contexts, is used to create a CLI node for specific overrides to one or more meters created on the SAP through the sap-ingress QoS policies.

The no form of the command is used to remove any existing meter overrides.

Default no meter-overrides

meter

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

Context config>service>epipe>sap>ingress>meter-override
config>service>vpls>sap>ingress>meter-override
config>service>ies>interface>sap>ingress>meter-override
config>service>vprn>interface>sap>ingress>meter-override - supported only on 7210 SAS-K 2F4T6C

Description **Platforms supported:** 7210 SAS-E, 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command, within the SAP ingress contexts, is used to create a CLI node for specific overrides to a specific meter created on the SAP through a sap-ingress QoS policies.

The no form of the command is used to remove any existing overrides for the specified meter-id.

Parameters *meter-id* — The meter-id parameter is required when executing the meter command within the meter-overrides context. The specified meter-id must exist within the sap-ingress QoS policy applied to the SAP. If the meter is not currently used by any forwarding class or forwarding type mappings, the meter will not actually exist on the SAP. This does not preclude creating an override context for the meter-id.

create — The create keyword is required when a meter *meter-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.

adaptation-rule

Syntax **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
no adaptation-rule

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter
config>service>vprn>interface>sap>ingress>meter-override>meter - supported only on 7210 SAS-K 2F4T6C

Description **Platforms supported:** 7210 SAS-E, 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command can be used to override specific attributes of the specified meter adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the meter 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 meter-override meter *meter-id* command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **meter-override** command is not specified, the default applies.

NOTE: When the meter mode in use is 'trtcm2', this parameter is interpreted as EIR value. For more information, see the description and relevant notes for meter modes in the 7210 SAS QoS user guide.

cir — The **cir** parameter defines the constraints enforced when adapting the CIR rate defined within the meter-override meter *meter-id* 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 meter, 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 meter will be equal to or less than the administrative rate specified using the **meter-override** 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 **meter-override** command.

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

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter

Description **Platforms Supported:** 7210 SAS-E.

This command, within the SAP ingress meter-overrides contexts, is used to override the sap-ingress QoS policy configured cbs parameter for the specified meter-id.

The no form of the command is used to restore the meter cbs setting to the meter defined value.

Default 32kbits

Parameters *size-in-kbytes* — The size parameter is an integer expression of the number of kilobytes reserved for the meter. 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 [4..2146959 | default]

cbs

Syntax **cbs** *size* **[kbits|bytes|kbytes]**
no cbs

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter
config>service>vprn>interface>sap>ingress>meter-override>meter - supported only on

7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command provides a mechanism to override the default CBS for the meter. The committed burst size parameter specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the CBS value then the packets are marked as in-profile by the meter to indicate that the traffic is complying meter configured parameters.

The **no** form of this command returns the CBS size to the default value.

Default 32 kbits

Parameters *size* — Specifies the value in either kbits, kilobytes or bytes.

Values kbits : [4..16384 | default]
bytes : [512..2097152]
kbytes : [1..2048]

mbs

Syntax **mbs** *size-in-kbits*
no mbs

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter

Description **Platforms Supported:** 7210 SAS-E.

This command, within the SAP ingress meter-overrides contexts, is used to override the sap-ingress QoS policy configured mbs parameter for the specified meter-id.

The no form of the command is used to restore the meter mbs setting to the meter defined value.

Default 128 size-in-kbits

Parameters *size-in-kbits* — The size parameter is an integer expression of the number of kilobits reserved for the meter. If a value of 10KBits 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 [4..2146959 | default]

mbs

Syntax **mbs** *size* **[kbits|bytes|kbytes]**
no mbs

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter

config>service>vprn>interface>sap>ingress>meter-override>meter - supported only on 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command provides a mechanism to override the default MBS for the meter. The maximum burst size parameter specifies the maximum burst size that can be transmitted by the source while still complying with the PIR. If the transmitted burst is lower than the MBS value then the packets are marked as in-profile by the meter to indicate that the traffic is complying meter configured parameters.

The **no** form of this command returns the MBS size to the default value.

Default 32Kbits

size — Specifies the value in either kbits, kilobytes or bytes.

Values kbits : [4..16384 | default] (for 7210 SAS-D)
bytes : [512..2097152]
kbytes : [1..2048]

mode

Syntax **mode** *mode*
no mode

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter
config>service>vprn>interface>sap>ingress>meter-override>meter - supported only on 7210 SAS-K 2F4T6C

Description **Platforms supported:** 7210 SAS-E, 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command within the SAP ingress meter-overrides contexts is used to override the sap-ingress QoS policy configured mode parameters for the specified meter-id.

The no mode command is used to restore the policy defined metering and profiling mode to a meter.

Parameters *mode* — Specifies the rate mode of the meter-override.

Values 7210 SAS-D- trtcm1|trtcm2|srtcm
7210 SAS-E- trtcm1|trtcm2

rate

Syntax **rate cir** *cir-rate* **[pir** *pir-rate***]**
no rate

Context config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter

config>service>vprn>interface>sap>ingress>meter-override>meter - supported only on 7210 SAS-K 2F4T6C

Description **Platforms supported:** 7210 SAS-E, 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command within the SAP ingress meter-overrides contexts is used to override the sap-ingress QoS policy configured rate parameters for the specified meter-id.

The no rate command is used to restore the policy defined metering and profiling rate to a meter.

Default **max —** The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

NOTE: When the meter mode is set to 'trtcm2' the PIR value is interpreted as the EIR value. For more information, see the 7210 SAS QoS user guide.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values [0..20000000 | max]

Default max

cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.
Fractional values are not allowed and must be given as a positive integer.

Values **[0..20000000 | max]**

Default 0

qos

Syntax **qos** *policy-id*
no qos

Context config>service>epipe>sap>ingress

Description **Platforms supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).

QoS ingress 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 ingress policies. The **qos** command only allows ingress policies to be associated on SAP ingress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type will return an error.

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

policy-id — The ingress policy ID to associate with SAP on ingress. The policy ID must already exist.

Values 1 — 65535

statistics

Syntax **statistics**

Context config>service>epipe>sap

Description **Platforms supported:** 7210 SAS-D and 7210 SAS-E.

This command enables the context to configure the counters associated with SAP ingress and egress.

egress

Syntax **egress**

Context config>service>epipe>sap>statistics

Description **Platforms supported:** 7210 SAS-D and 7210 SAS-E.

This command enables the context to configure the egress SAP statistics counter and set the mode of the counter.

This counter counts the number of packets forwarded through the SAP.

ingress

Syntax **ingress**

Context config>service>epipe>sap>statistics

Description **Platforms supported:** 7210 SAS-D and 7210 SAS-E.

This command enables the context to configure the ingress SAP statistics counter.

For 7210 E, by default, SAP ingress counters are associated with a SAP and cannot be disabled.

In 7210 SAS-D devices, for access-uplink SAPs the ingress counters are not enabled by default. For access SAPs if the ingress counter is enabled by default, it can be disabled.

The two types of ingress SAP counters are:

- A counter that counts the total packets or octets received on the SAP

- A counter associated with meters defined in the QoS policy of the SAP. This counter counts the in-profile and out-of-profile packets or octets received on the SAP.

forwarded-count

Syntax [no] **forwarded-count**

Context config>service>epipe>sap>statistics>egress
config>service>vpls>sap>statistics>egress
config>service>ies>sap>statistics>egress

Platform supported: 7210 SAS-D

Description This command associates a counter with the SAP. The counter counts the number of packets forwarded through the SAP.

A limited amount of such counters are available for use with access SAPs and access-uplink SAPs.

Use this command before enabling applicable accounting record collection on the SAP to associate a counter with the SAP.

The **no** form of this command disables the packet count

packets-forwarded-count

Syntax [no] **packets-forwarded-count**

Context config>service>epipe>sap>statistics>egress

Platform supported: 7210 SAS-E

Description This command associates a counter with the SAP. The counter counts the number of packets forwarded through the SAP.

A limited amount of such counters are available for use with access SAPs and access-uplink SAPs.

Use this command before enabling applicable accounting record collection on the SAP to associate a counter with the SAP.

The **no** form of this command disables the packet count.

received-count

Syntax [no] **recieved-count**

Context config>service>epipe>sap>statistics>ingress
config>service>vpls>sap>statistics>ingress
config>service>ies>sap>statistics>ingress

Platform supported: 7210 SAS-D

Description This command associates a counter with the SAP. It counts the number of packets and octets received on the SAP (ingress).

A limited amount of such counters are available for use with access-uplink SAPs.

Use this command before enabling applicable accounting record collection on the SAP.

The **no** form of this command disables counter.

counter-mode

Syntax **counter-mode** {in-out-profile-count| forward-drop-count} {**packet** | **octet**}

Context config>service>epipe>sap>statistics>ingress
config>service>vpls>sap>statistics>ingress

Platform supported: 7210 SAS-E

Description This command sets the mode of ingress counters associated with the SAP The mode of the counter can be set to any of the following:

- in-out-profile-count
- forward-drop-count
- packet
- octet

Note: On 7210 SAS-E devices the counter can only count packets or octets at a given time.

The mode of the counter cannot be changed if an accounting policy is already associated with a SAP. Changing the mode of the counter results in loss of previously collected counts and resets the counter. The **no** form of this command is not supported.

Default when either in-out-profile-count or forward-drop-count is in use packet

Parameters **forward-drop-count —** If the counter mode is specified as "forward-drop-count", one counter counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

in-out-profile-count — If the counter mode is specified as "in-out-profile-count", one counter counts the total in-profile packets and octets received on ingress of a SAP and another counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

octet — Sets the mode of ingress counters associated with the SAP to octets.

packet — Sets the mode of ingress counters associated with the SAP to packets.

counter-mode

Syntax **counter-mode** {in-out-profile-count| forward-drop-count}

Context config>service>epipe>sap>statistics>ingress
config>service>vpls>sap>statistics>ingress

Platform supported: 7210 SAS-D

Description This command allows the user to set the counter mode for the counters associated with sap ingress meters (a.ka. policers). A pair of counters is available with each meter. These counters count different events based on the counter mode value.

Note: The counter mode can be changed if an accounting policy is associated with a SAP. If the counter mode is changed the counters associated with the meter are reset and the counts are cleared. If an accounting policy is in use when the counter-mode is changed a new record will be written into the current accounting file.

Execute the following sequence of commands to ensure a new accounting file is generated when the counter-mode is changed:

1. Execute the command **config>service>epipe/vpls>sap> no collect-stats**, to disable writing of accounting records.
2. Change the counter-mode to the desired value, execute the command **config>service>epipe/vpls>sap>counter-mode {in-out-profile-count| forward-drop-count}**.
3. Execute the command **config>service>epipe/vpls>sap> collect-stats**, to enable writing of accounting records.

The **no** form of the command restores the counter mode to the default value.

Default in-out-profile-count

Parameters **forward-drop-count** — If the counter mode is specified as "forward-drop-count", one counter counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

in-out-profile-count — If the counter mode is specified as "in-out-profile-count", one counter counts the total in-profile packets and octets received on ingress of a SAP and another counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**no-endpoint**] [**create**]
spoke-sdp *sdp-id*[**:***vc-id*] **endpoint** *endpoint-name*

no spoke-sdp *sdp-id*[:*vc-id*]

Context config>service>epipe

Description **Platforms supported:** 7210 SAS-K 2F4T6C.

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 or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with an Epipe or VPL 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-enddevices can participate in the service.

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

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

Special Cases **Epipe —** At most, only one *sdp-id* can be bound to an Epipe service. Since an Epipe is a point-to-point service, it can have, at most, two end points. The two end points can be one SAP and one SDP or two SAPs. Vc-switching VLLs are an exception. If the VLL is a "vc-switching" VLL, then the two endpoints must both be SDPs.

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 — Removes the association of a spoke SDP with an explicit endpoint name.

endpoint *endpoint-name* **—** Specifies the name of the service endpoint.

control-word

Syntax [**no**] **control-word**

Context config>service>epipe>spoke-sdp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

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.

precedence

Syntax **precedence** [*precedence-value* | **primary**]
no precedence

Context config>service>epipe>spoke-sdp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

The **no** form of the command returns the precedence value to the default.

Default 4

Parameters *precedence-value* — Specifies the spoke SDP precedence.

Values 1 — 4

primary — Specifies to make this the primary spoke SDP.

pw-status-signaling

Syntax [**no**] **pw-status-signaling**

Context config>service>epipe>spoke-sdp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

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

vc-label

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

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

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures the egress VC label.

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

Values 16 — 1048575

vc-label

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

Context config>service>cpipe>spoke-sdp>ingress
config>service>epipe>spoke-sdp>ingress

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures the ingress VC label.

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

Values 2048 — 18431

vlan-vc-tag

Syntax **vlan-vc-tag** *0..4094*
no vlan-vc-tag [*0..4094*]

Context config>service>epipe>spoke-sdp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The **no** form of this command disables the command

Default no vlan-vc-tag

Parameters *0..4094* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

spoke-sdp-fec

Syntax **spoke-sdp-fec**
spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**]
spoke-sdp-fec *spoke-sdp-fec-id* **no-endpoint**
spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**] **endpoint** *name*

Context config>service>epipe

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command binds a service to an existing Service Distribution Point (SDP), using a dynamic MS-PW.

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

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

When using dynamic MS-PWs, the particular SDP to bind-to is automatically selected based on the Target Attachment Individual Identifier (TAII) and the path to use, specified under spoke-SDP FEC. The selected SDP will terminate on the first hop S-PE of the MS-PW. Therefore, an SDP must already be defined in the config>service>sdp context that reaches the first hop 7210 of the MS-PW. The 7210 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 PW with the PW ID FEC. However, the spoke-sdp-fec command enables PWs 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 *spoke-sdp-fec-id* — An unsigned integer value identifying the spoke-SDP.

Values 1 —4294967295

fec *fec-type* — An unsigned integer value for the type of the FEC used by the MS-PW.

Values 129 — 130

aii-type *aii-type* — An unsigned integer value for the Attachment Individual Identifier (AII) type used to identify the MS-PW endpoints.

Values 1 — 2

endpoint *endpoint-name* — Specifies the name of the service endpoint

no endpoint — Adds or removes a spoke SDP association.

auto-config

Syntax [**no**] **auto-config**

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

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command enables single sided automatic endpoint configuration of the spoke-SDP. The 7210 acts as the passive T-PE for signaling this MS-PW.

Automatic Endpoint Configuration allows the configuration of a spoke-SDP endpoint without specifying the TAII associated with that spoke-SDP. It allows a single-sided provisioning model where an incoming label mapping message with a TAII that matches the SAII of that spoke-SDP to be automatically bound to that endpoint. In this mode, the far end T-PE actively initiates MS-PW signaling and will send the initial label mapping message using T-LDP, while the 7210 T-PE for which auto-config is specified will act as the passive T-PE.

The **auto-config** command is blocked in CLI if signaling active has been enabled for this spoke-SDP. It it is only applicable to spoke SDPs configured under the Epipe, IES and VPRN interface context.

The **no** form of the command means that the 7210 T-PE either acts as the active T-PE (if signaling active is configured) or automatically determines which 7210 will initate MS-PW signaling based on the prefix values configured in the SAII and TAII of the spoke-SDP. If the SAII has the greater prefix value, then the 7210 will initiate MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAII has the greater value prefix, then the 7210 will assume that the far end T-PE will intiate MS-PW signaling and will wait for that label mapping message beforeresponding with a T-LDP label mapping message for the MS-PW in the reverse direction.

Default no auto-config

path

Syntax **path** *name*
no path

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

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command specifies the explicit path, containing a list of S-PE hops, that should be used for this spoke SDP. The path-name should correspond to the name of an explicit path configured in the **config>service>pw-routing** context.

If no path is configured, then each next-hop of the MS-PW used by the spoke-SDP will be chosen locally at each T-PE and S-PE.

Default no path

Parameters *path-name* — The name of the explicit path to be used, as configured under config>service>pw-routing.

precedence

Syntax **precedence** *prec-value*
precedence primary
no precedence

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

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

The **no** form of the command returns the precedence value to the default.

Default 42

Parameters *precedence-value* — Specifies the spoke SDP precedence.

Values 1 — 4

primary — Specifies to make this the primary spoke SDP.

pw-template-bind

Syntax **pw-template-bind** *policy-id*
no pw-template-bind

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

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command binds includes the parameters included in a specific PW Template to a spoke SDP.

The **no** form of the command removes the values from the configuration.

Default none

Parameters *policy-id* — Specifies the existing policy ID

Values 1 — 2147483647

retry-count

Syntax **retry-count** *retry-count*
no retry-count

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

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

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 *retry-count* — The maximum number of retries before putting the spoke-sdp into the shutdown state.

Values 10 — 10000

retry-timer

Syntax **retry-timer** *retry-timer*
no retry-timer

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

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

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 "AII unreachable".

The **no** form of this command reverts the timer to its default value.

Default 30

Parameters *retry-timer* — The initial retry-timer value in seconds.

Values 10 — 480

saii-type2

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

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

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 7210 T-PE. This value must correspont 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 7210 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 **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command enables a user to configure this 7210 as the active pr passive T-PE for signaling this MS-PW, or to automatically select whether this T-PE is active or passive based on the prefix. In an active role, this endpoint initiates MS-PW signaling without waiting for a T-LDP label mapping message to arrive from the far end T-PE. In a passive role, it will wait for the initial label mapping message from the far end before sending a label maping for this end of the PW. In auto maode, if the SAII has the greater prefix value, then the 7210 will initiate MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAII has the greater value prefix, then the 7210 will assume that the far end T-PE will intiate MS-PW signaling and will wait for that label mapping message beforeresponding with a T-LDP label mapping message for the MS-PW in the reverse direction.

The **no** form of the command means that the 7210 T-PE automatically selects the which 7210 will initate MS-PW signaling based on the prefix values configured in the SAII and TAII of the spoke-SDP, as described above.

Default auto

Parameters *signaling* — Configures this 7210 as the active T-PE for signaling this MS-PW.

Values auto, master

standby-signaling-slave

Syntax [**no**] **standby-signaling-slave**

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

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

When this command is enabled, the node will block the transmit forwarding direction of a spoke-SDP based on the setting of the 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. An existing spoke-sdp cannot be moved to an endpoint if the setting of standby-signaling-slave is not the same as at the endpoint level. If the standby-signaling-slave setting is changed at the endpoint level, then that change is automatically populated to the member spoke-sdps.This command is mutually exclusive with and endpoint that is part of an MC-LAG/MC-APS, ICB and MC endpoint, or for which standby-signaling-master has been enabled.

If the command is disabled, the node assumes the existing release 5.0 mode of behavior for the forwarding on the spoke-SDP.

Default standby-signaling-slave is disabled by default.

taii-type2

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

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

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

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 7210 T-PE. This value must correspont 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 7210 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

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 279
  - → VPLS Packet Walk-through for a Local service on page 280
  - → VPLS Packet Walk-through for a Distributed service on page 283
- VPLS Features on page 286
  - → VPLS Enhancements on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C on page 286
  - → VPLS Enhancements on 7210 SAS-K 2F4T6C on page 286
  - → VPLS over SAPs on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F4T1C, and 7210 SAS-K 2F4T6C on page 288
  - → VPLS MAC Learning and Packet Forwarding on page 288
  - → L2 Forwarding Table Management on page 294
  - → VPLS and Spanning Tree Protocol on page 299
- Routed VPLS on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C on page 323
  - → Routed VPLS - IES IP Interface Binding- 7210 SAS-D and 7210 SAS-K 2F2T1C on page 323
  - → Routed VPLS - Assigning a Service Name to a VPLS Service on page 324
  - → Routed VPLS - Service Binding Requirements on page 324
  - → Routed VPLS - ARP and VPLS FIB Interactions on page 327
  - → Routed VPLS - Specific ARP Cache Behavior on page 328

→ Routed VPLS - SAPs Only Supported on Standard Ethernet Ports on page 330

→ Routed VPLS - VPLS SAP Ingress IP Filter Override on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C on page 331

→ Routed VPLS support available and Caveats on page 337

- Epipe Emulation using Dot1q VLAN range SAP in VPLS with G.8032 on page 338

  → Epipe Emulation using Dot1q VLAN range SAP in VPLS with G.8032 -Configuration guidelines and restrictions on page 340

- Configuring a VPLS Service with CLI on page 343
- Service Management Tasks on page 390

VPLS Service Overview

Virtual Private LAN Service (VPLS) 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 (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.

7210 SAS-E, 7210 SAS-D, and 7210 SAS-K 2F2T1C do not support IP/MPLS uplinks and support only local VPLS service. These devices support provisioning of access SAPs or access-uplink SAPs to connect to the provider edge IP/MPLS routers, allowing users to provision an end-to-end VPLS VPN service using the concept of hierarchical VPLS.

7210 SAS-K 2F4T6C supports IP/MPLS uplinks and support both local VPLS service and distributed VPLS service. It provides an option to use VPLS service configured on the node with spoke-SDPs to connect to the provider edge IP/MPLS routers or use access SAPs or access-uplink SAPs to connect to the provider edge IP/MPLS routers, allowing users to provision an end-to-end VPLS VPN service using the concept of hierarchical VPLS. 7210 SAS-K 2F4T6C also supports mesh SDP for use in smaller network configurations.

VPLS Packet Walk-through for a Local service

This section provides an example of VPLS processing for a local VPLS service using QinQ access uplink SAPs to connect to the PE router. It is applicable to 7210 SAS-E, 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C (7210 SAS-K 2F4T6C when using access uplink SAPs). It describes VPLS processing of a customer packet sent across the network from site-A connected to PE-Router-A through a 7210 SAS to site-C connected through 7210 SAS to PE-Router-C (Figure 42). This section does not discuss the processing on the PE routers, but only on 7210 SAS routers.

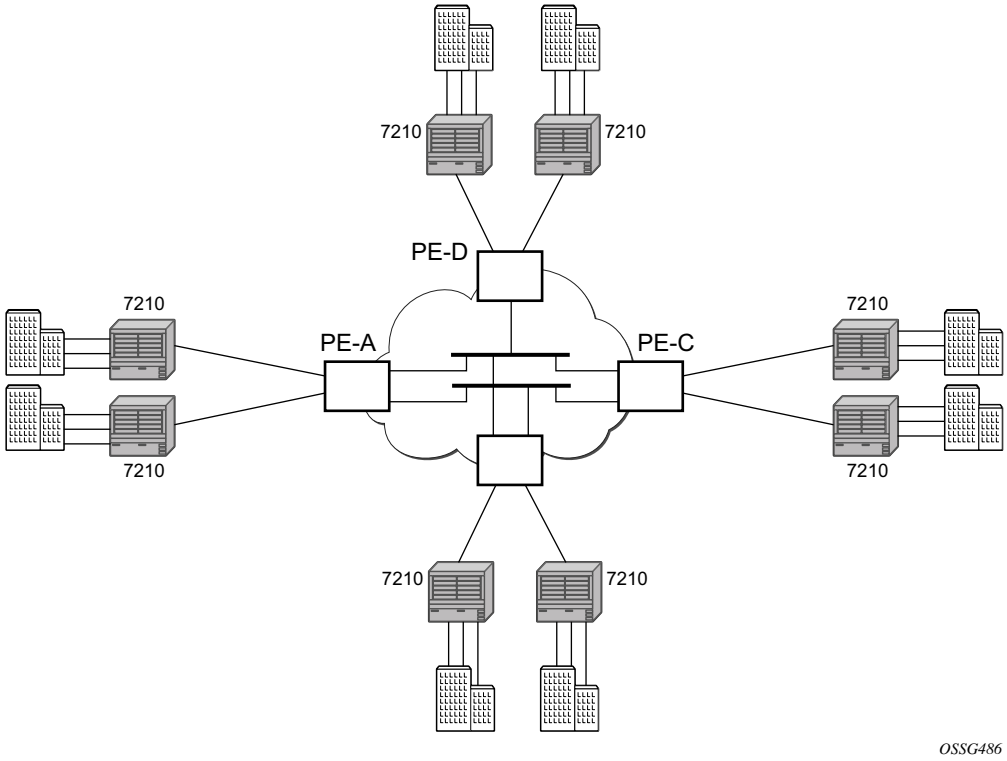

Figure 42: VPLS Service Architecture

1. 7210-A (Figure 43)
    a. Service packets arriving at 7210-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

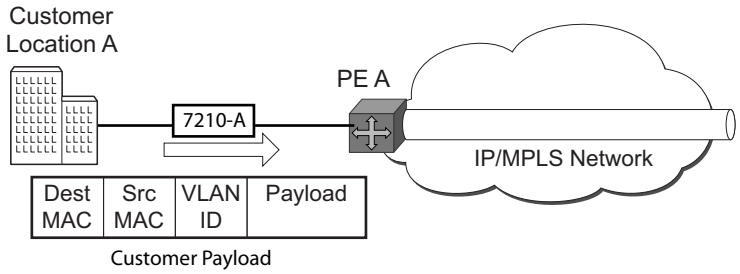

Figure 43: Access Port Ingress Packet Format and Lookup

b. 7210-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 44):

d. If the destination MAC address has already been learned by 7210 an existing entry in the FIB table identifies destination uplink QinQ SAP used for sending the packet towards the PE-Router-A.

e. The customer packet is sent on this uplink SAP once the IEEE 802.1Q tag is stripped and the uplink SAP tag is added to the packet.

   For an Unknown MAC Address (Figure 44):

f. If the destination MAC address has not been learned, 7210 will flood the packet to all the uplink SAPs participating in the service.

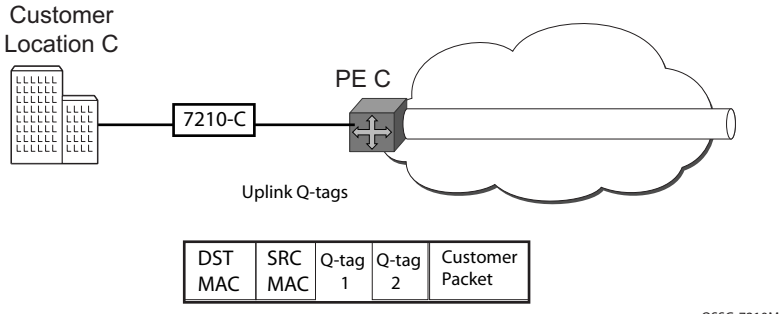

Figure 44: Network Port Egress Packet Format and Flooding

2. Core Router Switching

a. The PE router will encapsulate this packet in the appropriate MPLS header and transport it across the core network to the remote 7210-C.

3. 7210-C (Figure 43)

a. 7210-C associates the packet with the VPLS instance based on the VLAN tags in the received packet.

b. 7210-C learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the access uplink port 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 7210-C (unknown MAC address).

d. If the destination MAC address has been learned by 7210-C, an existing entry in the FIB table identifies the local access port and the IEEE 802.1Q tag (if any) to be added before sending the packet to customer Location-C. The egress Q tag may be different than the ingress Q tag.

e. If the destination MAC address has not been learned, 7210 will flood the packet to all the access SAPs participating in the service.

VPLS Packet Walk-through for a Distributed service

This section provides an example of VPLS processing for a local VPLS service using spoke-SDP/mesh-SDP as uplinks to connect to the PE routers. It is applicable to 7210 SAS-K 2F4T6C when using spoke-SDP/mesh-SDP as uplinks to connect to the PE routers. It describes VPLS processing of a customer packet sent across the network from site-A connected to PE-Router-A through a 7210 SAS to site-C connected through 7210 SAS to PE-Router-C (Figure 45). This section does not discuss the processing on the PE routers, but only on 7210 SAS routers.

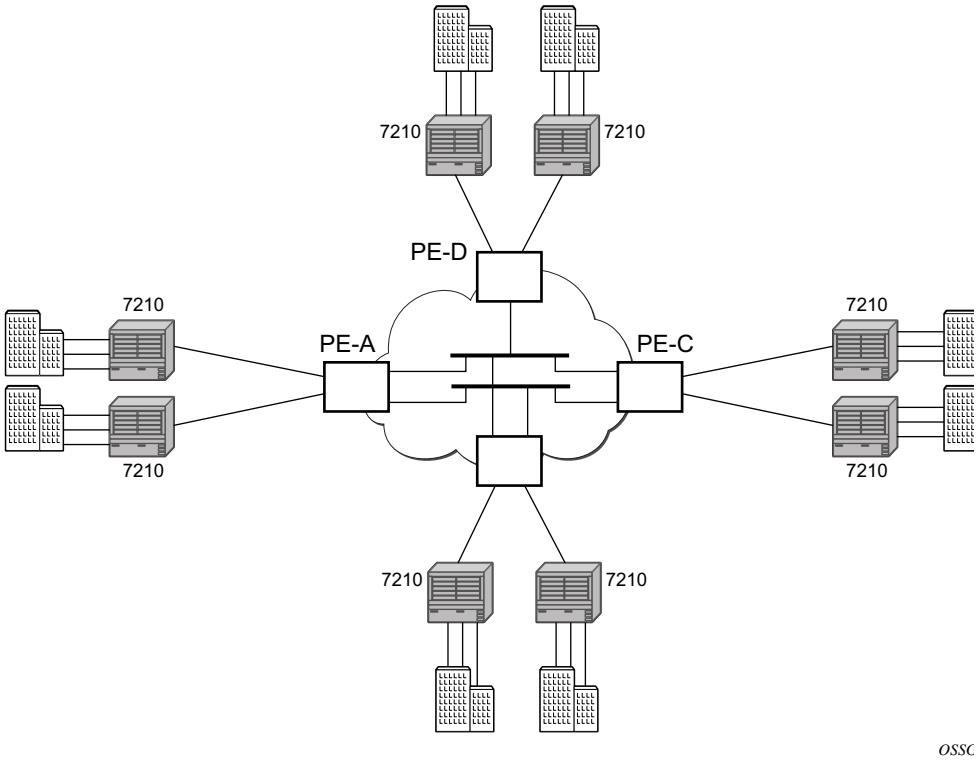

Figure 45: VPLS Service Architecture

1. 7210-A (Figure 46)
   a. Service packets arriving at 7210-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.

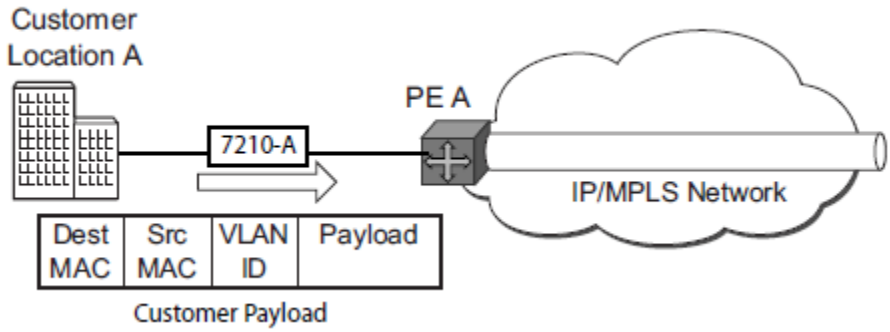

Figure 46: Access Port Ingress Packet Format and Lookup

b. 7210-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 46):

a. If the destination MAC address has already been learned by 7210, an existing entry in the FIB table identifies the far-end PE-Router and the service VC-label (inner label) used before sending the packet to PE-Router-A.

b. 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 a Unknown MAC Address (Figure 46):

a. If the destination MAC address has not been learned, 7210 will flood the packet to all the spoke SDPs participating in the service.

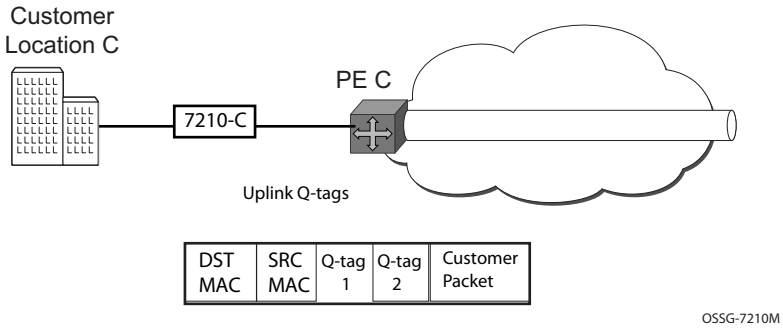

Figure 47: Network Port Egress Packet Format and Flooding

2. Core Router Switching

   a. The PE router will encapsulate this packet in the appropriate MPLS header and transport it across the core network to the remote 7210-C.

3. 7210-C (Figure 46)

   a. 7210-C associates the packet with the VPLS instance based on the VLAN tags in the received packet.

   b. 7210-C learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the spoke-SDP 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 7210-C (unknown MAC address).

   d. If the destination MAC address has been learned by 7210-C, an existing entry in the FIB table identifies the local access port and the IEEE 802.1Q tag (if any) to be added before sending the packet to customer Location-C. The egress Q tag may be different than the ingress Q tag.

   e. If the destination MAC address has not been learned, 7210 will flood the packet to all the access SAPs and possible other spoke-SDPs participating in the service.

VPLS Features

This section features:

- VPLS Enhancements on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C on page 286
- VPLS and Spanning Tree Protocol on page 299
- VPLS Access Redundancy on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C on page 307

VPLS Enhancements on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

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 including:
  - → Configurable FIB size limit
  - → FIB size alarms
  - → MAC learning disable
  - → Discard unknown
  - → Aging timers for learned MAC addresses.
- Implementation of Spanning Tree Protocol (STP) parameters on a per VPLS and per SAP basis.
- IGMP snooping on a per-SAP basis.

VPLS Enhancements on 7210 SAS-K 2F4T6C

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 including:
  - → Configurable FIB size limit
  - → FIB size alarms
  - → MAC learning disable
  - → Discard unknown
  - → Aging timers for learned MAC addresses
- Implementation of Spanning Tree Protocol (STP) parameters on a per VPLS and per SAP basis.
- IGMP snooping on a per-SAP and SDP basis.

VPLS over MPLS - supported on 7210 SAS-K 2F4T6C

The VPLS architecture proposed in draft-ietf-ppvpn-vpls-ldp-0x.txt specifies the use of provider equipment (PE) 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. The 7210 SAS M supports only H-VPLS. Multiple VPLS services can be offered over the same set of LSP tunnels. Signaling specified in RFC 4905 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/HVPLS is provided over MPLS by:

- Connecting 7210 SAS to bridging-capable provider edge (PE) routers through a mesh/ spoke SDP. The PE routers are connected using a full mesh of LSPs.
- 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 on page 288).
- Using a separate forwarding information base (FIB) per VPLS service.

VPLS over SAPs on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F4T1C, and 7210 SAS-K 2F4T6C

7210 SAS-D and 7210 SAS-E devices support QinQ SAPs or Dot1q SAPs that allows them to connect to upstream PE nodes providing IP/MPLS transport.

VPLS is provided over QinQ/Dot1q SAPs by:

- Connecting bridging-capable 7210 SAS devices.
- Replicating unknown and broadcast traffic in a service domain.
- Enabling MAC learning over QinQ/Dot1q SAPs and access ports (see VPLS MAC Learning and Packet Forwarding on page 288).
- Using a separate forwarding information base (FIB) per VPLS service.

VPLS MAC Learning and Packet Forwarding

The 7210 SAS edge devices perform the packet replication required for broadcast and multicast traffic across the bridged domain. MAC address learning is performed by the 7210 SAS device to reduce the amount of unknown destination MAC address flooding.

On 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C, user can configure Layer-2 uplinks (i.e. Dot1q SAPs or QinQ SAPs) for uplink connectivity and access SAPs for customer connectivity as logical bridge ports in a VPLS service.

On 7210 SAS-K 2F4T6C can configure access SAPs for customer connectivity and mesh-sdp/ spoke-sdp for uplink connectivity as logical bridge ports in a VPLS service.

Each 7210 SAS 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 using QinQ SAPs created on access uplink ports Unknown destination packets (for example, the destination MAC address has not been learned) are forwarded on all SAPs for that service until the target station responds and the MAC address is learned by the 7210 SAS associated with that service.

IGMP Snooping on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

In Layer 2 switches, multicast traffic is treated like an unknown MAC address or broadcast frame, which causes the incoming frame to be flooded out (broadcast) on every port within a VLAN. Although this is acceptable behavior for unknowns and broadcast frames, this flooded multicast traffic may result in wasted bandwidth on network segments and end stations, as IP multicast hosts can join and be interested in only specific multicast groups.

IGMP snooping entails using information in Layer 3 protocol headers of multicast control messages to determine the processing at Layer 2. By doing so, an IGMP snooping switch provides the benefit of conserving bandwidth on those segments of the network in which no node has expressed interest in receiving packets addressed to the group address.

IGMP snooping can be enabled in the context of VPLS services. The IGMP snooping feature allows for optimization of the multicast data flow to only the SAPs members of the group. The system builds a database of group members per service by listening to IGMP queries and reports from each SAP:

- When the switch receives an IGMP report from a host for a multicast group, the switch adds the host port number to the forwarding table entry.
- When it receives an IGMP leave message from a host, it removes the host port from the table entry, if no other group members are present. It also deletes entries if it does not receive periodic IGMP membership reports from the multicast clients.

The following are IGMP snooping features:

- IGMP v1, v2, and v3 are supported (RFC 1112, *Host Extensions for IP Multicasting*, and RFC 2236, *Internet Group Management Protocol, Version 2*).
- IGMP snooping can be enabled and disabled on individual VPLS service instances.
- IGMP snooping can be configured on individual SAPs that are part of a VPLS service. When IGMP snooping is enabled on a VPLS service, all its contained SAPs automatically have snooping enabled.
- Fast leave terminates the multicast session immediately, rather than using the standard group-specific query to check if other group members are present on the network.
- SAPs can be statically configured as multicast router ports. This allows the operator to control the set of ports to which IGMP membership reports are forwarded.
- Static multicast group membership on a per SAP basis can be configured.
- The maximum number of multicast groups (static and dynamic) that a SAP can join can be configured. An event is generated when the limit is reached.
- The maximum number of multicast groups (static and dynamic) that a VPLS instance simultaneously supports can be configured.

- Proxy summarization of IGMP messages reduces the number of IGMP messages processed by upstream devices in the network.
- IGMP filtering allows a subscriber to a service or the provider to block, receive, or transmit permission (or both) to individual hosts or a range of hosts.
  The following types of filters can be defined:
  → Filter group membership that report from a host or range of hosts. This filtering is performed by importing an appropriately-defined routing policy into the SAP.
  → Filters that prevent a host from transmitting multicast streams into the network. The operator can define a data-plane filter (ACL) that drops all multicast traffic, and apply this filter to a SAP.

IGMP Snooping on 7210 SAS-K 2F4T6C

In Layer 2 switches, multicast traffic is treated like an unknown MAC address or broadcast frame, which causes the incoming frame to be flooded out (broadcast) on every port within a VLAN. Although this is acceptable behavior for unknowns and broadcast frames, this flooded multicast traffic may result in wasted bandwidth on network segments and end stations, as IP multicast hosts can join and be interested in only specific multicast groups.

IGMP snooping entails using information in Layer 3 protocol headers of multicast control messages to determine the processing at Layer 2. By doing so, an IGMP snooping switch provides the benefit of conserving bandwidth on those segments of the network in which no node has expressed interest in receiving packets addressed to the group address.

IGMP snooping can be enabled in the context of VPLS services. The IGMP snooping feature allows for optimization of the multicast data flow to only those SAPs and SDPs that are members of the group. The system builds a database of group members per service by listening to IGMP queries and reports from each SAP and SDP:

- When the switch receives an IGMP report from a host for a multicast group, the switch adds the host port number to the forwarding table entry.
- When it receives an IGMP leave message from a host, it removes the host port from the table entry, if no other group members are present. It also deletes entries if it does not receive periodic IGMP membership reports from the multicast clients.

The following are IGMP snooping features:

- IGMP v1, v2, and v3 are supported (RFC 1112, Host Extensions for IP Multicasting, and RFC 2236, Internet Group Management Protocol, Version 2). On 7210 SAS-K12 IGMP v1, v2, v3 is supported with both with L2 uplinks and MPLS uplinks.

- IGMP snooping can be enabled and disabled on individual VPLS service instances.
- On 7210 SAS-K 2F4T6C, IGMP snooping can be configured on individual SAPs and SDPs that are part of a VPLS service. When IGMP snooping is enabled on a VPLS service, all its contained SAPs and SDPs automatically have snooping enabled.
- Fast leave terminates the multicast session immediately, rather than using the standard group-specific query to check if other group members are present on the network.
- On 7210 SAS-K 2F4T6C, SDPs and SAPs can be statically configured as multicast router ports. This allows the operator to control the set of ports to which IGMP membership reports are forwarded.
- On 7210 SAS-K 2F4T6C, Static multicast group membership on a per SAP and SDP basis can be configured.
- On 7210 SAS-K 2F4T6C, the maximum number of multicast groups (static and dynamic) that a SAP and SDP can join can be configured. An event is generated when the limit is reached.
- The maximum number of multicast groups (static and dynamic) that a VPLS instance simultaneously supports can be configured.
- Proxy summarization of IGMP messages reduces the number of IGMP messages processed by upstream devices in the network.
- IGMP filtering allows a subscriber to a service or the provider to block, receive, or transmit permission (or both) to individual hosts or a range of hosts. The following types of filters can be defined:
  - → Filter group membership that report from a particular host or range of hosts. This filtering is performed by importing an appropriately-defined routing policy into the SAP and SDP.
  - → Filters that prevent a host from transmitting multicast streams into the network. The operator can define a data-plane filter (ACL) that drops all multicast traffic, and apply this filter to a SAP and SDP.

Multicast VLAN Registration (MVR) support

Multicast VPLS Registration (MVR) is a bandwidth optimization method for multicast in a broadband services network. MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on one or more network-wide multicast VPLS instances.

MVR assumes that subscribers join and leave multicast streams by sending IGMP join and leave messages. The IGMP leave and join message are sent inside the VPLS to which the subscriber port is assigned. The multicast VPLS is shared in the network while the subscribers remain in separate VPLS services. Using MVR, users on different VPLS cannot exchange any information between them, but still multicast services are provided.

On the MVR VPLS, IGMP snooping must be enabled. On the user VPLS, IGMP snooping and MVR work independently. If IGMP snooping and MVR are both enabled, MVR reacts only to join and leave messages from multicast groups configured under MVR. Join and leave messages from all other multicast groups are managed by IGMP snooping in the local VPLS. This way, potentially several MVR VPLS instances could be configured, each with its own set of multicast channels.

MVR by proxy — In some situations, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behavior) but to another SAP. This is called MVR by proxy.

NOTE: In the diagram below, reference to SDP is applicable to 7210 SAS-K 2F4T6C.

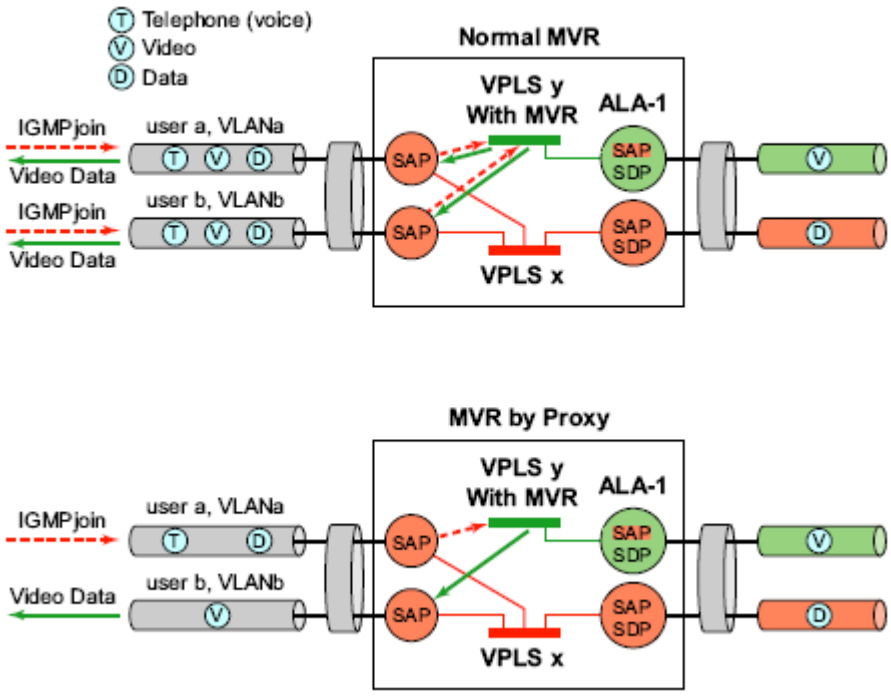

Figure 48: MVR and MVR by Proxy

Configuration Guidelines for MVR

In a MVR configuration, the svc-sap-type of the VPLS service that is the source, known as 'mvr vpls service' and the svc-sap-type of the VPLS service that is the sink, known as 'user vpls service' should match.

L2 Forwarding Table Management

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

FIB Size on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

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

- MAC FIB size limits — Allows users to specify the maximum number of MAC FIB entries learned locally for a SAP. If the configured limit is reached, then no new addresses will be learned from the SAP until at least one FIB entry is aged out or cleared.
  - → When the limit is reached on a SAP, 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, unknown destination MAC addresses are discarded.
  - → 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 at the VPLS service level allows users to disable the dynamic learning function on the service. Disable learning is not supported at the SAP level.
  - → Disable aging allows users to turn off aging for learned MAC addresses on a SAP of a VPLS service instance.

FIB Size on 7210 SAS-K 2F4T6C

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

- MAC FIB size limits - Allows users to specify the maximum number of MAC FIB entries 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 and spoke-SDP until at least one FIB entry is aged out or cleared.

→ When the limit is reached on a SAP and 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, unknown destination MAC addresses are discarded.

→ The log event SAP/SDP MAC limit reached is generated when the limit is reached. When the condition is cleared, the log event SAP/SDP MAC Limit Reached Condition Cleared is generated.

→ Disable learning at the VPLS service level allows users to disable the dynamic learning function on the service. Disable learning is not supported at the SAP level and spoke-SDP level.

→ Disable aging allows users to turn off aging for learned MAC addresses on a SAP and 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 Aging Timers on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

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). The age timer for the VPLS instance specifies the aging time for locally learned 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, in some circumstances, can take up to two times their respective age timer to be aged out.

Local and Remote Aging Timers on 7210 SAS-K 2F4T6C

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 time-period (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 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. MAC learning can be disabled for services.

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 using service ingress QoS policies. In a service ingress QoS policy, individual meters can be defined per forwarding class to provide rate-limiting/policing of broadcast traffic, MAC multicast traffic and unknown destination MAC traffic.

Note: In 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, only service ingress Queues are available and same functionality is achieved through queues. There is no support for meters and queues.

MAC Move

The MAC move feature is useful to protect against undetected loops in a VPLS topology, as well as the presence of duplicate MACs in a VPLS service.

If two clients in the VPLS have the same MAC address, the VPLS will experience a high relearn rate for the MAC. When MAC move is enabled, the 7210 SAS E shuts down the SAP and create an alarm event when the threshold is exceeded.

When MAC move is enabled, the 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C shuts down the SAP and create an alarm event when the threshold is exceeded.

When MAC move is enabled, the 7210 SAS-K 2F4T6C shuts down the SAP or spoke-SDP and create an alarm event when the threshold is exceeded.

Split Horizon SAP Groups on 7210 SAS-K

Note: Split Horizon group per service is supported only on 7210 SAS-K devices. On 7210 SAS-D and 7210 SAS-E, port-based split horizon groups are supported.

In many applications, the split-horizon group concept involving a group of SAPs is useful to prevent direct customer-to-customer traffic exchange (without the traffic being sent to the head-end service nodes). This extension is referred to as a split horizon SAP group. Traffic arriving on a SAP or a spoke SDP within a split horizon group will not be forwarded to other SAPs configured in the same split horizon group, but will be forwarded to other SAPs, which are not part of the split horizon group.

Split Horizon SAP Groups and Split Horizon Spoke SDP Groups on 7210 SAS-K 2F4T6C

Within the context of VPLS services, a loop-free topology inside a fully meshed VPLS core is achieved by applying a split-horizon forwarding concept. The 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 residential 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/spoke SDP group. Traffic arriving on a SAP or a spoke SDP within a split horizon group will not be forwarded to other SAPs and spoke SDPs configured in the same split horizon group, but will be forwarded to other SAPs/spoke SDPs, which are not part of the split horizon group.

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 7210 SAS participating in the service learns where the customer MAC addresses reside, on ingress SAPs.

Unknown destinations, broadcasts, and multicasts are flooded to all other SAPs in the service. If SAPs are connected, 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 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 parameters can affect the behavior. For information on command usage, descriptions, and CLI syntax, refer to Configuring a VPLS Service with CLI on page 343.

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 7210 SAS 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. Inter-working 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 7210 SAS supports one BDPU 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 7210 SAS implementation of Management VPLS (mVPLS) is used to group different VPLS instances under single RSTP instance. Introducing MSTP into the mVPLS allows the following:

- Inter-operation with traditional Layer 2 switches in access network.
- 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. 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 the following:

- Configuring mVPLS on VPLS-PEs (only PEs directly connected to GigE MAN network).
- 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 49 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

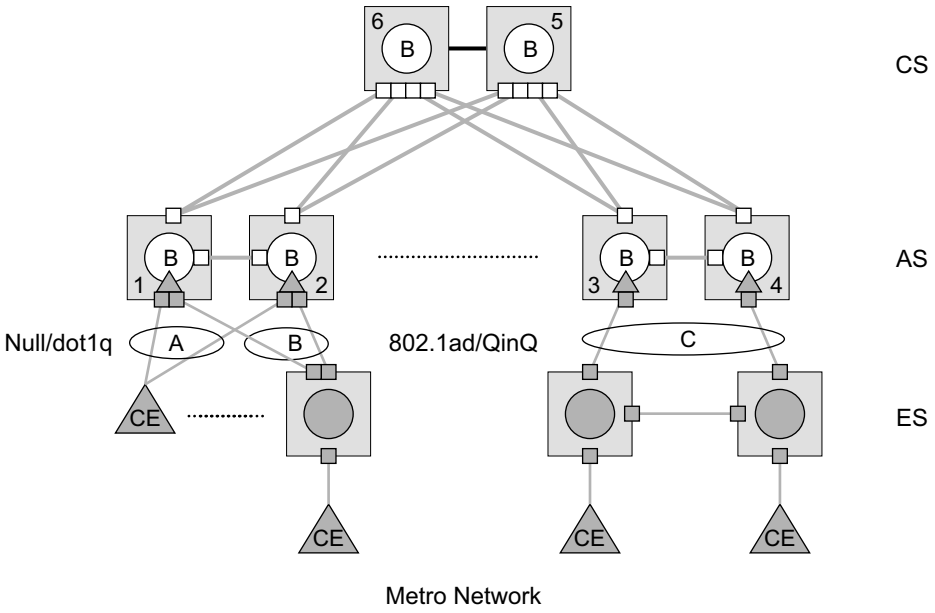

Figure 49: Access Resiliency

The following mechanisms are supported for the 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.

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.

MSTP General Principles

MSTP represents modification of RSTP that allows the grouping of different VLANs into multiple MSTIs. To enable different devices to participate in MSTIs, they must be consistently configured. A collection of interconnected devices that have the same MST configuration (region-name, revision and VLAN-to-instance assignment) comprises an MST region.

There is no limit to the number of regions in the network, but every region can support a maximum of 16 MSTIs. Instance 0 is a special instance for a region, known as the Internal Spanning Tree (IST) instance. All other instances are numbered from 1 to 4094. IST is the only spanning-tree instance that sends and receives BPDUs (typically BPDUs are untagged). All other spanning-tree instance information is included in MSTP records (M-records), which are encapsulated within MSTP BPDUs. This means that 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 7210 SAS Platform

The 7210 SAS 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 in MPLS-based VPLS service (7210 SAS-K 2F4T6C only)

To interconnect 7210 SAS devices (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 to guarantee that service tunnels will not be blocked in any circumstance without imposing artificial restrictions on the placement of the root bridge within the network. The modifications introduced are fully compliant with the 802.1D-2004 STP specification.

When running MSTP, spoke SDPs cannot be configured. Also, ensure that all bridges connected by mesh SDPs are in the same region. If not, the mesh will be prevented from becoming active (trap is generated).

To achieve this, all mesh SDPs are dynamically configured as either root ports or designated ports. The PE devices participating in each VPLS mesh determine (using the root path cost learned as part of the normal protocol exchange) which of the 7210 SAS devices is closest to the root of the network. This PE device is internally designated as the primary bridge for the VPLS mesh. Because 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 ensures that network ports always remain in forwarding state.

A combination of the above mentioned features ensure that network ports are never blocked and maintain interoperability with bridges external to the mesh running STP instances.

L2PT Termination

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

The 7210 SAS 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.

The 7210 SAS nodes support L2PT termination for STP BPDUs. More specifically:

- At ingress of every SAP 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. The type of protocol can be derived from LLC and SNAP encapsulation.
- In egress direction, PDUs of the corresponding protocol received on all VPLS ports will be intercepted and L2PT encapsulation will be performed for SAPs configured as L2PT termination points. Because of the implementation reasons, PDU interception and re-direction to CPM can be performed only at ingress. Therefore, to comply with the above requirement, as soon as at least 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) 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 to provide an inter-operable e2e solution.

To address these network designs, BPDU format translation is supported on 7210 SAS devices. If enabled on a given SAP, 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 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

L2PT termination for only STP (Spanning Tree Protocol) and PVST (Per VLAN Spanning Tree Protocol) are supported on 7210 SAS-E.

Cisco Discovery Protocol (CDP), Digital Trunking Protocol (DTP), Port Aggregation Protocol (PAgP), Uni-directional Link Detection (UDLD), Virtual Trunk Protocol (VTP), STP (Spanning Tree Protocol) and PVST (per-VLAN Spanning Tree protocol) are supported on 7210 SAS-D and 7210 SAS-K devices.

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/BPDU translation enabled.
- Forwarding occurs on the CPM and uses CPU processing cycles.

VPLS Access Redundancy on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

NOTE: On 7210 SAS-K 2F4T6C, these mechanisms are applicable when L2 uplinks are used.

An application of hierarchical VPLS is using 7210 SAS E nodes that are MPLS-enabled which must have Ethernet links to the closest PE node. To protect against failure of the PE node, 7210 SAS E can be dual-homed and have two SAPs on two PE nodes.

Listed below are several mechanisms to resolve a loop in an access network where 7210s are used:

- STP-based access, with or without mVPLS.
- Ethernet APS using G.8032.
- Non-STP-based access using mechanisms such as active/standby links and MC-LAG on the PE.

STP-Based Redundant Access to VPLS

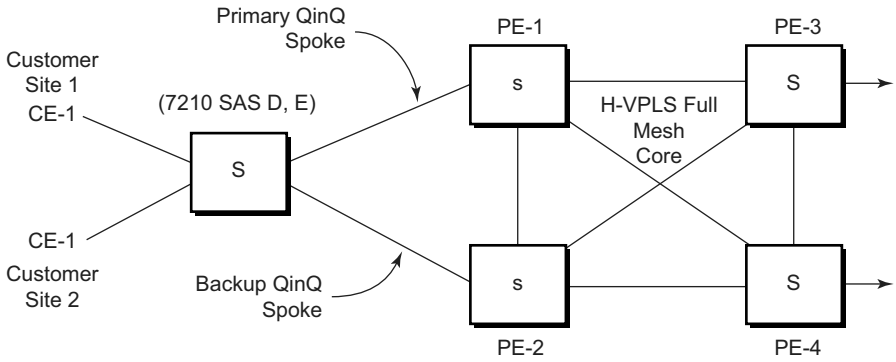

Figure 50: Dual Homed 7210 SAS E Acting as MTU-s in Two-Tier Hierarchy H-VPLS

In configuration shown in Figure 50, STP is activated on the MTU and two PEs to resolve a potential loop.

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. When TCN (topology change notification) is received with in the VPLS domain all MACs learned on SAPs are flushed except the SAP on which TCN was received.

VPLS Redundancy on 7210 SAS-K 2F4T6C using MPLS uplinks

The VPLS standard (RFC 4762, Virtual Private LAN Services Using LDP Signaling) includes provisions for hierarchical VPLS, using point-to-point spoke SDPs. Two applications have been identified for spoke SDPs:

- To connect to Multi-Tenant Units (MTUs) to PEs in a metro area network;
- To interconnect the VPLS nodes of two 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. In VPLS services, only two spoke-SDPs are allowed in an endpoint.

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 (on 7210 SAS-K 2F4T6C)

When two or more meshed VPLS instances are interconnected by redundant spoke SDPs (as shown in Figure 51), a loop in the topology results. 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 VPLS in this topology is not efficient, the node includes functionality which can associate several 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.

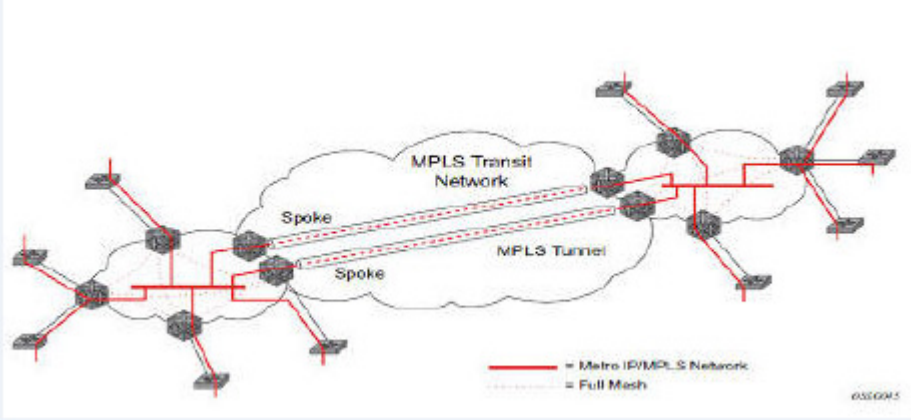

Figure 51: HVPLS with spoke redundancy

Spoke SDP Based Redundant Access (on 7210 SAS-K 2F4T6C)

This feature provides the ability to have a node deployed as MTUs (Multi-Tenant Unit Switches) multi-homed for VPLS to multiple routers deployed as PEs without requiring the use of mVPLS.

In the configuration example displayed in Figure 51, 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. If the primary and secondary spoke-SDPs have the same precedence value, the spoke-SDP with lower ID functions as the primary SDP.

The secondary spoke is in a blocking state (both on receive and transmit) if 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/sending traffic to/from the standby spoke. Optional revertive operation (with configurable switch-back delay) is applicable only when one of the spokes is configured with precedence of primary. If not, this action does not take place. 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 (on 7210 SAS-K 2F4T6C)

Note: MC-EP is not supported in 7210 SAS devices. This section serves to provide an example on how 7210 SAS devices can be used as a MTUs device in an MC-EP solution. In this solution, the 7750 SR routers provide the MC-EP functionality.

Inter-domain VPLS refers to a VPLS deployment where sites may be 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 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 pseudo wires to interconnect the VPLS instances located in the two domains.

MC-EP expands the single chassis endpoint based on active-standby pseudo wires for VPLS shown in Figure 52. In the solution depicted by the Figure 52, 7210 devices are used as MTUs.

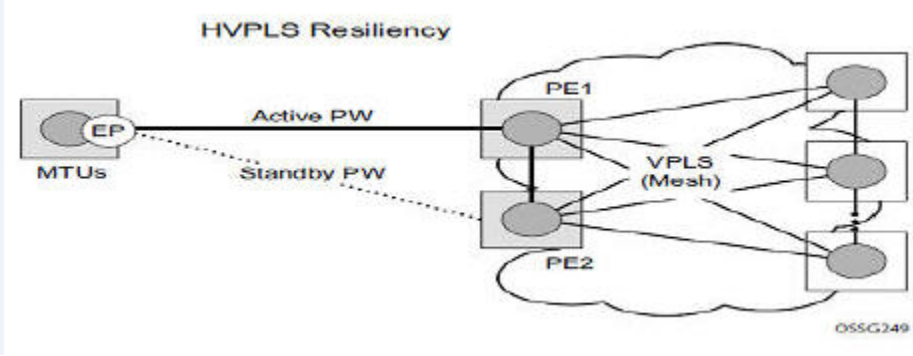

Figure 52: HVPLS Resiliency based on AS Pseudowires

The active-standby Pseudowires 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).

VPLS Access Redundancy when using MPLS uplinks on 7210 SAS-K 2F4T6C

A second application of hierarchical VPLS is using MTUs that are MPLS-enabled which must have spoke SDPs to the closest PE node. To protect against failure of the PE node, an MTU can be dual-homed.

Listed below are several mechanisms to resolve a loop in an access network where 7210s are used

- STP-based access, with or without mVPLS.
- Ethernet APS using G.8032.

STP-Based Redundant Access to VPLS using PWs on 7210 SAS-K 2F4T6C

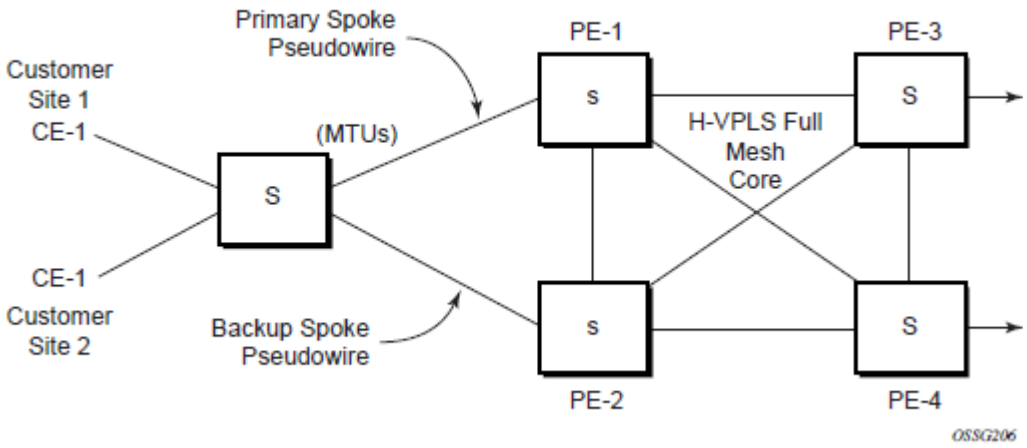

Figure 53: Dual Homed MTUs in Two-Tier Hierarchy H-VPLS

In configuration shown in Figure 53, STP is activated on the MTU and two PEs to resolve a potential loop.

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. Running STP in every VPLS in this topology is not efficient as 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. Thus, this applies 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.

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.

This is done by using “MAC-flush” messages defined by RFC 4762, *Virtual Private LAN Services Using LDP Signaling*. In the case where STP acts as a loop resolution mechanism, every Topology Change Notification (TCN) received in a context of STP instance is translated into an LDP-MAC address withdrawal message (also referred to as a MAC-flush message) requesting to clear all FDB entries except the ones learned from the originating PE. Such messages are sent to all PE peers connected through SDPs (mesh and spoke) in the context of VPLS service(s) managed by the given STP instance.

Redundant Access to VPLS Without STP and using MPLS uplinks (on 7210 SAS-K 2F4T6C)

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 to provide fast convergence.

Figure 51 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. Per 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) 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.

MAC Flush Message Processing in VPLS services with MPLS uplinks (on 7210 SAS-K 2F4T6C)

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 7210 SAS 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 learned from all other LDP peers except the originating PE. This type is also defined by RFC 4762 as an LDP MAC address withdrawal with an empty MAC address list.

Flush-all-mine message requests clearing all FDB entries learned from originating PE. This means that this message has exactly other effect then flush-all-but-mine message. This type is not included in RFC 4762 definition and it is implemented using vendor specific TLV.

The advantages and disadvantages of the individual types should be apparent from examples in the previous section. The description here focuses on summarizing actions taken on reception and conditions individual messages are generated.

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

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

The flush-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. Note, that enabling “send-flush-on-failure” the flag is blocked in VPLS service managed by mVPLS. This is to prevent that both messages are sent at the same time.

Dual Homing to a VPLS Service (on 7210 SAS-K 2F4T6C)

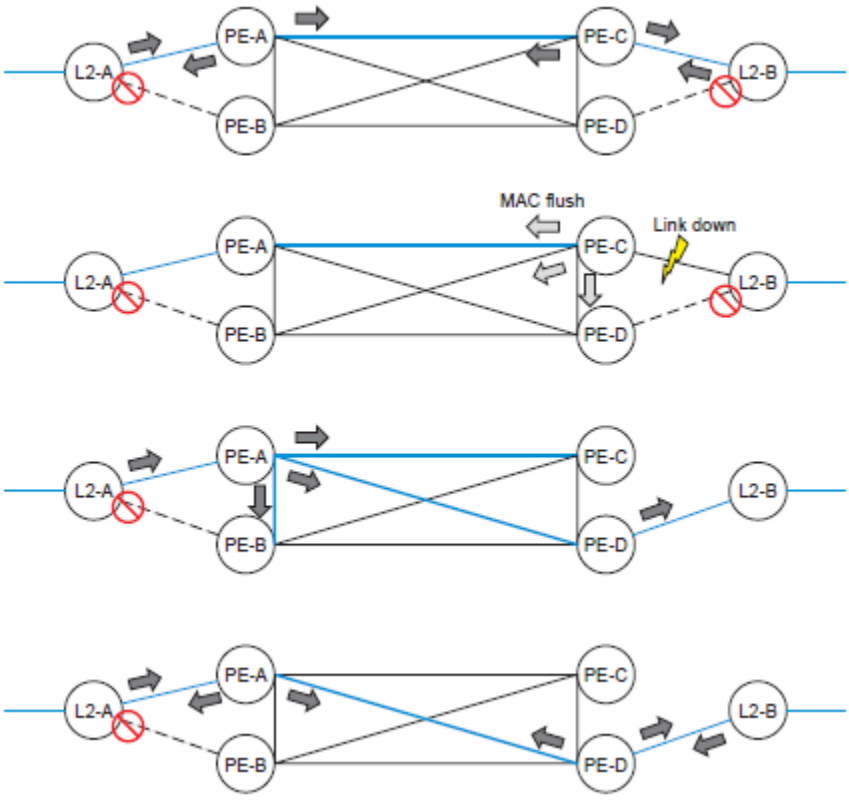

Figure 54: Dual Homed CE Connection to VPLS

Figure 54 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-l2vpnvpls-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 54) 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.

VPLS Service Considerations

This section describes various 7210 SAS 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 that pass Ethernet frames. The following SAP encapsulations are supported on the VPLS service on an access port:

- Ethernet null (Supported on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C)
- Ethernet Dot1q (Supported on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C)
- Ethernet QinQ (Supported on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C. It is not supported on 7210 SAS-E.)

The following encapsulations are supported on an access-uplink port:

- Ethernet QinQ (Supported on 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C)

VLAN Processing

The SAP encapsulation definition on Ethernet ingress ports defines which VLAN tags are used to determine the service that the packet belongs:

1. Null encapsulation defined on ingress — Any VLAN tags are ignored and the packet goes to a default service for the SAP.
2. Dot1q encapsulation defined on ingress — Only first VLAN tag is considered.
3. QinQ encapsulation defined on ingress— Both VLAN tags are considered.
   Note that the SAP can be defined with a wild-card for the inner label (for example, “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 100.1, 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.* definition.

In situations 2 and 3 above, traffic encapsulated with tags for which there is no definition are discarded.

IP Interface in a VPLS Service on 7210 SAS-E

NOTE: 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C do not support this feature.

The 7210 SAS-E supports host IP interface in a VPLS service. The IP interface is non-rout able and support only IP ping, IP trace-route and management ACLs. SAP ingress QoS policies applied on a SAP in a VPLS service are applicable to the packets destined to the IP interface.

IP Interface in a VPLS Service on 7210 SAS-E -Configuration Notes

1. The system uses the chassis MAC address by default for IP packets originating on the VPLS IP interface. The user is provided with an option to configure a MAC address.
2. The system adds the MAC address for a VPLS IP interface to the Layer 2 forwarding table. The Layer 2 forwarding table is a hash table. In case of hash collisions, the system removes the dynamically learned Layer 2 forwarding entry. If there are multiple such entries installed, the oldest entry is removed. Layer 2 packets will be forwarded as before. If the system is unable to locate a dynamic entry for removal, the IP interface is set to an operational down state. In such a case, the **show router vpls-management interface detail** command, displays **DownReason** as **HashCollision**. Additionally, the system generates a log message of the format MAC address **aa:bb:cc:dd:ee:ff configured for VPLS IP interface <interface-name> resulted in hash collision**. On receiving a log message, it is recommended that the user configure a new MAC address for the IP interface. If the user configured MAC address results in a hash collision, the CLI command fails, allowing the user to try with another MAC address.

BGP-AD VPLS

BGP-AD and Target LDP (T-LDP) Interaction on 7210 SAS-K 2F4T6C

NOTE: This feature is supported only on 7210 SAS-K 2F4T6C devices, it is not supported on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C.

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 triggers 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 55.

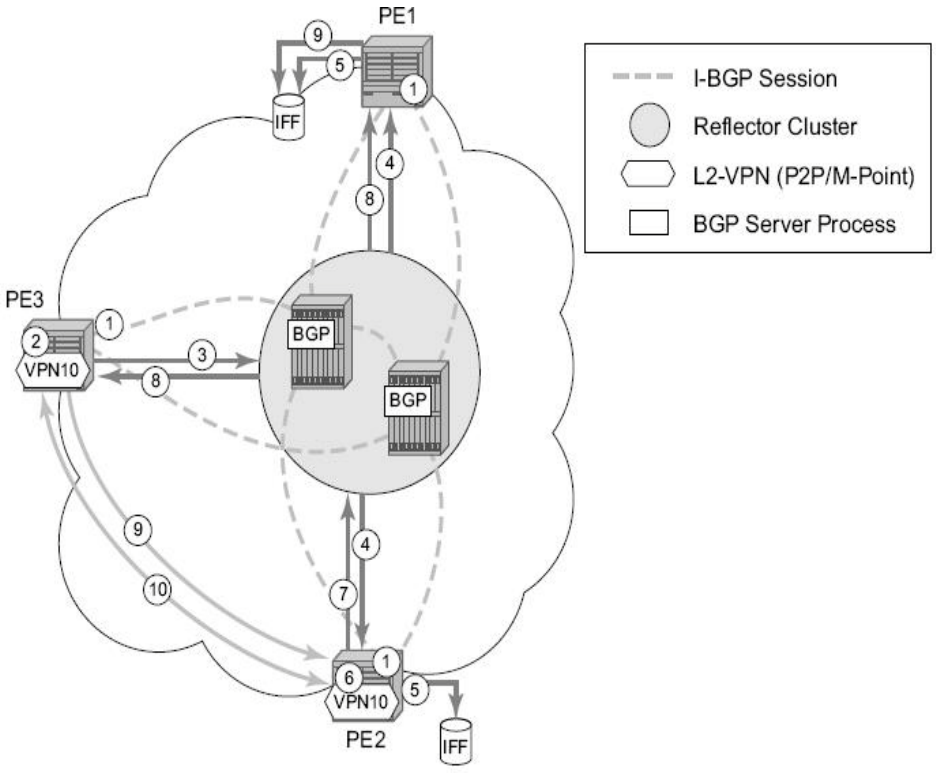

Figure 55: BGP-AD and T-LDP Interaction

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.

BGP-AD VPLS - SDP Usage

Service Access Points (SAP) are linked to transport tunnels using Service Distribution Points (SDP). The service architecture of the 7210 platform 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 7210 SAS OS MPLS Guide for more information about MPLS, LDP, and RSVP.

BGP-AD VPLS - 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.

BGP-AD VPLS - 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 remains down until a manually configured SDP becomes active.

When multiple SDPs exist between two endpoints, the tunnel manager selects the appropriate SDP. The algorithm preferred SDPs with the best (lower) metric. Should there be multiple SDPs with equal metrics, the operational state of the SDPs with the best metric is considered. If the operational state is the same, the SDP with the higher sdp-id is used. If an SDP with a preferred metric is found with an operational state that is not active, the tunnel manager flags it as ineligible and restarts the algorithm.

BGP-AD VPLS - 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 “pwtemplate” 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 are updated. The type of change made to the “pwtemplate” influences 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.

BGP-AD VPLS - Mixing Statically Configured and Auto-Discovered Pseudowires in a VPLS service

The services implementation allows for manually provisioned and auto-discovered pseudowire (SDP bindings) to co-exist 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.

BGP-AD VPLS - 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 co-exist 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 or 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 or standby pseudowires and to instantiate the related endpoint.

Routed VPLS on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

Routed VPLS (R-VPLS) allows a VPLS instance to be associated with an IES IP interface on 7210 SAS-D and 7210 SAS-K 2F2T1C. Only IPv4 addressing and forwarding is supported on these platforms.

Routed VPLS (R-VPLS) allows an VPLS instance to be associated with an IES or VPRN IP interface. Only IPv4 addressing and forwarding is supported on these platforms.

NOTE, the following sections on Routed VPLS, the reference to VPRN applies to platforms which support this feature. For now, it is supported only on 7210 SAS-K 2F4T6C devices.

Within an R-VPLS service, traffic with a destination MAC matching that of the associated IP interface is routed based on the IP forwarding table; all other traffic is forwarded based on the VPLS forwarding table.

On 7210 SAS-D and 7210 SAS-K 2F2T1C, RVPLS service can be associated with an IPv4 interface and supports only static routing. It is primarily designed for use of in-band management of the node. It allows for inband management of the 7210 nodes in a ring deployment using a single IPv4 subnet, reducing the number of IP subnets needed.

On 7210 SAS-K 2F4T6C, RVPLS service can be associated with an IPv4 interface and supports static routing and other routing protocols. It can be used to provide a service to the customer or for inband management of the node.

NOTE: 7210 SAS-E does not support Routed VPLS.

Routed VPLS - IES IP Interface Binding- 7210 SAS-D and 7210 SAS-K 2F2T1C

A standard IP interface within an existing IES or VPRN service context may be bound to a service name. 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. In other words, Routed VPLS allows the binding of IP interfaces in IES or VPRN services to be bound to VPLS services.

Routed 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. 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. In other words, Routed VPLS allows the binding of IP interfaces in IES or VPRN services to be bound to VPLS services.

Routed VPLS - 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 scans the existing IES or VPRN services for an IP interface bound to the specified service name. If an IP interface is found, the IP interface is 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 does 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 is adhered and become operational up. There is no need to toggle the shutdown or 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.

Routed VPLS - Service Binding Requirements

If the defined service name is created on the system, the system checks to ensure that the service type is VPLS. If the created service type is VPLS, the IP interface is eligible to enter the operationally upstate.

NOTE: On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, it is required that the user uses the R-VPLS tag while creating the VPLS service or the service cannot be bound to an IP interface. This is not required for 7210 SAS-D.

Routed VPLS - Bound Service Name Assignment

If a bound service name is assigned to a service within the system, the system first checks to ensure the service type is VPLS. Secondly the system ensures that the service is not already bound to another IP interface through the service name. If the service type is not VPLS or the service is already bound to another IP interface through the service ID, the service name assignment fails.

A single VPLS instance cannot be bound to two separate IP interfaces.

Routed VPLS - 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 any time. 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 scans for a VPLS service context using the name and takes the following actions:

- If the name is not currently in use by a service, the IP interface is 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 is 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 is placed in the operationally down: VPLS service allow-ip-int-binding flag not set state. There is no need to toggle the shutdown or 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 is eligible to be placed in the operationally up state depending on other operational criteria being met.

Routed VPLS - 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.

Routed VPLS - 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 is dependent upon the operational state of the VPLS service.

The VPLS service itself remains down until at least one virtual port (that is, SAP) is operational.

Routed VPLS - IP Interface MTU and Fragmentation on 7210 SAS-D

In 7210 SAS-D Access-Uplink mode, VPLS service MTU is not supported. The user must ensure that the port MTU is configured appropriately so that the largest packet traversing through any of the SAPs (virtual ports) of the VPLS service can be forwarded out of any of the SAPs. VPLS services do not support fragmentation and can discard packets larger than the configured port MTU.

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 port MTU of all the SAPs configured in the service. The port MTU excluding the Layer 2 Header and tags for all the ports which have SAPs configured in this VPLS service are considered and the minimum value among those are computed (which is called computed MTU). The operational value of the IP interface is set as follows:

- If the configured (administrative) value of IP MTU is greater than the computed MTU, then the operational IP MTU is set to the computed MTU.
- If the configured (administrative) value of IP MTU is lesser than or equal to the computed MTU, then operational IP MTU is set to the configured (administrative) value of IP MTU.

Routed VPLS - IP Interface MTU and Fragmentation on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

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. The service 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 SAPs are created in the system, the SAPs cannot become operational unless the configured port MTU minus the SAP service delineation overhead is greater than or equal to the configured VPLS service MTU. Thus, an operational SAP 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.

Routed VPLS - 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 returns the local IP address on the VPLS service of the remote router is returned. If the next-hop is not currently in the ARP cache, the system generates 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 dynamically populated ARP entries age out according to the ARP aging timer.

NOTE: In 7210 SAS-D, static ARP, entries cannot be used. Static ARP is supported on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices.

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 is flooded with all packets destined to that MAC (routed or bridged) to all SAPs within the VPLS service context. Once the MAC is known (populated in the VPLS FIB), all packets destined to the MAC (routed or bridged) is targeted to the specific SAP where the MAC has been learned. As with ARP entries, static MAC entries may be created in the VPLS FIB. Dynamically learned MAC addresses can 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 SAP than the current SAP in the FIB. In this case, the system automatically moves the MAC FIB entry to the new VPLS SAP.

NOTES:

- In 7210 SAS-D, whenever a MAC entry is removed from the VPLS FIB (either explicitly by the user or due to MAC aging or mac-move), ARP entries which match this MAC address is removed from the ARP cache. Though the VPLS FIB entries are not removed; an ARP entry ages out and is removed from the ARP cache. This restriction is not applicable to 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.
- In 7210 SAS-D, if If the VPLS FIB limit is reached and we are no longer able to learn new MAC address, ARP will also not be learnt. This restriction does not apply to 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

Routed VPLS - Specific ARP Cache Behavior

In typical routing behavior, the system uses the IP route table to select the egress interface, 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 (VPLS service SAPs).

The following tables describes how the ARP cache and MAC FIB entry states interact.

Table 16: Routing behavior in RVPLS and interaction ARP Cache and MAC FIB

ARP Cache Entry	MAC FIB Entry	Routing or System behavior
ARP Cache Miss (No Entry)	Known or Unknown	Triggers a request to control plane ARP processing module, to send out an ARP request, out of all the SAPs. (also known as virtual ports) of the VPLS instance.

Table 16: Routing behavior in RVPLS and interaction ARP Cache and MAC FIB

ARP Cache Entry	MAC FIB Entry	Routing or System behavior
ARP Cache Hit	Known	Forward to specific VPLS virtual port or SAP.
	Unknown	This behavior cannot happen typically in 7210 SAS-D, as and when a L2 entry is removed from the FDB, the matching MAC address is also removed from the ARP cache. On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, the packet is sent out of all the SAPs of the VPLS instance.

Routed VPLS - 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 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 service. SAPs within the service can be created on standard Ethernet ports.

Routed VPLS - 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 - VPLS Feature Support and 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). The following restrictions apply to both network mode and access-uplink mode unless called out separately:

- On 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C, spoke or mesh SDP bindings cannot be configured.
- On 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C, the VPLS service type cannot be M-VPLS.
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, the VPLS Service type must be 'r-vpls' and any other VPLS service is not allowed.
- MVR from Routed VPLS and to another SAP is not supported.
- Default QinQ SAPs is not supported in R-VPLS service.
- The “allow-ip-int-binding” command cannot be used in a VPLS service acting as the G.8032 control instance.
- IPv4 filters (ingress and egress) can be used with the R-VPLS SAPs. Additionally, IP ingress override filters are supported which affects the behavior of the IP filters attached to the R-VPLS SAPs. Please see below for more information about use of ingress override filters.
- MAC filters (ingress and egress) are not supported for use with R-VPLS SAPs.
- VPLS IP interface is not allowed in a R-VPLS service. The converse also holds.

- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, during creation of the VPLS service the keyword 'rvpls' must be used. It lets the software know that this is a VPLS service to which an IP interface will be associated.
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, the VPLS service can be configured either access SAP or Access-Uplink SAPs.
- On 7210 SAS-D, VPLS service can use the following 'svc-sap-type' values: any, dot1q-preserve and null-star. Only specific SAP combinations are allowed for a given svc-sap-type, except that default QinQ SAPs cannot be used in a R-VPLS service. The allowed SAP combinations are similar to that available in a plain VPLS service and is as given in the table above in the services Chapter (with the exception noted before).
- On 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, VPLS service can use the following 'svc-sap-type' values: any. Only specific SAP combinations are allowed for a given svc-sap-type, except that default QinQ SAPs cannot be used in a R-VPLS service. The allowed SAP combinations are similar to that available in a plain VPLS service and is as given in the table above in the services Chapter (with the exception noted before).
- G.8032 or mVPLS/STP based protection mechanism can be used with R-VPLS service. A separate G.8032 control instance or a separate mVPLS/STP instance needs to be used and the R-VPLS SAPs needs to be associated with these control instances such that the R-VPLS SAP's forwarding state is driven by the control instance protocols.
- IGMP snooping is not supported in a VPLS service. IP multicast is not supported in the R-VPLS service.
- On 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C, DHCP snooping is not supported for the SAPs configured in the routed VPLS service. Instead, DHCP relay can be enabled on the IES service associated with the routed VPLS service.

Routed VPLS - VPLS SAP Ingress IP Filter Override on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

When an IP Interface is attached to a VPLS service context, the VPLS SAP provisioned IP filter for ingress routed packets may be optionally overridden 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 packet types.

If a filter for a given packet type (IPv4) is not overridden, the SAP specified filter is applied to the packet (if defined).

The following tables lists ACL Lookup behavior with and without Ingress Override filter attached to an IES interface in a R-VPLS service:

Table 17: ACL Lookup behavior with Ingress Override filter attached to an IES interface in a R-VPLS service.

Type of traffic	SAP Ingress IPv4 Filter	SAP Egress IPv4 Filter	Ingress Override IPv4 Filter
Destination MAC != IES IP interface MAC	Yes	Yes	No
Destination MAC = IES IP interface MAC and Destination IP on same subnet as IES interface	No	No	Yes
Destination Mac = IES IP interface mac and destination IP not on same subnet as IES IP interface and route to destination IP does not exist	No	No	No for 7210 SAS-D and 7210 SAS-E devices Yes for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices
Destination Mac = IES IP interface mac and destination IP not on same subnet as IES IP interface and route to destination IP exists	No	No	Yes
Destination MAC = IES IP interface MAC and IP TTL = 1	No	No	No for 7210 SAS-D and 7210 SAS-E devices Yes for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices
Destination MAC = IES IP interface MAC and IPv4 packet with Options	No	No	No for 7210 SAS-D and 7210 SAS-E devices Yes for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices

Table 17: ACL Lookup behavior with Ingress Override filter attached to an IES interface in a R-VPLS service.

Type of traffic	SAP Ingress IPv4 Filter	SAP Egress IPv4 Filter	Ingress Override IPv4 Filter
Destination MAC = IES IP interface MAC and IPv4 Multicast packet	No	No	No for 7210 SAS-D and 7210 SAS-E devices Yes for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices

Table 18: ACL Lookup behavior without Ingress Override filter attached to an IES interface in a R-VPLS service

Type of traffic	SAP Ingress IPv4 Filter	SAP Egress IPv4 Filter
Destination MAC != IES IP interface MAC	Yes	Yes
Destination MAC = IES IP interface MAC and Destination IP on same subnet as IES IP interface	Yes	No
Destination Mac = IES IP interface mac and destination IP not on same subnet as IES IP interface and route to destination IP does not exist	No for 7210 SAS-D and 7210 SAS-E devices Yes for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices	No
Destination Mac = IES IP interface MAC and destination IP not on same subnet as IES IP interface and route to destination IP exists	Yes	No
Destination MAC = IES IP interface MAC and IP TTL = 1	No for 7210 SAS-D and 7210 SAS-E devices Yes for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices	No

Table 18: ACL Lookup behavior without Ingress Override filter attached to an IES interface in a R-VPLS service

Type of traffic	SAP Ingress IPv4 Filter	SAP Egress IPv4 Filter
Destination MAC = IES IP interface MAC and IPv4 packet with Options	No for 7210 SAS-D and 7210 SAS-E devices Yes for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices	No
Destination MAC = IES IP interface MAC and IPv4 Multicast packet	No for 7210 SAS-D and 7210 SAS-E devices Yes for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C devices	No

Routed VPLS - QoS Support for VPLS SAPs and IP interface

- SAP ingress classification (IPv4 and MAC criteria) is supported for SAPs configured in the service. SAP ingress policies cannot be associated with IES IP interface.
- On 7210 SAS-D, egress port based queuing and shaping are available. It is shared among all the SAPs on the port.
- On 7210 SAS-D, Port based Egress Marking is supported for both routed packets and bridged packets. The existing access egress QoS policy can be used for Dot1p marking and DSCP marking.
- On 7210 K-2F2T1C and 7210 SAS-K 2F4T6C, per SAP egress queuing, shaping and scheduling is available. Per SAP egress Dot1p marking is supported for both routed packet and bridged packets.
- On 7210 K-2F2T1C and 7210 SAS-K 2F4T6C, IES IP interface bound to routed VPLS services, IES IP interface on access SAPs and IES IP interface on Access-Uplink SAPs are designed for use with inband management of the node. Consequently, they share a common set of queues for CPU bound management traffic. All CPU bound traffic is policed to pre-defined rates before being queued into CPU queues for application processing. The system uses meters per application or a set of applications. It does not allocate meters per IP interface. The possibility of CPU overloading has been reduced by use of these mechanisms. Users must use appropriate security policies either on the node or in the network to ensure that this does not happen.

Routed VPLS - Routing Related Protocols on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

Routed VPLS is supported only in the base routing instance on 7210 SAS-D and 7210 SAS-K 2F2T1C. It is supported in both base routing instance and VPRN service on 7210 SAS-K 2F4T6C. Only IPv4 addressing support is available for IES interfaces associated with Routed VPLS service. The following lists the support available for routing protocols on IP interfaces bound to a VPLS service on different platforms.

Table 19: Routing Protocols on IP interfaces bound to a VPLS service.

Services	7210 SAS-D and 7210 SAS-K 2F2T1C	7210 SAS-K 2F4T6C
Static-routing	Supported	Supported
BGP	Not Supported	Not Supported
OSPF	Not Supported	Supported

Table 19: Routing Protocols on IP interfaces bound to a VPLS service.

Services	7210 SAS-D and 7210 SAS-K 2F2T1C	7210 SAS-K 2F4T6C
ISIS	Not Supported	Supported
BFD	Not Supported	Not Supported
VRRP	Not Supported	Not Supported
ARP and Proxy-Arp	ARP is supported	Both are supported
DHCP Relay (Note-1)	Supported	Supported

NOTE 1: DHCP relay can be configured for the IES interface associated with the Routed VPLS service. DHCP snooping cannot be configured on the VPLS SAPs in the routed VPLS Service.

Routed VPLS - Spanning Tree and Split Horizon

A routed VPLS context supports all spanning tree and split horizon capabilities that a non-routed VPLS service supports.

Routed VPLS support available and Caveats

Routed VPLS supported functionality and restrictions for both access-uplink and network mode is given below. The following is applicable to both the modes, unless called out explicitly.

- On 7210 SAS-D and 7210 SAS-K 2F2T1C, Routed VPLS can be bound to only IES IP interface. It cannot be bound to VPRN IP interface.
- On 7210 SAS-K 2F4T6C, routed VPLS can be bound to IES or VPRN IP interface.
- On 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C, only IPv4 addressing and forwarding is supported with Routed VPLS services. IPv6 addressing and forwarding is not supported.
- On 7210 SAS-D, static ARP cannot be configured with an IES IP interface associated with an R-VPLS, though static MAC can be configured in an R-VPLS service.
- On 7210 SAS-D and 7210 SAS-K 2F2T1C, only static routes are supported. No dynamic routing protocols are supported.
- On 7210 SAS-D, whenever a VPLS FIB entry is removed either due to user action, aging or mac-move, the corresponding ARP entry whose MAC address matches that of the MAC in the FIB is removed from the ARP cache.
- On 7210 SAS-D, multiple SAPs configured on the same port cannot be part of the same R-VPLS Service. In other words, a single service can only be configured with a single SAP on a given port. This restriction does not apply to 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. In other words, on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C multiple SAPs configured on the same port can be part of the same service.
- Service MTU configuration is not supported in the R-VPLS service. This does not apply to 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C. In other words, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C supports service MTU configuration for RVPLS.

Epipe Emulation using Dot1q VLAN range SAP in VPLS with G.8032

NOTE: This feature is supported on 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C with access-uplink ports. It is not supported on 7210 SAS-K 2F4T6C with network ports.

On the node where the service originates, in addition to the access dot1q range SAP, the service needs to be configured with access-uplink SAPs on the two G.8032 ring ports. G.8032 mechanism is used to for breaking the loop in the ring and VPLS service protection. The intermediate nodes on the ring needs to use VPLS service with access-uplink SAPs on the ring ports and use the same G.8032 instance for protection, as one is used for service protection on the originating node.

The Figure 56 shows how two business offices, served by an operator are connected in a ring network deployment using Dot1q range SAPs and a VPLS service with G.8032 for protection.

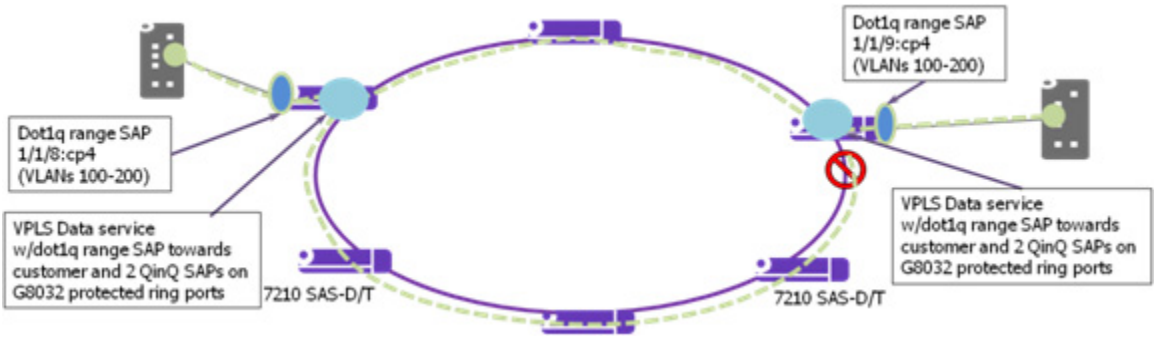

Figure 56: Epipe Emulation in a ring using VPLS with G.8032

The following are the requirements to provide for an Epipe service connectivity between two business sites:

- Transport all the VLANs used by the internal enterprise network of the businesses.
- Support high availability for the service between the business sites by protecting against failure of the links or nodes in the ring.

To achieve connectivity between two business sites in access-uplink/L2 mode is to configure SAPs for each of the individual VLANs used in the enterprise network in a VPLS service and use G.8032 for protection. The number of VLANs that was supported is limited by the number of SAPs supported on the platform.

The 7210 SAS platforms, currently support the use of Dot1q range SAPs with only Epipe services in either network/MPLS mode or access-uplink/L2 mode. Dot1q range SAPs allows operators to

transport a range of VLANs by providing similar service treatment (service treatment refers to forwarding decision along with encapsulation used, QoS and ACL processing, accounting, etc.) to all the VLANs configured in the range. It simplifies service configuration and allows operators to scale the number of VLANs that can be handled by the node. This took care of the need to support hundreds of VLANs using a single SAP or a small number of SAPs. When MPLS the mode is deployed in ring topology, operators have the option of using different redundancy mechanisms such as FRR, primary/secondary LSPs, Active/Standby PWs, to improve Epipe service availability. No such option is available to protect Epipe service in L2 mode when deployed in a ring topology. Additionally many operators prefer G.8032 based ring protection mechanism, since a single control instance on the ring can potentially protect all the VPLS services on the ring.

This feature allows operators to deploy Epipe services in a ring topology when using L2 mode, by emulating an Epipe service using a VPLS service with G.8032 protection and at the same time provides the benefits of using dot1q range SAPs. The user should ensure that the VPLS service is a point-to-point service. This is achieved by configuring a VPLS service with an access dot1q range SAP used at the customer hand-off on one node in the ring and an access dot1q range SAP in a customer hand-off of a VPLS service on another node (that is, at the other end of the Epipe), such that there are only two endpoints for the service in the network.

On the node where the service originates, in addition to the access dot1q range SAP, the service needs to be configured with access-uplink SAPs on the two G.8032 ring ports. G.8032 mechanism is used to for breaking the loop in the ring and VPLS service protection. The intermediate nodes on the ring needs to use VPLS service with access-uplink SAPs on the ring ports and use the same G.8032 instance for protection, as one is used for service protection on the originating node.

Epipe Emulation using Dot1q VLAN range SAP in VPLS with G.8032 -Configuration guidelines and restrictions

The VPLS service with dot1-range SAPs use svc-sap-type of dot1q-range and supports limited functionality in comparison to a normal VPLS service, The following paragraph provide more details of the feature functionality, configuration guidelines and restrictions:

- The user can define access dot1q range SAPs, which specifies a group of VLANs which receive similar service treatment, that is, forwarding behavior, SAP ingress QoS treatment and SAP (behavior like that available in Epipe service) and allows it to be configured in a VPLS service.
  - → On the node, where the service originates, in addition to the access dot1q range SAP, the service should be configured with Q1.* SAPs on the two G.8032 ring ports. The access or access-uplink Q1.*SAPs can be used, but the access-uplink SAPs are recommended for use.The user cannot configure any other SAPs in the same VPLS service.
  - → There is no special configuration required on intermediate nodes, that is, the ring nodes which do not terminate or originate the service. The nodes should be configured for providing transit VPLS service and the VPLS service must use the same G.8032 instance for protection as is used by the service on originating and terminating node.
  - → The Epipe service on 7210, currently does not check if the inner tag received on a Q1.* SAP is within the range of the configured VLANs. VPLS service too has the same behavior.
- Support for SAP Ingress QoS, Ingress and Egress ACLs, accounting, and other services, for dot1q range SAP configured in a VPLS service matches the support available in Epipe service.
- G.8032 mechanism is used for loop detection in ring network and service protection. A separate VPLS service representing the G.8032 control instance must be configured and the state should be associated with this service.
  - → Use of dot1q range SAPs to provide service on the interconnection node, in a G.8032 major-ring/sub-ring deployment, when using the virtual channel, is not supported. This restriction is not applicable when the interconnection node in a G.8032 major-ring/sub-ring is configured without a virtual channel.
- mVPLS/xSTP support is available for use with Q1.* SAP on the ring ports to break the loop. This is a add-on to the G.8032 support.
- Broadcast, Unknown Unicast and Multicast (BUM) traffic is flooded in the service.
- Learning is enabled on the service by default, to avoid the need to flood the service traffic out of one of the ring ports, after network MAC addresses are learnt. The user has an option to disable learning per service. Learning enable/disable per SAP is not supported.
- MAC limiting is available per service. MAC limiting per SAP is not supported.

- CFM OAM is supported. The support for UP MEPs on the dot1q range SAP in the service to be used for fault management and performance management using the CFM/Y.1731 OAM tools is available.
  - → On 7210 SAS-D, only UP MEP is allowed to be configured only on the dot1q VLAN range SAPs. CFM/Y.1731 tools can be used for trouble shooting and performance measurements. User must pick a VLAN value from the range of VLANs configured for the dot1-range SAP using the CLI command config>eth-cfm>domain>association>bridge-identifier VLAN and enable the use of using the CLI command primary-vlan-enable under the MEP CLI context. It is used as the VLAN tag in the packet header for all the CFM/Y.1731 messages sent out in the context of the UP MEP. This is not supported on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.
  - → Down MEPs and MIPs are not allowed to be configured.
  - → Fault propagation is not supported with UP MEPs for dot1q range SAP in access-uplink mode.
- CFM support is not available for SAPs on the ring ports.
- IGMP snooping and MVR is not supported.

Epipe Emulation using Dot1q VLAN range SAP in VPLS with G.8032 -Configuration guidelines and restrictions

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 344
- Common Configuration Tasks on page 346
  - → Configuring VPLS Components on page 346
    - – Creating a VPLS Service on page 347
    - – Configuring a VPLS SAP on page 356
- Configuring VPLS Redundancy on page 369
    - – Creating a Management VPLS for SAP Protection on page 369
    - – Configuring Load Balancing with Management VPLS on page 371
    - – Creating a Management VPLS for SAP Protection, using MPLS uplinks (on 7210 SAS-K 2F2T1C) on page 381
    - – Creating a Management VPLS for Spoke SDP Protection, using MPLS uplinks (7210 SAS-K 2F4T6C) on page 382
    - – Configuring Load Balancing with Management VPLS using MPLS uplinks (on 7210 SAS-K 2F4T6C) on page 383
    - – Configuring Selective MAC Flush, when using MPLS uplinks (on 7210 SAS-K 2F4T6C) on page 388
    - – Configuring AS Pseudo-wire in VPLS, when using MPLS uplinks (on 7210 SAS-K 2F4T6C) on page 388
- Configuring BGP Auto-Discovery on page 368
- Configuring AS Pseudo-wire in VPLS on page 423
- Service Management Tasks on page 390
  - → Modifying VPLS Service Parameters on page 390
  - → Deleting a Management VPLS on page 391
  - → Disabling a Management VPLS on page 392
  - → Deleting a VPLS Service on page 393
  - → Disabling a VPLS Service on page 393
  - → Re-Enabling a VPLS Service on page 393

Basic Configuration

The following fields require specific input (there are no defaults) to configure a basic VPLS service:

- Customer ID (refer to Configuring Customers on page 79)
- For a local service, configure two SAPs, specifying local access ports and encapsulation values.
- For a distributed service, configure a SAP and an SDP (only for 7210 SAS devices that support MPLS uplinks) for each far-end node.

The following example displays a sample configuration of a local VPLS service on ALA-1.

The **svc-sap-type** option must be specified at service creation. The default option is **null-star** on 7210 SAS-E and 'any' on other platforms. The following example displays a configuration of a local VPLS service with a null-star SAP type on ALA-1. This service comprises of a null access SAP and an Uplink LAG SAP.

```
*A:ALA-1>config>service# info
----------------------------------------------
...
       vpls 7 customer 7 create
            stp
                shutdown
            exit
            sap 1/1/21 create
            exit
            sap lag-1:700 create
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

The following example displays a sample configuration of a distributed VPLS service between ALA-1, ALA-2, and ALA-3.

```
vpls 100 customer 1 svc-sap-type any create
            shutdown
            bgp
            exit
            stp
                shutdown
            exit
        exit
        epipe 101 customer 1 svc-sap-type any create
            shutdown
        exit
----------------------------------------------
*A:ALA-1>config>service#
```

```
*A:ALA-2>config>service# info
----------------------------------------------
...
vpls 9000 customer 6 create
     description "This is a distributed VPLS."
     stp
     shutdown
     exit
     sap 1/1/5:16 create
     description "VPLS SAP"
     exit
     spoke-sdp 2:22 create
     exit
     no shutdown
exit
...
----------------------------------------------
*A:ALA-2>config>service#
```

Common Configuration Tasks

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

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)
3. Modify STP default parameters (optional) (see VPLS and Spanning Tree Protocol on page 299)
4. Enable service

Configuring VPLS Components

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

- Creating a VPLS Service on page 347
  - → Enabling MAC Move on page 350
- Configuring a VPLS SAP on page 356
  - → Local VPLS SAPs on page 356
  - → Configuring SAP-Specific STP Parameters on page 360
  - → STP SAP Operational States on page 364
- Configuring VPLS Redundancy on page 369

Creating a VPLS Service

Use the following CLI syntax to create a VPLS service (for 7210 SAS-E):

```
config>service# vpls service-id [customer customer-id] [cre-
   ate] [vpn <vpn-id>] [m-vpls] [svc-sap-type {null-
   star|dot1q|dot1q-preserve}] [customer-vid <vlan-id>]
```

The following example displays a VPLS configuration (for 7210 SAS-E):

```
*A:ALA-1>config>service>vpls# info
----------------------------------------------
...
        vpls 1000 customer 1 create
            description "This is a VPLS with NULL SAP"
            stp
                shutdown
            exit
            no shutdown
        exit
        vpls 2000 customer 6 svc-sap-type dot1q create
            description "This is a Distributed VPLS with DOT1Q SAP"
            stp
                shutdown
            exit
            no shutdown
        exit
        vpls 3000 customer 8 svc-sap-type dot1q-preserve customer-vid 300 create
            description "This is a VPLS with QinQ Uplink SAP"
            stp
                shutdown
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-1>config>service>vpls#
```

Use the following CLI syntax to create a VPLS service (for 7210 SAS-D):

CLI Syntax: `config>service# vpls service-id [customer customer-id] [cre-ate] [vpn <vpn-id>] [m-vpls] [svc-sap-type {null-star|dot1q-pre-serve|any }] [customer-vid <vlan-id>]`

The following example displays a VPLS configuration (for 7210 SAS-D):

```
*A:ALA-1>config>service>vpls# info
----------------------------------------------
    ...
         vpls 1000 customer 1 create
              description "This is a VPLS with NULL SAP"
              stp
                   shutdown
              exit
              no shutdown
```

```
            exit
            vpls 2000 customer 6 svc-sap-type any create
                 description "This is a Distributed VPLS with ANY SAP"
                 stp
                      shutdown
                 exit
                 no shutdown
            exit
            vpls 3000 customer 8 svc-sap-type dot1q-preserve customer-vid 300 create
                 description "This is a VPLS with QinQ Uplink SAP"
                 stp
                      shutdown
                 exit
                 no shutdown
            exit
        ...
-------------------------------------------
*A:ALA-1>config>service>vpls#
```

Use the following CLI syntax to create a VPLS service (for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C):

CLI Syntax:

```
config>service# vpls service-id [customer customer-id] [cre-ate] [vpn
<vpn-id>] [m-vpls] [svc-sap-type {dot1q-range |any }] [customer-vid
<vlan-id>]
```

The following example displays a VPLS configuration (for 7210 SAS-2F2T1C):

```
*A:sask_duti>config>service>vpls$ info
-------------------------------------------
            shutdown
            stp
                no shutdown
            exit
-------------------------------------------
*A:sask_duti>config>service>vpls$ info detail
-------------------------------------------
            shutdown
            no description
            service-mtu-check
            service-mtu 1514
            no def-mesh-vc-id
            no disable-learning
            no disable-aging
            no discard-unknown
            fdb-table-size 250
            fdb-table-high-wmark 95
            fdb-table-low-wmark 90
            no host-connectivity-verify
            no shcv-policy-ipv4
            no send-flush-on-failure
```

```
            local-age 300
            no mfib-table-size
            mfib-table-high-wmark 95
            mfib-table-low-wmark 90
            remote-age 900
            no propagate-mac-flush
            stp
                priority 32768
                hello-time 2
                forward-delay 15
                max-age 20
                hold-count 6
                mode rstp
                mst-max-hops 20
                no mst-name
                mst-revision 0
                no shutdown
            exit
            igmp-snooping
                shutdown
                query-interval 125
                robust-count 2
                report-src-ip 0.0.0.0
                no query-src-ip
                mvr
                    shutdown
                    no description
                    no group-policy
                exit
            exit
            mac-move
                move-frequency 2
                retry-timeout 10
                shutdown
            exit
            static-mac
            exit
            no vsd-domain
-----------------------------------------------
*A:sask_duti>config>service>vpls$
```

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 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
move-frequency *frequency*
retry-timeout *timeout*
no shutdown

The following example displays mac-move information.

```
*A:ALA-1# show service id 6 all
....
*A:ALA-1#
-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 1150               Mac Move          : Disabled
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Table Size        : 1000               Total Count       : 1000
Learned Count     : 1000               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      : True
===============================================================================
....
*A:ALA-1#
```

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 parameters, mentioned below, must be done in the constraints of the following two formula:

2 x (*Bridge_Forward_Delay* - 1.0 seconds) >= *Bridge_Max_Age*
Bridge_Max_Age >= 2 x (*Bridge_Hello0_Time* + 1.0 seconds)

The following STP parameters can be modified at VPLS level:

- Bridge STP Admin State on page 351
- Mode on page 352
- Bridge Priority on page 352
- Max Age on page 353
- Forward Delay on page 353
- Hello Time on page 354
- MST Instances on page 355
- MST Max Hops on page 355
- MST Name on page 355
- MST Revision on page 355

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 device. When STP on the VPLS is administratively enabled, but the administrative state of a SAP is down, BPDUs received on such a SAP are discarded.

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

Mode

To be compatible with the different iterations of the IEEE 802.1D standard, the 7210 SAS devices support several variants of the Spanning Tree protocol:

- `rstp` — Rapid Spanning Tree Protocol (RSTP) compliant with IEEE 802.1D-2004 - default mode.
- `dot1w` — Compliant with IEEE 802.1w.
- `comp-dot1w` — Operation as in RSTP but backwards compatible with IEEE 802.1w (this mode was introduced for interoperability with some MTU types).
- `mstp` — Compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q REV/D5.0-09/2005. This mode of operation is only supported in an mVPLS.

See section Spanning Tree Operating Modes on page 299 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.

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 363).

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 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 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 353.

CLI Syntax: `config>service>vpls service-id# stp`
`hello-time hello-time`
Range: 1 to 10 seconds
Default: 2 seconds
Restore Default: `no hello-time`

Hold Count

The **hold-count** command configures the peak number of BPDUs that can be transmitted in a period of one second.

CLI Syntax: `config>service>vpls service-id# stp`
`hold-count count-value`
Range: 1 to 10
Default: 6
Restore Default: `no hold-count`

MST Instances

You can create up to 15 MST-instances. They can range from 1 to 4094. By changing path-cost and priorities, you can make sure that each instance will form it's own tree within the region, thus making sure different VLANs follow different paths.

You can assign non overlapping VLAN ranges to each instance. VLANs 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 a VPLS SAP

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

Use the following CLI syntax to create:

- Local VPLS SAPs on page 356
- Distributed VPLS service with SAPs (on 7210 SAS-K 2F4T6C) on page 357
- Configuring Default QinQ SAPs to Pass all Traffic from Access to Access-uplink Port without any Tag Modifications on page 359
- Distributed VPLS SAPs on page 372
- Configuring Default QinQ SAPs to Pass all Traffic from Access to Access-uplink Port without any Tag Modifications on page 373
- Configuring SAP-Specific STP Parameters on page 360
- STP SAP Operational States on page 364
- Configuring VPLS SAPs with per service Split Horizon on page 381

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.

```
*A:ALA-1>config>service# info
----------------------------------------------
        vpls 1000 customer 1 create
            description "This is a Local VPLS with NULL SAP"
            stp
                shutdown
            exit
            sap 1/1/1 create
            exit
            sap 1/1/2 create
            exit
            sap 1/1/3:500.* create
            exit
            no shutdown
        exit
        vpls 2000 customer 6 create
            description "This is a Local VPLS with DOT1Q SAP"
            stp
                shutdown
            exit
            sap 1/1/4:100 create
            exit
            sap 1/1/5:200 create
```

```
            exit
            sap 1/1/3:900.* create
            exit
            no shutdown
        exit
        vpls 3000 customer 8 create
            description "This is a Local VPLS"
            stp
                shutdown
            exit
            sap 1/1/4:300 create
            exit
            sap 1/1/5:300 create
            exit
            sap 1/1/3:1200 create
            exit
            no shutdown
        exit
-------------------------------------------
*A:ALA-1>config>service#
```

Distributed VPLS service with SAPs (on 7210 SAS-K 2F4T6C)

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.

The following example displays a configuration of VPLS SAPs configured for ALA-1, ALA-2, and ALA-3.

```
*A:ALA-3>config>service# info
-------------------------------------------
vpls 100 customer 1 svc-sap-type any create
            shutdown
            no description
            service-mtu-check
            service-mtu 1514
            no def-mesh-vc-id
            no disable-learning
            no disable-aging
            no discard-unknown
            fdb-table-size 250
            fdb-table-high-wmark 95
            fdb-table-low-wmark 90
            no host-connectivity-verify
            no shcv-policy-ipv4
            no send-flush-on-failure
            local-age 300
            no mfib-table-size
            mfib-table-high-wmark 95
            mfib-table-low-wmark 90
            remote-age 900
            no propagate-mac-flush
            bgp
```

```
                no route-distinguisher
                no vsi-export
                no vsi-import
                no route-target
            exit

-------------------------------------------
*A:ALA-3>config>service#
```

Configuring Default QinQ SAPs to Pass all Traffic from Access to Access-uplink Port without any Tag Modifications

The following example displays the VPLS SAP configuration of Default QinQ SAPs:

```
ALA-1>config>service# vpls 9 customer 1 svc-sap-type null-star create
            shutdown
            stp
                shutdown
            exit
            sap 1/1/5:*.* create
                statistics
                    ingress
                        received-count
                    exit
                exit
            exit
            sap 1/1/6:*.* create
                statistics
                    ingress
                        received-count
                    exit
                exit
            exit
        exit
```

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 360
- SAP Virtual Port Number on page 361
- SAP Priority on page 361
- SAP Path Cost on page 362
- SAP Edge Port on page 362
- SAP Auto Edge on page 363
- SAP Link Type on page 363
- MST Instances on page 363

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 it's own virtual port number that is unique to every other SAP defined on the VPLS. The virtual port number is assigned at the time that the SAP is added to the VPLS.

Since the order in which SAPs are added to the VPLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance. To achieve consistency after a reboot, the virtual port number can be specified explicitly.

CLI Syntax: `config>service>vpls>sap# stp`
`port-num` *`number`*
Range: 1 — 2047
Default: (automatically generated)
Restore Default: `no port-num`

SAP Priority

SAP priority allows a configurable "tie breaking" parameter to be associated with a SAP. When configuration BPDUs are being received, the configured SAP priority will be used in some circumstances to determine whether a SAP will be designated or blocked.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16 bit value. In the latest STP standard (802.1D-2004) only the upper 4 bits of the port priority field are used to encode the SAP priority. The remaining 4 bits are used to extend the port ID field into a 12 bit virtual port number field. The virtual port number uniquely references a SAP within the STP instance. See SAP Virtual Port Number on page 361 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 incremental with the configured path cost for that SAP. When BPDUs are sent out other egress SAPs, the newly calculated root path cost is used.

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 7210 SAS E 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 353). 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 362).

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`

MST Instances

The SAP mst-instance command is used to create MST instances at the SAP level. MST instance at a SAP level can be created only if MST instances are defined at the service level.

The parameters that can be defined per instance are mst-path-cost and mst-port-priority.

- mst-path-cost — Specifies path-cost within a given MST instance. The path-cost is proportional to link speed.
- mst-port-priority — Specifies the port priority within a given MST instance.

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 364
- Operationally Discarding on page 364
- Operationally Learning on page 364
- Operationally Forwarding on page 365

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, 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 353.

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

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.

IEEE 802.1d (referred as dot1d) and Cisco's per VLAN Spanning Tree (PVST) BPDU encapsulations are supported on a per SAP basis. The STP is associated with a VPLS service like PVST is per VLAN. The difference between the two encapsulations is in the Ethernet and LLC framing and a type-length-value (TLV) field trailing the BPDU.The encapsulation format cannot be configured by the user, the system automatically determines the encapsulation format based on the BPDUs received on the port.

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

Table 20: SAP BPDU Encapsulation States

Field	dot1d encap-type null	dot1d encap-type dot1q	PVST encap-type null	PVST encap-type dot1q
Destination MAC	01:80:c2:00:00:00	01:80:c2:00:00:00	N/A	01:00:0c:cc:cc:cd
Source MAC	Sending Port MAC	Sending Port MAC	N/A	Sending Port MAC
EtherType	N/A	0x81 00	N/A	0x81 00
Dot1p and CFI	N/A	0xe	N/A	0xe
Dot1q	N/A	VPLS SAP ID	N/A	VPLS SAP 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	Standard 802.1d	Standard 802.1d	N/A	Standard 802.1d

Table 20: SAP BPDU Encapsulation States (Continued)

TLV: Type & Len	N/A	N/A	N/A	58 00 00 00 02
TLV: VLAN	N/A	N/A	N/A	VPLS SAP encap value
Padding	As Required	As Required	N/A	As Required

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. When a PVST-encapsulated BPDU is received, the SAP converts 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 are 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.When a dot1d-encapsulated BPDU is received, 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.

Configuring VPLS SAPs with per service Split Horizon

Note: Split Horizon group per service is supported only on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

To configure a VPLS service with a split horizon group, add the split-horizon-group parameter when creating the SAP. Traffic arriving on a SAP within a split horizon group will not be copied to other SAPs in the same split horizon group.

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

```
*A:ALA-1>config>service# info
----------------------------------------------
...
vpls 800 customer 6001 vpn 700 create
     description "VPLS with split horizon for DSL"
     stp
     shutdown
     exit
     sap 1/1/3:100 split-horizon-group DSL-group1 create
     description "SAP for residential bridging"
     exit
     sap 1/1/3:200 split-horizon-group DSL-group1 create
     description "SAP for residential bridging"
     exit
     split-horizon-group DSL-group1
     description "Split horizon group for DSL"
     exit
     no shutdown
exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

Configuring SDP Bindings (7210 SAS-K 2F4T6C)

Note: SDPs are supported only on 7210 SAS-K 2F4T6C. It is not supported on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C.

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 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 (7210 SAS-K 2F4T6C) on page 368).

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 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 data-link specific information of the VC.

Configuring VPLS Spoke SDPs with Split Horizon (7210 SAS-K 2F4T6C)

Note: Split Horizon group with spoke-SDPs is supported only on 7210 SAS-K12.

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 369
- Configuring Load Balancing with Management VPLS on page 371
- Creating a Management VPLS for SAP Protection, using MPLS uplinks (on 7210 SAS-K 2F2T1C) on page 381
- Creating a Management VPLS for Spoke SDP Protection, using MPLS uplinks (7210 SAS-K 2F4T6C) on page 382
- Configuring Load Balancing with Management VPLS using MPLS uplinks (on 7210 SAS-K 2F4T6C) on page 383
- Configuring Selective MAC Flush, when using MPLS uplinks (on 7210 SAS-K 2F4T6C) on page 388
- Configuring AS Pseudo-wire in VPLS, when using MPLS uplinks (on 7210 SAS-K 2F4T6C) on page 388

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 57. The tasks below should be performed on both nodes providing the protected VPLS service.

1. Create an access uplink SAPs to the peer node.
2. Create a management VPLS.
3. Define a SAP in the m-vpls on the port towards the 7210 SAS E. Note that the port must be dot1q. The SAP corresponds to the (stacked) VLAN on the 7210 SAS E in which STP is active.
4. Optionally modify STP parameters for load balancing.
5. Create access uplink SAPs in the m-vpls using the access uplink SAPs defined in Step 1.
6. Enable the management VPLS service and verify that it is operationally up.
7. Create a list of VLANs on the port managed by this management VPLS.
8. Create one or more user VPLS services with SAPs on VLANs in the range defined by Step 6.

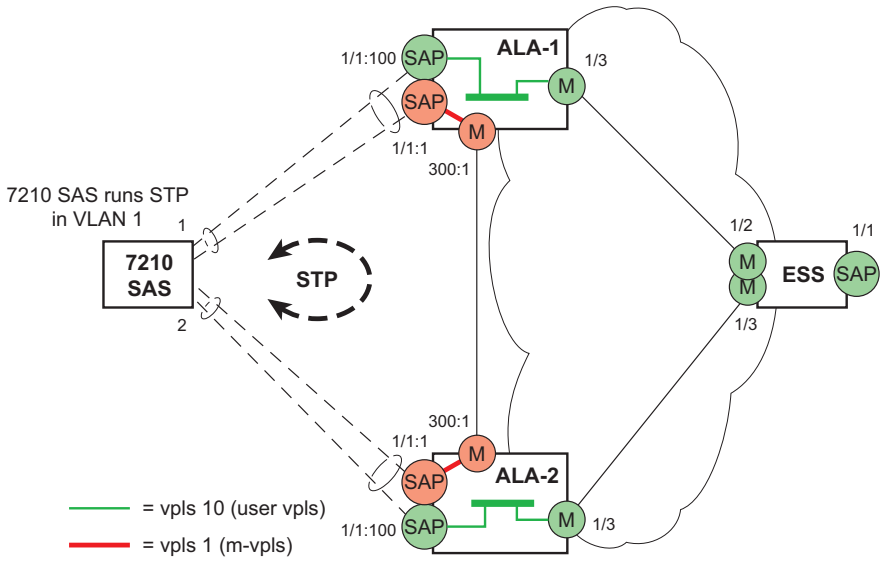

Figure 57: Example Configuration for Protected VPLS SAP

CLI Syntax: config>service# vpls *service-id* [customer *customer-id*] [create] [m-vpls]
description *description-string*
sap *sap-id* create
managed-vlan-list
range *vlan-range*
stp
no shutdown

The following example displays a VPLS configuration:

```
*A:ALA-1>config>service# info
----------------------------------------------
        vpls 2000 customer 6 m-vpls create
            stp
                no shutdown
            exit
            sap 1/1/1:100 create
            exit
            sap 1/1/2:200 create
            exit
            sap 1/1/3:300 create
               managed-vlan-list
                  range 1-50
            exit
            no shutdown
        exit
----------------------------------------------
*A:ALA-1>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 SAPs (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 SAPs. Load balancing can be achieved in SAP protection scenarios.

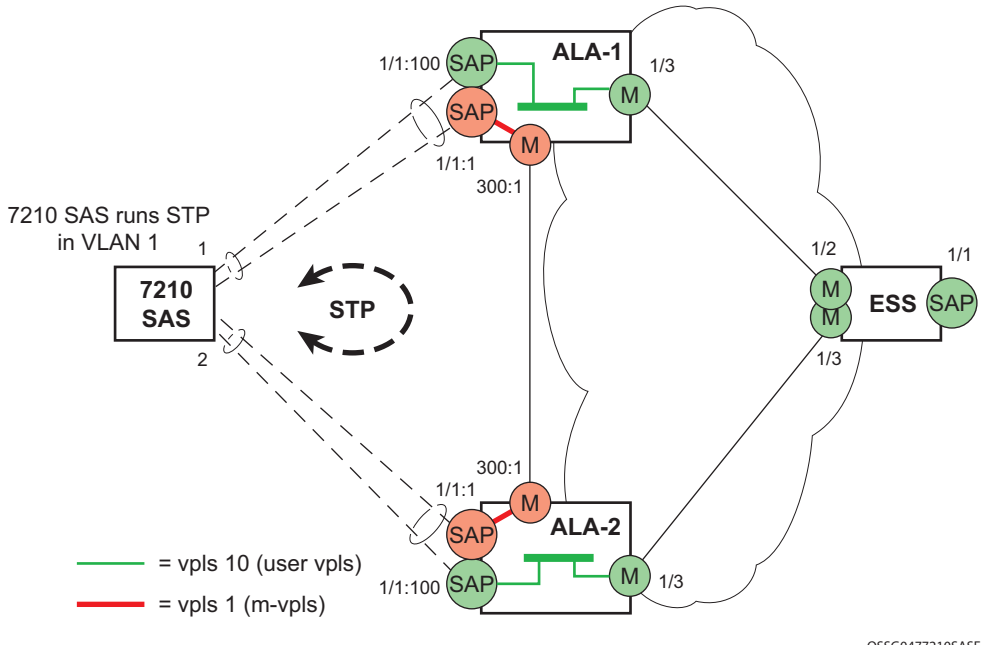

Figure 58: Example Configuration for Load Balancing Across with Management VPLS

Note: the STP path costs in each peer node should be reversed.

CLI Syntax: (for 7210 SAS-E) `config>service# vpls` *`service-id`* `[customer` *`customer-id`*`] [create] [m-vpls] [svc-sap-type {null-star | dot1q | dot1q-preserve}] [customer-vid vlan-id]`

```
description description-string
sap sap-id create
   managed-vlan-list
      range vlan-range
stp
no shutdown
```

CLI Syntax: (for 7210 SAS-D)`config>service# vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-star | any | dot1q-preserve}] [customer-vid vlan-id]`

```
description description-string
sap sap-id create
```

```
                managed-vlan-list
                   range vlan-range
                stp
                no shutdown
```

The following example displays a VPLS configuration (for 7210 SAS-E):

```
*A:ALA-1>config>service# info
----------------------------------------------
        vpls 100 customer 1 m-vpls svc-sap-type dot1q create
            stp
                no shutdown
            exit
            sap 1/1/2:100.* create
                managed-vlan-list
                    range 1-10
                exit
                stp
                    path-cost 1
                exit
            exit
            sap 1/1/3:500.* create
                shutdown
                managed-vlan-list
                    range 1-10
                exit
            exit
            no shutdown
        exit
        vpls 200 customer 6 m-vpls svc-sap-type dot1q create
            stp
                no shutdown
            exit
            sap 1/1/2:1000.* create
                managed-vlan-list
                    range 110-200
                exit
            exit
            sap 1/1/3:2000.* create
                managed-vlan-list
                    range 110-200
                exit
                stp
                    path-cost 1
                exit
            exit
            no shutdown
        exit
        vpls 101 customer 1 svc-sap-type dot1q create
            stp
                shutdown
            exit
            sap 1/1/1:100 create
            exit
            sap 1/1/2:1.* create
            exit
            sap 1/1/3:1.* create
            exit
            no shutdown
```

```
        exit
        vpls 201 customer 1 svc-sap-type dot1q create
            stp
                shutdown
            exit
            sap 1/1/1:200 create
            exit
            sap 1/1/2:110.* create
            exit
            sap 1/1/3:110.* create
            exit
            no shutdown
        exit
-------------------------------------------

*A:ALA-1>config>service#
```

Configuring a BGP-Auto-Discovery on 7210 SAS-K 2F4T6C

This section provides important information to explain the different configuration options used to populate the required BGP AD and generate the LDP generalized pseudowire-ID FEC fields. There are a large number of configuration options that are available with the this feature. Not all these configurations option are required to start using BGP AD. At the end of this section, it will be apparent that a very simple configuration will automatically generate the required values used by BGP and LDP. In most cases, deployments will provide full mesh connectivity between all nodes across a VPLS instance. However, capabilities are available to influence the topology and build hierarchies or hub and spoke models.

BGP-AD automatically creates SDP-bindings using a template to configure SDP-binding configuration parameters. L2-auto-bind is a command used to initiate a template that is used by BGP-AD for PW instantiation under related VPLS instances.

The template may be referenced in the “service vpls bgp-ad” object and used subsequently to instantiate PWs to a remote PE and VSI instance advertised through BGP Auto-Discovery. Changes to these dynamically created objects cannot be performed directly through CLI or SNMP. There are two possible methods to initiate the change:

- Configure a new “l2-auto-bind” association under service>vpls>bgp-ad. This method is used when the existing policy is used by multiple VPLS services and only one or a few require the change.
- Change the parameters of the current template. This method is used when a change in parameter is required for the majority of VPLS services that use the template.

Changes are not automatically propagated to the instantiated objects and must be done through one of two tool commands:

CLI Syntax: `tools>perform>service# eval-pw-template` *`policy-id`* `[allow-service-impact]`

CLI Syntax: `tools>perform>service>id# eval-pw-template` *`policy-id`* `[allow-service-impact]`

This command forces evaluation of changes that were made in the l2-auto-bind template indicated in the command. This command can be applied to an individual VPLS service or all VPLS services that reference the template if no service is specified.

The parameters are divided into three classes.

Class 1: Modified at create time only.

Class 2: Modified only when the object is administratively shutdown.

Class 3: No restrictions.

Parameters that fall into class 1 will destroy existing objects and re-create objects with the new values. Parameters in class 2 will momentarily shutdown the object, change the parameter, then re-enable the object. Class 3 can be changed without affecting the operational status of the objects of service.

For the l2-auto-bind template, the parameters are treated as follows:

Class 1: Adding or removing a split-horizon-group, switching between a manual and auto SDP

Class 2: Changing the vc-type {ether|vlan}

Class 3: All other changes.

The keyword allow-service-impact enables service impacting changes. If this keyword is not configured, an error message is generated if the parameter changes are service impacting.

Configuration Steps

Using Figure 59, 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.

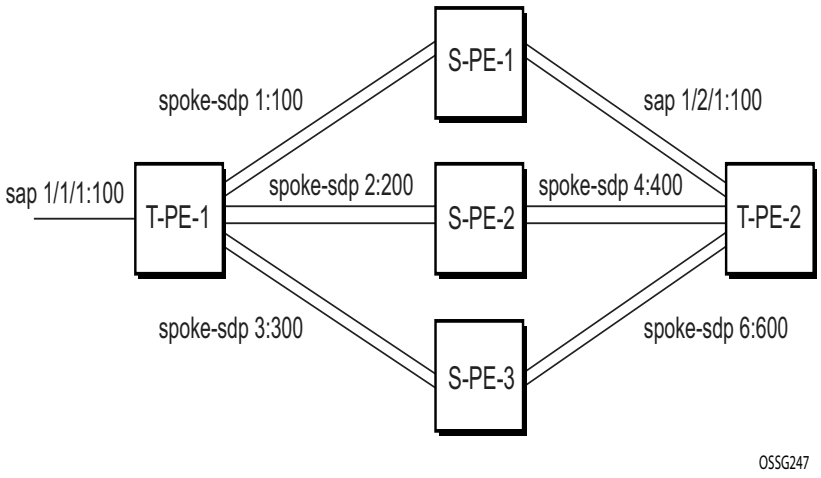

Figure 59: BGP AD Configuration Example

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] pw-template-bi* - Configure pw-template bind policy
  [no] route-target    - Configure route target
  [no] shutdown        - Administratively enable/disable BGP auto-discovery
       vpls-id         - Configure VPLS-ID
  [no] vsi-export      - VSI export route policies
       vsi-id          + Configure VSI-id
  [no] vsi-import      - VSI import route policies
```

Figure 60: 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.

```
*A:H-SASK# show service l2-route-table

===============================================================================
Services: L2 Route Information - Summary
===============================================================================
```

```
Svc Id     L2-Routes (RD-Prefix)                 Next Hop        Origin
               Sdp Bind Id                           PW Temp Id
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of L2 Route Entries: 0
===============================================================================

===============================================================================
Services: L2 Multi-Homing Route Information - Summary
===============================================================================
Svc Id     L2-Routes (RD-Prefix)        Next Hop        SiteId     State  DF
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of L2 Multi-Homing Route Entries: 0
===============================================================================

=======================================================
Services: L2 Bgp-Vpls Route Information - Summary
=======================================================
Svc Id     L2-Routes (RD)        Next Hop        Ve-Id
               Sdp Bind Id           PW Temp Id
-------------------------------------------------------
-------------------------------------------------------
No. of L2 Bgp-Vpls Route Entries: 0
=======================================================

=======================================================
Services: L2 Bgp-Vpws Route Information - Summary
=======================================================
Svc Id     L2-Routes (RD)        Next Hop        Ve-Id
               Sdp Bind Id           PW Temp Id
---------------------------------
```

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.

```
*A:H-SASK# show service sdp-using

===============================================================================
SDP Using
===============================================================================
SvcId      SdpId               Type   Far End              Opr   I.Label E.Label
                                      :GlobalId            State
-------------------------------------------------------------------------------
1          1:1                 Spok   4.4.4.4              Up    130769  130872
2          2:2                 Spok   4.4.4.4              Up    130770  130873
3          3:3                 Spok   4.4.4.4              Up    130771  130874
4          4:4                 Spok   4.4.4.4              Up    130772  130875
5          5:5                 Spok   4.4.4.4              Up    130773  130876
6          6:6                 Spok   4.4.4.4              Up    130774  130877
7          7:7                 Spok   4.4.4.4              Up    130775  130878
8          8:8                 Spok   4.4.4.4              Up    130776  130879
```

```
9          9:9                Spok    4.4.4.4              Up     130777  130880
10         10:10              Spok    4.4.4.4              Up     130778  130881
11         11:11              Spok    4.4.4.4              Up     130779  130882
12         12:12              Spok    4.4.4.4              Up     130780  130883
```

BGP AD advertises 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,)

BGP- AD- 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 TAII from the VSI-ID contained in the last 4 bytes of the received BGP NLRI.

Figure 61 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 61: BGP AD Triggering LDP Functions

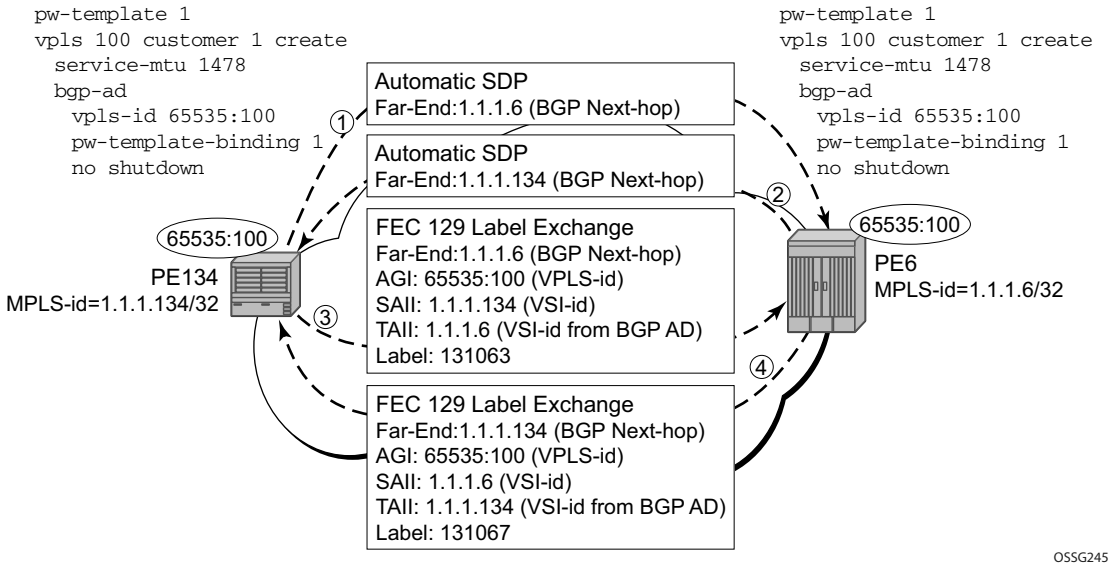

The first command shows the LDP peering relationships that have been established (Figure 62). The type of adjacency is displayed in the “Adj Type” column. In this case the type is “Both” meaning link and targeted sessions have been successfully established.

Figure 62: Show Router LDP Session Output

```
PERs6# show router ldp session
LDP Sessions
Peer LDP Id          Adj Type   State         Msg Sent  Msg Recv  Up Time
1.1.1.134:0          Both       Established   21482     21482     0d 15:38:44
No. of Sessions: 1
```

0988

The second command shows the specific LDP service label information broken up per FEC element type, 128 or 129, basis (Figure 63). The information for FEC element 129 includes the AGI, SAII and the TAII.

Figure 63: Show Router LDP Bindings FEC-Type Services

```
PERs6# show router ldp bindings fec-type services
LDP LSR ID: 1.1.1.6
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up,  D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, C - Cpipe Service
        TLV - (Type, Length: Value)
LDP Service FEC 128 Bindings
Type   VCId       SvcId      SDPId  Peer            IngLbl  EgrLbl  LMTU  RMTU
No Matching Entries Found

LDP Service FEC 129 Bindings
AGI                                SAII                    TAII
Type               SvcId      SDPId  Peer            IngLbl  EgrLbl  LMTU  RMTU
65535:100                          1.1.1.6                 1.1.1.134
V-Eth              100        17406  1.1.1.134       131063U 131067S 1464  1464
No. of FEC 129s: 1
```

0989

BGP-AD- Pseudowire Template

The pseudowire template is defined under the top-level service command (**config>service> pw-template**) and specifies whether to use an automatically generated SDP or manually configured SDP. It also provides the set of parameters required for establishing the pseudowire (SDP binding) as displayed below:

```
PERs6>config>service# pw-template 1 create
```

```
 -[no] pw-template <policy-id> [use-provisioned-sdp | prefer-provisioned-sdp]
<policy-id> : [1..2147483647]
<use-provisioned-s*> : keyword
<prefer-provisioned*> : keyword


[no] accounting-pol*    - Configure accounting-policy to be used
[no] auto-learn-mac*    - Enable/Disable automatic update of MAC protect list
[no] block-on-peer-*    - Enable/Disable block traffic on peer fault
[no] collect-stats      - Enable/disable statistics collection
[no] control word       - Enable/Disable the use of Control Word
[no] disable-aging      - Enable/disable aging of MAC addresses
[no] disable-learni*    - Enable/disable learning of new MAC addresses
[no] discard-unknow*    - Enable/disable discarding of frames with unknown source
                          MAC address
     egress             + Spoke SDP binding egress configuration
[no] force-qinq-vc-*    - Forces qinq-vc-type forwarding in the data-path
[no] force-vlan-vc-*    - Forces vlan-vc-type forwarding in the data-path
[no] hash-label         - Enable/disable use of hash-label
     igmp-snooping      + Configure IGMP snooping parameters
     in gr ess          + Spoke SDP binding ingress configuration
[no] l2pt-terminati*    - Configure L2PT termination on this spoke SDP
[no] limit-mac-move     - Configure mac move
[no] mac-pinning        - Enable/disable MAC address pinning on this spoke SDP
[no] max-nbr-mac-ad*    - Configure the maximum number of MAC entries in the FDB
                          from this SDP
[no] restrict-prote*    - Enable/disable protected src MAC restriction
[no] sdp-exclude        - Configure excluded SDP group
[no] sdp-include        - Configure included SDP group
[no] split-horizon-*    + Configure a split horizon group
     stp                + Configure STP parameters
     vc-type            - Configure VC type
[no] vlan-vc-tag        - Configure VLAN VC tag
```

A **pw-template-binding** command configured within the VPLS service under the **bgp-ad** sub-command is a pointer to the pw-template that should be used. If a VPLS service does not specify an import-rt list, then that binding applies to all route targets accepted by that VPLS. The **pw-template-bind** command can select a different template on a per import-rt basis. It is also possible to specify specific pw-templates for some route targets with a VPLS service and use the single **pw-template-binding** command to address all unspecified but accepted imported targets.

Figure 64: PW-Template-Binding CLI Syntax

```
PERs6>config>service>vpls>bgp-ad# pw-template-binding
  - pw-template-binding <policy-id> [split-hozion-group <group-name>] [import-
rt
    {ext-community, ...(upto 5 max)}]
  - no pw-template-binding <policy-id>

 <policy-id>          : [1..2147483647]
 <group-name>         : [32 chars max]
 <ext-community>      : target:{<ip-addr:comm-val>|<as-number:ext-comm-val>}
                        ip-addr      - a.b.c.d
                        comm-val     - [0..65535]
                        as-number    - [1..65535]
                        ext-comm-val - [0..4294967295]
```

0990

It is important understand the significance of the split-horizon-group used by the pw-template. Traditionally, when a VPLS instance was manually created using mesh-sdp bindings, these were automatically placed in a common split-horizon-group to prevent forwarding between the pseudowire in the VPLS instances. This prevents loops that would have otherwise occurred in the Layer 2 service. When automatically discovering VPLS service using BGP AD the service provider has the option of associating the auto-discovered pseudowire with a split-horizon group to control the forwarding between pseudowires.

Configuring a VPLS Management Interface

Use the following CLI syntax to create a VPLS management interface.

```
CLI Syntax: config>service>vpls# interface ip-int-name
            address ip-address[/mask] [netmask]
            arp-timeout seconds
            description description-string
            mac ieee-address
            no shutdown
            static-arp ip-address ieee-address
```

The following example shows the configuration.

```
A:ALA-49>config>service>vpls>interface# info detail
----------------------------------------------
            no description
            mac 14:31:ff:00:00:00
            address 123.231.10.10/24
            no arp-timeout
            no shutdown
----------------------------------------------
A:ALA-49>config>service>vpls>interface#
```

Creating a Management VPLS for SAP Protection, using MPLS uplinks (on 7210 SAS-K 2F2T1C)

NOTE: This is supported on 7210 SAS-K 2F4T6C.

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. The tasks below should be performed on both nodes providing the protected VPLS service.

Before configuring a management VPLS, first read Configuring VPLS Redundancy on page 369 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 7210 SAS M. Note that the port must be dot1q. The SAP corresponds to the (stacked) VLAN on the 7210 SAS M in which STP is active.
4. Optionally modify STP parameters for load balancing (see Configuring Load Balancing with Management VPLS on page 414).
5. Create an SDP in the m-vpls using the SDP defined in Step 1. Ensure that this SDP runs over a protected LSP.
6. Enable the management VPLS service and verify that it is operationally up.
7. Create a list of VLANs on the port managed by this management VPLS.
8. Create one or more user VPLS services with SAPs on VLANs in the range defined by Step 7.

CLI Syntax: `config>service# vpls service-id [customer customer-id] [create] [m-vpls]`

```
description description-string
sap sap-id create
   managed-vlan-list
      range vlan-range
         stp
         no shutdown
```

The following example displays a VPLS configuration:

```
*A:ALA-1>config>service# info
```

```
-----------------------------------------------
        vpls 2000 customer 6 m-vpls create
            stp
                no shutdown
            exit
            sap 1/1/1:100 create
            exit
            sap 1/1/2:200 create
            exit
            sap 1/1/3:300 create
               managed-vlan-list
                  range 1-50
            exit
            no shutdown
        exit
-----------------------------------------------
*A:ALA-1>config>service#
```

Creating a Management VPLS for Spoke SDP Protection, using MPLS uplinks (7210 SAS-K 2F4T6C)

NOTE: This is supported on 7210 SAS-K 2F4T6C.

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 59. The tasks below should be performed on all four nodes providing the protected VPLS service.

Before configuring a management VPLS, please first read Configuring VPLS Redundancy on page 369 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-spoke 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.
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.

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
                 spoke-sdp sdp-id:vc-id create
                 stp
                 no shutdown
```

The following example displays a VPLS configuration:

```
*A:ALA-A1>config>service# info
----------------------------------------------
...
sdp 1 mpls create
            shutdown
            no description
            signaling tldp
            no ldp
            no sr-isis
            no sr-ospf
            no bgp-tunnel
            no path-mtu
            no adv-mtu-override
            keep-alive
                shutdown
                hello-time 10
                hold-down-time 10
                max-drop-count 3
                timeout 5
                no message-length
            exit

...
----------------------------------------------
*A:ALA-A1>config>service#
```

Configuring Load Balancing with Management VPLS using MPLS uplinks (on 7210 SAS-K 2F4T6C)

NOTE: This is supported on 7210 SAS-K 2F4T6C.

This is done by creating two management VPLS instances, where both instances have different active 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 spokes.

Load balancing can be achieved in both the SAP protection and spoke SDP protection scenarios.

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
                spoke-sdp sdp-id:vc-id create
                stp
                path-cost
                stp
                no shutdown
```

This following output shows example configurations for load balancing across two protected VPLS spoke SDPs:

The configuration on ALA-A (7210 SAS) is shown below.

```
# MVPLS 100 configs
*A:ALA-A# configure service vpls 100
*A:ALA-A>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 100"
     stp
     no shutdown
     exit
     sap lag-3:100 create
     description "Default sap description for service id 100"
     managed-vlan-list
          range 101-110
     exit
     exit
     spoke-sdp 1201:100 create
     stp
     path-cost 100
     exit
     exit
     spoke-sdp 1401:100 create
     exit
     no shutdown
----------------------------------------------
*A:ALA-A>config>service>vpls#

# UVPLS 101 configs
*A:ALA-A>config>service# vpls 101
```

```
*A:ALA-A>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 101"
     sap lag-3:101 create
     description "Default sap description for service id 101"
     exit
     spoke-sdp 1201:101 create
     exit
     spoke-sdp 1401:101 create
     exit
     no shutdown
----------------------------------------------
*A:ALA-A>config>service>vpls#

# MVPLS 200 configs
*A:ALA-A# configure service vpls 200
*A:ALA-A>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 200"
     stp
     no shutdown
     exit
     sap lag-3:200 create
     description "Default sap description for service id 200"
     managed-vlan-list
     range 201-210
     exit
     exit
     spoke-sdp 1202:200 create
     exit
     spoke-sdp 1402:200 create
     stp
     path-cost 100
     exit
     exit
     no shutdown
----------------------------------------------
*A:ALA-A>config>service>vpls#

# UVPLS 201 configs
*A:ALA-A>config>service# vpls 201
*A:ALA-A>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 201"
     sap lag-3:201 create
     description "Default sap description for service id 201"
     exit
     spoke-sdp 1202:201 create
     exit
     spoke-sdp 1402:201 create
     exit
     no shutdown
----------------------------------------------
*A:ALA-A>config>service>vpls# exit all
```

The configuration on ALA-B (7210), the top left node is shown below. It is configured such that it becomes the root bridge for MVPLS 100 and MVPLS 200.

```
# MVPLS 100 configs
*A:ALA-B# configure service vpls 100
*A:ALA-B>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 100"
     stp
     priority 0
     no shutdown
     exit
     spoke-sdp 1201:100 create
     exit
     spoke-sdp 2301:100 create
     exit
     no shutdown
----------------------------------------------
*A:ALA-B>config>service>vpls#

# UVPLS 101 configs
*A:ALA-B>config>service# vpls 101
*A:ALA-B>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 101"
     spoke-sdp 1201:101 create
     exit
     spoke-sdp 2301:101 create
     exit
     no shutdown
----------------------------------------------
*A:ALA-B>config>service>vpls#

# MVPLS 200 configs
*A:ALA-B# configure service vpls 200
*A:ALA-B>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 200"
     stp
     priority 0
     no shutdown
     exit
     spoke-sdp 1202:200 create
     exit
     spoke-sdp 2302:200 create
     exit
     no shutdown
----------------------------------------------
*A:ALA-B>config>service>vpls#

# UVPLS 201 configs
*A:ALA-B>config>service# vpls 201
*A:ALA-B>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 201"
     spoke-sdp 1202:201 create
     exit
     spoke-sdp 2302:201 create
     exit
     no shutdown
----------------------------------------------
*A:ALA-B>config>service>vpls#
```

The configuration on ALA-C (7210), the top right node is shown below.

```
# MVPLS 100 configs
*A:ALA-C# configure service vpls 100
*A:ALA-C>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 100"
     stp
     priority 4096
     no shutdown
     exit
     spoke-sdp 1401:100 create
     exit
     spoke-sdp 2301:100 create
     exit
     no shutdown
----------------------------------------------
*A:ALA-C>config>service>vpls#

# UVPLS 101 configs
*A:ALA-C>config>service# vpls 101
*A:ALA-C>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 101"
     spoke-sdp 1401:101 create
     exit
     spoke-sdp 2301:101 create
     exit
     no shutdown
----------------------------------------------
*A:ALA-C>config>service>vpls#

# MVPLS 200 configs
*A:ALA-C# configure service vpls 200
*A:ALA-C>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 200"
     stp
     priority 4096
     no shutdown
     exit
     spoke-sdp 1402:200 create
     exit
     spoke-sdp 2302:200 create
     exit
     no shutdown
----------------------------------------------
*A:ALA-C>config>service>vpls#

# UVPLS 201 configs
*A:ALA-C>config>service# vpls 201
*A:ALA-C>config>service>vpls# info
----------------------------------------------
     description "Default tls description for service id 201"
     spoke-sdp 1402:201 create
     exit
     spoke-sdp 2302:201 create
     exit
```

```
    no shutdown
----------------------------------------------
*A:ALA-C>config>service>vpls#
```

Configuring Selective MAC Flush, when using MPLS uplinks (on 7210 SAS-K 2F4T6C)

NOTE: This is supported on 7210 SAS-K 2F4T6C.

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 AS Pseudo-wire in VPLS, when using MPLS uplinks (on 7210 SAS-K 2F4T6C)

NOTE: This is supported on 7210 SAS-K 2F4T6C.

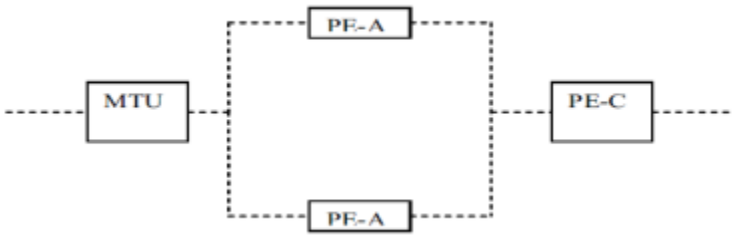

Figure 65: Sample Topology-AS Pseudo-wire in VPLS

In Figure 65, Pseudo-wire is configured on MTU. A sample configuration on the MTU is listed below:

```
*A:MTU>config>service>vpls>endpoint# back
*A:MTU>config>service>vpls# info
----------------------------------------------
     send-flush-on-failure
     stp
     shutdown
     exit
     endpoint "vpls1" create
          description "vpls1_endpoint"
          revert-time 60
          ignore-standby-signaling
          no suppress-standby-signaling
          block-on-mesh-failure
     exit
     sap 1/1/3 create
     exit
     spoke-sdp 301:1 endpoint "vpls1" create
     stp
     shutdown
     exit
          block-on-mesh-failure
     exit
     spoke-sdp 302:1 endpoint "vpls1" create
     stp
     shutdown
     exit
          block-on-mesh-failure
     exit
     no shutdown
----------------------------------------------
*A:MTU>config>service>vpls#
```

Service Management Tasks

This section discusses the following service management tasks:

- Modifying VPLS Service Parameters on page 390
- Modifying Management VPLS Parameters on page 391
- Deleting a Management VPLS on page 391
- Disabling a Management VPLS on page 392
- Deleting a VPLS Service on page 393
- Disabling a VPLS Service on page 393
- Re-Enabling a VPLS Service on page 393

Modifying VPLS Service Parameters

You can change existing service parameters. The changes are applied immediately. To display a list of services, use the **show service service-using vpls** command. Enter the parameter such as description SAP and then enter the new information.

The following displays a modified VPLS configuration.

```
*A:ALA-1>config>service>vpls# info
----------------------------------------------
            description "This is a different description."
            disable-learning
            disable-aging
            discard-unknown
            local-age 500
            stp
                shutdown
            exit
            sap 1/1/5:22 create
                description "VPLS SAP"
            exit
            exit
            no shutdown
----------------------------------------------
*A:ALA-1>config>service>vpls#
```

Modifying Management VPLS Parameters

To modify the range of VLANs on an access port 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 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] 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.

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 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
       sap sap-id
       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 shutdown.

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

- VPLS Service Configuration Commands on page 396
- Global Commands for 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C on page 396
- Global Commands for 7210 SAS-K 2F4T6C on page 397
- VPLS service - IP interface (Host only) commands on 7210 SAS-E on page 399
- VPLS service xSTP commands on page 399
- VPLS Service SAP DHCP Snooping commands on page 400
- SAP Commands on page 400
- VPLS SAP Configuration - QoS and Filter commands for 7210 SAS-D and 7210 SAS-E on page 402
- VPLS SAP Configuration - QoS and Filter commands for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C on page 402
- VPLS Service SAP IGMP Snooping and MVR commands on page 404
- VPLS SAP Meter Override Commands on page 406
- SAP Commands on page 400
- VPLS SAP Configuration - QoS and Filter commands for 7210 SAS-D and 7210 SAS-E on page 402
- VPLS SAP Configuration - QoS and Filter commands for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C on page 402
- VPLS Service SAP IGMP Snooping and MVR commands on page 404
- VPLS SAP Meter Override Commands on page 406
- VPLS service SAP xSTP commands on page 406
- VPLS Mesh SDP Commands for 7210 SAS-K 2F4T6C on page 408
- VPLS Spoke SDP Commands for 7210 SAS-K 2F4T6C on page 410
- Routed VPLS Commands applicable only to 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C on page 412
- Show Commands on page 413
- Clear Commands on page 414
- Debug Commands on page 414

VPLS Service Configuration Commands

Global Commands for 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

```
config
    — service
        — vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type {null-star | dot1q | dot1q-preserve}] [customer-vid vlan-id] (7210 SAS-E)
        — vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-star|dot1q-preserve|dot1q-range|any}] [customer-vid vlan-id] (7210 SAS-D)
        — vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {any|dot1q-range}] [r-vpls] (7210 SAS-K 2F2T1C)
        — no vpls service-id
            — description description-string
            — no description
            — [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
            — no local-age
            — [no] mac-move
                — move-frequency frequency
                — no move-frequency
                — retry-timeout timeout
                — no retry-timeout
                — [no] shutdown
            — remote-age aging-timer
            — no remote-age
            — service-mtu octets
            — no service-mtu
```

Global Commands for 7210 SAS-K 2F4T6C

```
config
    — service
        — vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {any|dot1q-range}]
          [r-vpls]
        — no vpls service-id
            — 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
                — 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-ad
                — [no] shutdown
                — vpls-id vpls-id
                — vsi-id
                    — prefix low-order-vsi-id
                    — no prefix
            — description description-string
            — no description
            — [no] disable-aging
            — [no] disable-learning
            — [no] discard-unknown
            — endpoint endpoint-name [create]
            — no endpoint
                — block-on-mesh-failure
                — [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
                — revert-time revert-time | infinite
                — no revert-time
                — static-mac ieee-address [create]
                — no static-mac
                — [no] suppress-standby-signaling
            — [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
            — no local-age
            — [no] mac-move
                — move-frequency frequency
                — no move-frequency
                — retry-timeout timeout
```

```
            — no retry-timeout
            — [no] shutdown
— [no] propagate-mac-flush
— remote-age aging-timer
— no remote-age
— [no] send-flush-on-failure
— service-mtu octets
— no service-mtu
— no service-name
— [no] shutdown
— split-horizon-group group-name [create]
— description description-string
— no description
```

VPLS service - IP interface (Host only) commands on 7210 SAS-E

NOTE: This is not a Routed VPLS IP interface.

config
- **service**
  - **vpls** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**m-vpls**] [**svc-sap-type** {**null-star** | **dot1q** | **dot1q-preserve**}] [**customer-vid** *vlan-id*] (for 7210 SAS-E)
  - **no vpls** *service-id*
    - [**no**] **interface** *ip-int-name* [**create**]
      - **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** *ip-address ieee-address*
      - **no static-arp** *ip-address* [*ieee-address*]

VPLS service xSTP commands

config
- **service**
  - **vpls** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**m-vpls**] [**svc-sap-type** {**null-star** | **dot1q** | **dot1q-preserve**}] [**customer-vid** *vlan-id*] (for 7210 SAS-E)
  - **vpls** *service-id* [customer *customer-id*] [**create**] [**m-vpls**] [**svc-sap-type {null-star|dot1q-preserve|dot1q-range|any**}] [**customer-vid** *vlan-id*] (for 7210 SAS-D)
  - **vpls** *service-id* **[customer** *customer-id***] [create] [m-vpls] [svc-sap-type {any|dot1q-range}] [r-vpls] (for 7210 SAS-K)**
  -
  - **no vpls** *service-id*
    - **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-age*
      - **no max-age**
      - **The max info age for the STP instance in seconds. Allowed values are integers in the range 6 to 40. mode** {**rstp** | **comp-dot1w** | **dot1w** | **mstp**}
      - **no The max info age for the STP instance in seconds. Allowed values are integers in the range 6 to 40. mode**
      - [**no**] **mst-instance** *mst-inst-number*
        - **mst-port-priority** *bridge-priority*
        - **no mst-port-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 Service SAP DHCP Snooping commands

config
- **service**
  - **vpls** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**m-vpls**] [**svc-sap-type** {**null-star** | **dot1q** | **dot1q-preserve**}] [**customer-vid** *vlan-id*] (for 7210 SAS-E)
  - **vpls** *service-id* [customer *customer-id*] [**create**] [**m-vpls**] [**svc-sap-type {null-star|dot1q-preserve|dot1q-range|any**}] [**customer-vid** *vlan-id*] (for 7210 SAS-D)
  - **vpls** *service-id* **[customer** *customer-id***] [create] [m-vpls] [svc-sap-type {any|dot1q-range}] [r-vpls] (for 7210 SAS-K)**
  - **no vpls** *service-id*
    - **sap** *sap-id* [**split-horizon-group** *group-name*] [**eth-ring** *ring-index*]
    - **no sap** *sap-id*
      - **dhcp**
        - **description** *description-string*
        - **no description**
        - [**no**] **option**
          - **action** [*dhcp-action*]
          - **no action**
          - **[no] 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**
        - [**no**] **shutdown**
        - [**no**] **snoop**

SAP Commands

config
- **service**
  - **vpls** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**m-vpls**] [**svc-sap-type** {**null-star** | **dot1q** | **dot1q-preserve**}] [**customer-vid** *vlan-id*] (for 7210 SAS-E)

- **vpls** *service-id* [customer *customer-id*] [**create**] [**m-vpls**] [**svc-sap-type {null-star|dot1q-preserve|dot1q-range|any}**] [**customer-vid** *vlan-id*] (for 7210 SAS-D)
- **vpls** *service-id* **[customer** *customer-id***] [create] [m-vpls] [svc-sap-type {any|dot1q-range}] [r-vpls] (for 7210 SAS-K)**
- **no vpls** *service-id*
  - **sap** *sap-id* [**split-horizon-group** *group-name*] [**eth-ring** *ring-index*]
  - **no sap** *sap-id*
    - **accounting-policy** *acct-policy-id*
    - **no accounting-policy**
    - **bpdu-translation** {**auto** | **pvst** | **stp**}
    - **no bpdu-translation**
    - [**no**] **collect-stats**
    - **description** *description-string*
    - **no description**
    - [**no**] **disable-aging**
    - [**no**] **disable-learning**
    - [**no**] **discard-unknown-source**
    - **eth-cfm**
      - **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**
          - **client-meg-level** [*level* [*level...*]]
          - **no client-meg-level**
          - [**no**] **description**
          - **interval** {**1**| **60**}
          - **no interval**
          - **priority** *priority-value*
          - **no priority**
          - **no send-ais-on-port-down**
          - **send-ais-on-port-down**
        - [**no**] **ccm-enable**
        - **ccm-ltm-priority** *priority*
        - **no ccm-ltm-priority**
        - **description** *description-string*
        - **no description**
        - [**no**] **eth-test-enable**
          - **bit-error-threshold** *bit-errors*
          - **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
          - **no test-pattern**
        - **mac-address** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
        - **mac-address** *mac-address*
        - **no mac-address**
        - **one-way-delay-threshold** *seconds*
        - [**no**] **shutdown**
    - **l2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp]**
    - **no l2pt-termination**
    - **limit-mac-move** [**blockable** | **non-blockable**]
    - **no limit-mac-move**
    - [**no**] **mac-pinning**
    - **mac-swap-enable**
    - **no mac-swap-enable**
    - **max-nbr-mac-addr** *table-size*

```
— no max-nbr-mac-addr
— tod-suite tod-suite-name
— no tod-suite
```

VPLS SAP Configuration - QoS and Filter commands for 7210 SAS-D and 7210 SAS-E

```
config
    — service
        — vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-star|dot1q-preserve|dot1q-range|any}] [customer-vid vlan-id] - SAS-D
        — vpls service-id [customer customer-id] [create] [m-vpls] [customer-vid vlan-id] [svc-sap-type {null-star|dot1q-preserve|dot1q}] - SAS-E
        — no vpls service-id
            — sap sap-id [create] [eth-ring ring-index]
            — no sap sap-id
                — egress
                    — aggregate-meter-rate rate-in-kbps [burst burst-in-kbits] [enable-stats] (supported only on 7210 SAS-D)
                    — no aggregate-meter-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]
                — ingress
                    — aggregate-meter-rate rate-in-kbps [burst burst-in-kbits] (supported only on 7210 SAS-D)
                    — no aggregate-meter-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]
                    — qos policy-id
                    — no qos
```

VPLS SAP Configuration - QoS and Filter commands for 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

```
config
    — service
        — vpls
        — no vpls
            — sap sap-id [create] [split-horizon-group group-name]
            — no sap sap-id
                — egress
                    — agg-shaper-rate agg-rate
                    — no agg-shaper-rate
                    — dot1p-inner dot1p-inner
                    — no dot1p-inner
```

```
            — no dot1p-outer
            — dot1p-outer dot1p-outer
            — 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 policy-id
            — no qos
      — ingress
            — agg-shaper-rate agg-rate
            — no agg-shaper-rate
            — aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
            — no aggregate-meter-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]
            — qos policy-id
            — no qos
```

VPLS Service SAP IGMP Snooping and MVR commands

```
config
    — service
        — [no] vpls
            — igmp-snooping
                — mvr
                    — description description-string
                    — no description
                    — group-policy policy-name
                    — no group-policy
                    — no shutdown
                    — shutdown
                — query-interval seconds
                — no query-interval
                — no query-src-ip
                — query-src-ip ip-address
                — no report-src-ip
                — report-src-ip ip-address
                — robust-count count
                — no robust-count
                — no shutdown
                — shutdown
            — sap sap-id [split-horizon-group group-name] [eth-ring ring-index]
            — no sap sap-id
                — igmp-snooping
                    — disable-router-alert-check
                    — 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
                    — [no] mrouter-port
                    — mvr
                        — from-vpls service-id
                        — no from-vpls
                        — to-sap sap-id
                        — no to-sap
                    — query-interval seconds
                    — no query-interval
                    — query-response-interval seconds
                    — no query-response-interval
                    — robust-count count
                    — no robust-count
                    — [no] send-queries
                    — static
                        — [no] group group-address
                            — [no] source ip-address (applicable only in access-uplink mode)
                            — [no] starg
```

— **version** *version*
— **no version**
— **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**

VPLS SAP Meter Override Commands

```
config
    — service
        — vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type {null-star | dot1q | dot1q-preserve}] [customer-vid vlan-id]_(for 7210 SAS-E)
        —
        — vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-star|dot1q-preserve|dot1q-range|any}] [customer-vid vlan-id] (for 7210 SAS-D)
        — no vpls service-id
            — sap sap-id [split-horizon-group group-name] [eth-ring ring-index]
            — no sap sap-id
                — ingress
                    — meter-override
                        — meter meter-id [create]
                        — no meter meter-id
                            — adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
                            — cbs size-in-kbytes (supported on 7210 SAS-E)
                            — no cbs
                            — cbs size [kbits|bytes|kbytes] (supported on 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C)
                            — no cbs
                            — mbs size-in-kbits (supported on 7210 SAS-E)
                            — no mbs
                            — mbs size [kbits|bytes|kbytes] (supported on 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C)
                            — no mbs
                            — no mode
                            — no mode
                            — rate cir cir-rate [pir pir-rate]
```

VPLS service SAP xSTP commands

```
config
    — service
        — vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type {null-star | dot1q | dot1q-preserve}] [customer-vid vlan-id]_ (for 7210 SAS-E)
        —
        — vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-star|dot1q-preserve|dot1q-range|any}] [customer-vid vlan-id] (for 7210 SAS-D)
        — vpls service-id [customer customer-id] [create] [svc-sap-type {any|dot1q-range}] [r-vpls]
        — no vpls service-id
            — sap sap-id [split-horizon-group group-name] [eth-ring ring-index]
            — no sap sap-id
                — 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-port-priority stp-priority
                  — no mst-port-priority
            — path-cost sap-path-cost
            — no path-cost
            — [no] port-num virtual-port-number
            — priority stp-priority
            — no priority
            — no root-guard
            — root-guard
            — [no] shutdown
      — managed-vlan-list
            — [no] default-sap
            — [no] range vlan-range
```

VPLS SAP Statistics commands for 7210 SAS-E and 7210 SAS-D

```
config
    — service
        — vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type {null-star | dot1q | dot1q-preserve}] [customer-vid vlan-id] (for 7210 SAS-E)
        — vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-star|dot1q-preserve|dot1q-range|any}] [customer-vid vlan-id] (for 7210 SAS-D)
        — no vpls service-id
            — sap sap-id [split-horizon-group group-name] [eth-ring ring-index] [create]
            — no sap sap-id
                — statistics
                    — egress
                        — [no] packets-forwarded-count (supported only on 7210 SAS-E)
                        — forwarded-count (supported only on 7210 SAS-D)
                    — ingress
                        — counter-mode {in-out-profile-count|forward-drop-count} {packet | octet}(supported only on 7210 SAS-E)
                        — counter-mode {in-out-profile-count|forward-drop-count}(supported only on 7210 SAS-D)
                        — [no] received-count (supported only on 7210 SAS-D)
```

VPLS Mesh SDP Commands for 7210 SAS-K 2F4T6C

```
config
    — service
        — [no] vpls
            — mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
            — no mesh-sdp sdp-id[:vc-id]
                — accounting-policy acct-policy-id
                — no accounting-policys
                — [no] collect-stats
                — [no] control-word
                — description description-string
                — no description
                — egress
                    — no vc-label [egress-vc-label]
                — eth-cfm
                    — 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
                            — 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
                        — [no] description description-string
```

```
                    — [no] eth-test-enable
                        — bit-error-threshold bit-errors
                        — test-pattern {all-zeros | all-ones} [crc-enable]
                        — no test-pattern
                    — mac-address {allDef | macRemErrXcon | remErrX-
                      con | errXcon | xcon | noXcon}
                    — mac-address mac-address
                    — no mac-address
                    — one-way-delay-threshold seconds
                    — [no] shutdown
— [no] force-vlan-vc-forwarding
— igmp-snooping
        — [no] disable-router-alert-check
        — 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
        — [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 grp-ip-address
                — [no] starg
        — version version
        — no version
— ingress
        — vc-label egress-vc-label
— [no] mac-pinning
— [no] static-mac ieee-address
— [no] shutdown
— vlan-vc-tag 0..4094
— no vlan-vc-tag [0..4094]
— dhcp
        — description description-string
        — no description
        — no snoop
        — snoop
```

VPLS Spoke SDP Commands for 7210 SAS-K 2F4T6C

```
config
    — service
        — vpls
        — no vpls
            — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] [split-horizon-group group-name] endpoint endpoint-name
            — no spoke-sdp sdp-id[:vc-id]
                — accounting-policy acct-policy-id
                — no accounting-policy
                — [no] block-on-mesh-failure
                — bpdu-translation {auto | pvst | stp}
                — no bpdu-translation
                — [no] collect-stats
                — [no] control-word
                — description description-string
                — no description
                — [no] disable-aging
                — [no] disable-learning
                — [no] discard-unknown-source
                — dhcp
                    — description description-string
                    — no description
                    — no snoop
                    — snoop
                — eth-cfm
                    — 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
                        — 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
                        — [no] description description string[
                        — no] eth-test-enable
                        — bit-error-threshold bit-errors
                        — test-pattern {all-zeros | all-ones} [crc-enable]
                        — no test-patternmac-address {allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon}
                        — mac-address mac-address
                        — no mac-addressone-way-delay-threshold seconds
                        — [no] shutdown
                    — mip [mac mac address]
                    — mip default-mac
                    — no mip
                — egress
                    — vc-label egress-vc-label
```

```
            — no vc-label [egress-vc-label]
        — [no] force-vlan-vc-forwarding
        — igmp-snooping
            — [no] disable-router-alert-check
            — 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
            — [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] starg
            — version version
            — no version
        — [no] ignore-standby-signaling
        — ingress
            — vc-label egress-vc-label
            — no vc-label [egress-vc-label]
        — [no] l2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp]
        — limit-mac-move [blockable | non-blockable]
        — no limit-mac-move
        — [no] mac-pinning
        — max-nbr-mac-addr table-size
        — no max-nbr-mac-addr
        — precedence precedence-value | primary
        — no precedence
        — [no] shutdown
        — [no] static-mac ieee-address
        — statistics
            — ingress
            — [no] drop-count-extra-vlan-tag-pkts
        — 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 root-guard
            — root-guard
            — [no] shutdown
        — vlan-vc-tag 0..4094
```

— **no vlan-vc-tag** [*0..4094*]

Routed VPLS Commands applicable only to 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

config
— **service**
— **vpls** *service-id* **[customer** *customer-id***] [vpn** *vpn-id***] [r-vpls] [create]**
— **[no] allow-ip-int-binding**

Show Commands

```
show
    — service
        — fdb-info
        — fdb-mac ieee-address [expiry]
        — id service-id
            — all
            — arp [ip-address]|[mac ieee-address]|[sap sap-id]|[interface ip-int-name] | [summary]
            — base
            — bgp-ad
            — dhcp
                — statistics [sap sap-id] [interface interface-name]
                — summary [interface interface-name | saps]
            — fdb [sap sap-id]| [mac ieee-address | endpoint endpoint | detail] [expiry]
            — igmp-snooping
                — all
                — base
                — mrouters
                — mvr
                — mrouters [detail]
                — port-db sap sap-id [detail]
                — port-db sap sap-id group grp-address
                — proxy-db [detail]
                — proxy-db [group grp-ip-address]
                — querier
                — static [sap sap-id]
                — statistics[sap sap-id | sdp sdp-ic:vc-id]
            — l2pt disabled
            — l2pt [detail]
            — mac-move
            — mfib [brief]
            — mfib [group grp-address | mstp-configuration
            — sap [sap-id [detail | stp]]
            — stp [detail]
        — sap-using [sap sap-id]
        — sap-using [ingress | egress] filter filter-id
        — sap-using [ingress] qos-policy qos-policy-id
        — service-using [vpls]
```

Clear Commands

```
clear
    — service
        — id service-id
            — fdb {all | mac ieee-address | sap sap-id |}
            — igmp-snooping
                — port-db sap sap-id [group grp-address]
                — querier
            — statistics [all | sap sap-id | sdp sdp-id:vc-id]
                —
            — stp
                — detected-protocols [all | sap sap-id]
        — statistics
            — id service-id
                — counters
                — stp
            — sap sap-id {all | counters | stp }
```

Debug Commands

```
debug
    — service
        — id service-id
            — [no] event-type {config-change | svc-oper-status-change | sap-oper-status-change | sdpbind-oper-status-change}
            — [no] sap sap-id
```

VPLS Service Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls
config>service>vpls>igmp-snooping
config>service>vpls>sap
config>service>vpls>sap>stp
config>service>vpls>mesh-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>stp

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

description

Syntax **description** *description-string*
no description

Context config>service>vpls
config>service>vpls>split-horizon-group
config>service>vpls>igmp-snooping>mvr
config>service>vpls>sap
config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>pw-template>split-horizon-group

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default No description associated with the configuration context.

Parameters — *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

dhcp

Syntax **dhcp**

Context config>service>vpls>sap
config>service>ies>sap
config>service>vprn>sap - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables the context to configure DHCP parameters.

option

Syntax [**no**] **option**

Context config>service>vpls>sap>dhcp
config>service>ies>sap>dhcp
config>service>vprn>sap>dhcp - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.

The **no** form of this command returns the system to the default.

Default no option

action

Syntax **action** [*dhcp-action*] {***replace*** | ***drop*** | ***keep***}
no action

Context config>service>vpls>sap>dhcp>option
config>service>ies>sap>dhcp>option
config>service>vprn>sap>dhcp>option - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the Relay Agent Information Option (Option 82) processing.

The **no** form of this command returns the system to the default value.

Default The default is to keep the existing information intact.

Parameters *dhcp-action* — Specifies the DHCP option action.

replace — In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.

keep — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on towards the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

gi-address

Syntax **gi-address** *ip-address* [*src-ip-addr*]
no gi-address

Context config>service>ies>if>dhcp
config>service>vprn>if>dhcp - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the gateway interface address for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. The GI address is needed, when the router functions as a DHCP relay, to distinguish between the different subscriber interfaces and potentially between the group interfaces defined.

By default, the GI address used in the relayed DHCP packet is the primary IP address of a normal IES interface. Specifying the GI address allows the user to choose a secondary address.

Default no gi-address

Parameters *ip-address* — Specifies the host IP address to be used for DHCP relay packets.

src-ip-address — Specifies the source IP address to be used for DHCP relay packets.

circuit-id

Syntax **circuit-id** [**ascii-tuple** | **ifindex** | **sap-id** | **vlan-ascii-tuple**]
no circuit-id

Context config>service>vpls>sap>dhcp>option

config>service>ies>if>dhcp>option
config>service>vprn>if>dhcp>option - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by "|".

In order to send a tuple in the circuit ID, the **action replace** command must be configured in the same context.

If disabled, the **circuit-id** sub-option of the DHCP packet will be left empty.

The **no** form of this command returns the system to the default.

Default circuit-id

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 "|".

ifindex — Specifies that the interface index will be used. (The If Index of a router interface can be displayed using the command `show>router>interface>detail`)

sap-id — Specifies that the SAP identifier will be 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-encapsulated ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

remote-id

Syntax [**no**] **remote-id** [**mac** | **string** *string*]

Context config>service>vpls>sap>dhcp>option
config>service>ies>sap>dhcp>option
config>service>vprn>sap>dhcp>option - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies what information goes into the remote-id sub-option in the DHCP Relay packet.

If disabled, the **remote-id** sub-option of the DHCP packet will be left empty.

The **no** form of this command returns the system to the default.

Default no remote-id

Parameters **mac —** This keyword specifies the MAC address of the remote end is encoded in the sub-option.

string *string* **—** Specifies the remote-id.

vendor-specific-option

Syntax [**no**] **vendor-specific-option**

Context config>service>vpls>sap>dhcp>option
config>service>ies>if>dhcp>option
config>service>vprn>if>dhcp>option - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the vendor specific sub-option of the DHCP relay packet.

client-mac-address

Syntax [**no**] **client-mac-address**

Context config>service>vpls>sap>dhcp>option>vendor
config>service>vpls>sap>dhcp>option>vendor
config>service>vpls>sap>dhcp>option>vendor - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables the sending of the MAC address in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the vendor specific sub-option of the DHCP relay packet.

pool-name

Syntax [**no**] **pool-name**

Context config>service>vprn>if>dhcp>option>vendor - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>ies>if>dhcp>option>vendor

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command sends the pool name in the Nokia vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending.

DHCP COMMANDS

sap-id

Syntax [**no**] **sap-id**

Context config>service>vpls>sap>dhcp>option>vendor
config>service>ies>sap>dhcp>option>vendor
config>service>vprn>sap>dhcp>option>vendor - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables the sending of the SAP ID in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the vendor specific sub-option of the DHCP relay packet.

service-id

Syntax [**no**] **service-id**

Context config>service>vpls>sap>dhcp>option>vendor
config>service>ies>sap>dhcp>option>vendor
config>service>vprn>sap>dhcp>option>vendor - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables the sending of the service ID in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the vendor specific sub-option of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

Context config>service>vpls>sap>dhcp>option>vendor
config>service>ies>sap>dhcp>option>vendor
config>service>vprn>sap>dhcp>option>vendor - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the string in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command returns the default value.

Parameters *text* — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

system-id

Syntax [**no**] **system-id**

Context config>service>vpls>sap>dhcp>option>vendor
config>service>ies>sap>dhcp>option>vendor
config>service>vprn>sap>dhcp>option>vendor - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies whether the system-id is encoded in the vendor specific sub-option of Option 82.

snoop

Syntax [**no**] **snoop**

Context config>service>vpls>sap>dhcp
config>service>vpls>mesh-sdp>dhcp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>spoke-sdp>dhcp- Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables DHCP snooping of DHCP messages on the SAP. Enabling DHCP snooping on VPLS interfaces (SAPs) is required where DHCP messages where Option 82 information is to be inserted. This includes interfaces in the path to receive messages from either DHCP servers or from subscribers.

Use the **no** form of the command to disable DHCP snooping on the specified VPLS SAP.

Default no snoop

VPLS Service Commands

vpls

Syntax

vpls *service-id* [**customer** *customer-id*] [**create**] [**m-vpls**] **[svc-sap-type {null-star|dot1q-preserve|dot1q-range|any}] [customer-vid** *vlan-id***] - 7210 SAS-D**
vpls *service-id* [**customer** *customer-id*] [**create**] [**m-vpls**] **[customer-vid** *vlan-id***] [svc-sap-type {null-star|dot1q-preserve|dot1q}] - 7210 SAS-E**
vpls *service-id* [**customer** *customer-id*] **[create] [m-vpls] [svc-sap-type {any|dot1q-range}] [r-vpls] - 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C**
no vpls *service-id*

Context config>service

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command creates or edits a Virtual Private LAN Services (VPLS) instance. The **vpls** command is used to create or maintain a VPLS service. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites together acting like a zero-hop, Layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context. When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

More than one VPLS service may be created for a single customer ID.

By default, no VPLS instances exist until they are explicitly created.

- 7210 SAS-E, 7210 SAS-D, and 7210 SAS-K 2F2T1C supports local service.
- 7210 SAS-K 2F4T6C supports both local and distributed service.

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs defined within the service ID have been shutdown and deleted, and the service has been shutdown.

Parameters *service-id* — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values *service-id*: 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

m-vpls — Specifies a management VPLS.

create — This keyword is mandatory while creating a VPLS service. Keyword used to create the service instance. The create keyword requirement can be enabled or disabled in the environment>create context.

customer-vid vlan-id — Defines the dot1q VLAN ID to be specified while creating the local Dot1q SAP for svc-sap-type dot1q-preserve.

Values 1 — 4094

svc-sap-type- — This parameter specifies the type of access SAPs and access-uplink SAPs allowed in the service.

dot1q-range -Specifies that the access SAP in the service can use VLAN ranges as the SAP tags. The VLAN ranges are configured using the CLI command *configure> connection-profile*. On ingress of the access dot1q SAP using VLAN ranges, the received tag on the SAP is preserved. A VPLS service with svc-sap-type set to dot1q-range can be used for Epipe emulation with G8032 for protection. For more information about the capabilities and restrictions, see Epipe Emulation using Dot1q VLAN range SAP in VPLS with G.8032 on page 338. This is supported on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

dot1q - Specifies that the allowed SAP in the service are Dot1q SAPs and dot1q explicit null SAPs. This is supported only on 7210 SAS-E.

dot1q-preserve - Specifies that the allowed SAP in the service are Dot1q. The Dot1q ID is not stripped after packets matches the SAP. This command is supported only on 7210 SAS-E and 7210 SAS-D. It is not supported on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

null-star - Specifies that the allowed SAP in the service, which can be null SAPs, Dot1q default, Q.* SAP, 0.* SAP or Default QinQ SAP. This command is supported only on 7210 SAS-E and 7210 SAS-D. It is not supported on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

any - Specifies that the service can be configured with null, dot1q and QinQ SAPs. It is supported only on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

“any”

r-vpls — Allows this VPLS instance to be associated with an IP interface to provide R-VPLS functionality. It is supported only on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

Default “any” - 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C
“null-star” - 7210 SAS-E

bgp

Syntax **bgp**

Context config>service>vpls

Platforms Supported: 7210 SAS-K 2F4T6C.

This command enables the context to configure the BGP related parameters to BGP AD.

block-on-mesh-failure

Syntax [**no**] **block-on-mesh-failure**

Context config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>endpoint

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command enables blocking (brings the entity to an opererationally down state) after all configured SDPs or endpoints are in opererationally 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 - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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.

l2pt-termination

Syntax **l2pt-termination** [cdp] [dtp] [pagp] [stp] [udld] [vtp]
no l2pt-termination

Context config>service>vpls>sap
config>service>vpls>spoke-sdp Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables Layer 2 Protocol Tunneling (L2PT) termination on a given SAP. L2PT termination is supported for STP/CDP/DTP/PAGP/UDLD and VTP PDUs.

This feature can be enabled only if STP is disabled in the context of the given VPLS service.

NOTE: CDP, DTP, PAGP, STP, UDID and VTP is supported only on 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C. It is not supported on 7210 SAS-E

Default no l2pt-termination

Parameters *cdp* — Specifies the Cisco discovery protocol.

dtp — Specifies the dynamic trunking protocol.

pagp — Specifies the port aggregation protocol.

stp — Specifies all spanning tree protocols: stp, rstp, mstp, pvst (default).

udld — Specifies unidirectional link detection.

vtp — Specifies the VLAN trunking protocol.

disable-aging

Syntax [**no**] **disable-aging**

Context config>service>vpls
config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command disables MAC address aging across a VPLS service or on a VPLS service SAP.

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 a VPLS service instance, the local age timer is applicable to both the local learned and remote learned MAC entries in the VPLS forwarding database (FDB). The disable-aging command at the service level 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 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 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 - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command disables learning of new MAC addresses in the VPLS forwarding database (FDB) for the service instance.

When **disable-learning** is enabled, new source MAC addresses will not be entered in the VPLS service forwarding database.

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

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

By default, packets with unknown destination MAC addresses are flooded. If discard-unknown is enabled at the VPLS level, packets with unknown destination MAC address will be dropped instead (even when configured FIB size limits for VPLS or SAP are not yet reached).

The **no** form of this command allows flooding of packets with unknown destination MAC addresses in the VPLS.

Default **no discard-unknown** — Packets with unknown destination MAC addresses are flooded.

endpoint

Syntax **endpoint** *endpoint-name* [**create**]
no endpoint

Context config>service>vpls

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures a service endpoint.

Parameters *endpoint-name* — Specifies an endpoint name up to 32 characters in length.

create — This keyword is mandatory while creating a service endpoint.

description

Syntax **description** *description-string*
no description

Context config>service>vpls>endpoint

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

ignore-standby-signaling

Syntax [**no**] **ignore-standby-signaling**

Context config>service>vpls>endpoint
config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

When this command is enabled, the node will ignore standby-bit received from TLDP peers for the given spoke SDP and performs internal tasks withot 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 disabled

split-horizon-group

Syntax **split-horizon-group** *group-name* [**create**]

Context config>service>vpls
config>service>pw-template - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command is used to create a new split horizon group for the VPLS instance. Traffic arriving on a SAP or spoke SDP within this split horizon group will not be copied to other SAPs or spoke SDPs in the same split horizon group.

A split horizon group must be created before SAPs and spoke SDPs can be assigned to the group. The split horizon group is defined within the context of a single VPLS instance. The same group-name can be re-used in different VPLS instances.

he **no** form of the command removes the group name from the configuration.

Parameters *group-name* — Specifies the name of the split horizon group to which the SAP or Spoke-SDP belongs.

create — Mandatory keyword to create a split-horizon group.

fdb-table-high-wmark

Syntax [**no**] **fdb-table-high-wmark** *high-water-mark*

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the value to send logs and traps when the threshold is reached.

Parameters *high-water-mark* — 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** *low-water-mark*

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the value to send logs and traps when the threshold is reached.

Parameters *low-water-mark* — 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** *table-size*
no fdb-table-size [*table-size*]

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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 maxium FDB table size to default.

Default **250** — Forwarding table of 250 MAC entries.

Parameters *table-size* — Specifies the maximum number of MAC entries in the FDB.

Values 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C- [1..30719]
7210 SAS-E - [1..16383]

bgp-ad

Syntax **[no] bgp-ad**

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures BGP auto-discovery.

vpls-id

Syntax **vpls-id** *vpls-id*

Context config>service>vpls>bgp-ad

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures the VPLS ID component that will be signaled in one of the extended community attributes (ext-comm). Values and format (6 bytes, other 2 bytes of type-subtype will be automatically generated).

Parameters *vpls-id* — Specifies a globally unique VPLS ID for BGP auto-discovery in this VPLS service.

Values vpls-id : <ip-addr:comm-val>|<as-number:ext-comm-val>
ip-addr a.b.c.d
comm-val 0 — 65535

vsi-export

Syntax **vsi-export policy-name** *[policy-name...(up to 5 max)]*
no vsi-export

Context config>service>vpls>bgp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command specifies the name of the VSI export policies to be used for BGP auto-discovery, if this feature is configured in the 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 policy-name [policy-name...(up to 5 max)]**
no vsi-import

Context config>service>vpls>bgp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command specifies the name of the VSI import policies to be used for BGP auto-discovery, if this feature is configured in the 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-target

Syntax **route-target {ext-community|{[export ext-community][import ext-community]}}**
no route-target

Context config>service>vpls>bgp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures the route target (RT) component that will be signaled in the related MPBGP attribute to be used for BGP auto-discovery, if this feature is configured in the VPLS service.

If this command is not used, the RT is built automatically using the VPLS ID. The ext-comm can have the same two formats as the VPLS ID, a two-octet AS-specific extended community, IPv4 specific extended community.

The following rules apply:

- If BGP AD VPLS-id is configured & no RT is configured under BGP node - RT = VPLS-ID.
- If BGP AD VPLS-id is not configured then an RT value must be configured under BGP node. (this is the case when only BGP VPLS is configured)
- If BGP AD VPLS-id is configured and an RT value is also configured under BGP node, the configured RT value prevails

Parameters ***export*** *ext-community* — •Specify communities allowed to be sent to remote PE neighbors.

import *ext-community* — •Specify communities allowed to be accepted from remote PE neighbors.

pw-template-binding

Syntax **pw-template-binding policy-id** *[split-horizon-group group-name] [import-rt {extcommunity,...(up to 5 max)}]*
no pw-template-bind *policy-id*

Context config>service>vpls>bgp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

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 BGP-AD,f this feature is configured in the VPLS service.

The tools perform commands can be used to control the application of changes in pw-template for BGP-AD.

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 — 42949672954byte-asnumber 0 — 4294967295

local-age

Syntax **local-age** *aging-timer*
no local-age

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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). MACs associated with a SAP are classified as local MACs, and MACs associated with are remote MACsQinQ / access uplink SAPs.

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). The **local-age** timer specifies the aging time for local learned MAC addresses.

The **no** form of this command returns the local aging timer to the default value.

Default **local age 300** — Local MACs aged after 300 seconds.

Parameters *aging-timer* — The aging time for local MACs expressed in seconds.

Values 60 — 86400

mac-move

Syntax [**no**] **mac-move**

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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.

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** context. This means that when the re-learn rate has exceeded the limit, another (blockable) SAP will be disabled instead.

The **mac-move** command enables the feature at the service level for SAPs, as only those objects can be blocked by this feature.

The operation of this feature is the same on the SAP. For example, if a MAC address moves from SAP to SAP, one will be blocked to prevent thrashing.

mac-move will disable a VPLS port when the number of relearns detected has reached the number of relearns needed to reach the move-frequency in the 5-second interval. For example, when the move-frequency is configured to 1 (relearn per second) mac-move will disable one of the VPLS ports when

5 relearns were detected during the 5-second interval because then the average move-frequency of 1 relearn per second has been reached. This can already occur in the first second if the real relearn rate is 5 relearns per second or higher.

The **no** form of this command disables MAC move.

move-frequency

Syntax **move-frequency** *frequency*
no move-frequency

Context config>service>vpls>mac-move

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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

retry-timeout

Syntax **retry-timeout** *timeout*
no retry-timeout

Context config>service>vpls>mac-move

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.

It is recommended that the retry-timeout value is larger or equal to 5s * cumulative factor of the highest priority port so that the sequential order of port blocking will not be disturbed by re-initializing lower priority ports.

A zero value indicates that the SAP will not be automatically re-enabled after being disabled. If, after the SAP is reenabled it is disabled again, the effective retry timeout is doubled in order to avoid thrashing.

The **no** form of the command reverts to the default value.

Default 10 (when mac-move is enabled)

Parameters *timeout* — Specifies the time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.

Values 0 — 120

mfib-table-high-wmark

Syntax [**no**] **mfib-table-high-wmark** *high-water-mark*

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the multicast FIB high watermark. When the percentage filling level of the multicast FIB exceeds the configured value, a trap is generated and/or a log entry is added.

Parameters *high-water-mark* — Specifies the multicast FIB high watermark as a percentage.

Values 1 — 100

Default 95%

mfib-table-low-wmark

Syntax [**no**] **mfib-table-low-wmark** *low-water-mark*

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the multicast FIB low watermark. When the percentage filling level of the Multicast FIB drops below the configured value, the corresponding trap is cleared and/or a log entry is added.

Parameters *low-water-mark* — Specifies the multicast FIB low watermark as a percentage.

Values 1 — 100

Default 90%

mfib-table-size

Syntax **mfib-table-size** *size*
no mfib-table-size

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the maximum number of (s,g) entries in the multicast forwarding database (MFIB) for this VPLS instance.

The *mfib-table-size* parameter specifies the maximum number of multicast database entries for both learned and static multicast addresses for the VPLS instance. When a table-size limit is set on the mfib of a service which is lower than the current number of dynamic entries present in the mfib then the number of entries remains above the limit.

The **no** form of this command removes the configured maxium MFIB table size.

Default none

Parameters *size* — The maximum number of (s,g) entries allowed in the Multicast FIB.

Values 7210 SAS-E - [1..2047]
7210 SAS- D [1..1024]
7210 SAS-K 2F2T1C - [1..1000]
7210 SAS-K 2F4T6C - [1..1015]

remote-age

Syntax **remote-age** *seconds*
no remote-age

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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). MACs associated with a SAP are classified as local MACs.

Like in a layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The **remote-age** timer specifies the aging time for remote learned MAC addresses. To reduce the amount of signaling required between switches configure this timer larger than the **local-age** timer.

The **no** form of this command returns the remote aging timer to the default value.

Default **remote age 900** — Remote MACs aged after 900 seconds

Parameters *seconds* — The aging time for remote MACs expressed in seconds.

Values 60 — 86400

send-flush-on-failure

Syntax [**no**] **send-flush-on-failure**

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command enables sending out “flush-all-from-ME” messages to all LDP peers included in affected VPLS, in the event of physical port failures or “oper-down” events of individual SAPs. This feature provides an LDP-based mechanism for recovering a physical link failure in a dual-homed connection to a VPLS service. This method provides an alternative to RSTP solutions where dual homing redundancy and recovery, in the case of link failure, is resolved by RSTP running between a PE router and CE devices. If the endpoint is configured within the VPLS and send-flush-on-failure is enabled, flush-all-from-me messages will be sent out only when all spoke SDPs associated with the endpoint go down.

This feature cannot be enabled on management VPLS.

Default no send-flush-on-failure

service-mtu

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

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

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 (i.e., 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.

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.

The **no** form of this command returns the default **service-mtu** for the indicated service type to the default value.

Note: To disable service MTU check execute the command no service-mtu-check. Disabling service MTU check allows the packets to pass to the egress if the packet length is lesser than or equal to the MTU configured on the port.

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

split-horizon-group

Syntax [**no**] **split-horizon-group** [*group-name*] [**create**]

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command creates a new split horizon group for the VPLS instance. Traffic arriving on a SAP or spoke SDP within this split horizon group will not be copied to other SAPs or spoke SDPs in the same split horizon group.

A split horizon group must be created before SAPs and spoke SDPs can be assigned to the group.

The split horizon group is defined within the context of a single VPLS. The same group name can be re-used in different VPLS instances.

The **no** form of the command removes the group name from the configuration.

group-name — Specifies the name of the split horizon group to which the SAP or spoke-SDP belongs.

create — Mandatory keyword to create a split-horizon group.

root-guard

Syntax [**no**] **root-guard**

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp - Platforms Supported - 7210 SAS-K 2F4T6C

Description This command specifies whether this port is allowed to become an STP root port. It corresponds to the restrictedRole parameter in 802.1Q. If set, it can cause lack of spanning tree connectivity.

Default no root-guard

tod-suite

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

Context config>service>vpls>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command applies a time-based policy (filter or QoS policy) to the service SAP. The suite name must already exist in the **config>cron** context.

Default no tod-suite

Parameters *tod-suite-name* — Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP. The suite can be applied to more than one SAP.

vsi-id

Syntax **vsi-id**

Context config>service>vpls>bgp-ad

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command enables the context to configure the Virtual Switch Instance Identifier (VSI-ID).

prefix

Syntax **prefix low-order-vsi-id**
no prefix

Context config>service>vpls>bgp-ad>vsi-id

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command specifies the low-order 4 bytes used to compose the Virtual Switch Instance Identifier (VSI-ID) to use for NLRI in BGP auto-discovery in this VPLS service.

If no value is set, the system IP address will be used.

Default no prefix

Parameters *low-order-vsi-id* — Specifies a unique VSI ID.

Values 0— 4294967295

service-name

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

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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 o identify and reference a given service once it is initially created.

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

allow-ip-int-binding

Syntax [no] **allow-ip-int-binding**

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

The allow-ip-int-binding command that sets a flag on the VPLS service that enables the ability to attach an IP interface to the VPLS service in order to make the VPLS service routable. When the allow-ip-int-binding command is not enabled, the VPLS service cannot be attached to an IP interface.

On 7210 SAS-D and 7210 SAS-K 2F2T1C, the VPLS can be bound to only IES IP interface. On 7210 SAS-K 2F4T6C the VPLS can be bound to IES or VPRN IP interface. Only IPv4 addressing and forwarding is supported in the RVPLS context.

When attempting to set the allow-ip-int-binding VPLS flag, the system first checks to see if the correct configuration constraints exist for the VPLS service and the network ports. In Release 5.0 the following VPLS features must be disabled or not configured for the allow-ip-int-binding flag to set:

- SAP ingress QoS policies applied to the VPLS SAPs cannot have MAC match criteria defined
- The VPLS service type cannot be M-VPLS.
- MVR from Routed VPLS and to another SAP is not supported

Once the VPLS allow-ip-int-binding flag is set on a VPLS service, the above features cannot be enabled on the VPLS service.

In the event that a service name is applied to a VPLS service and that service name is also bound to an IP interface but the allow-ip-int-binding flag has not been set on the VPLS service context, the system attempt to resolve the service name between the VPLS service and the IP interface will fail. After the allow-ip-int-binding flag is successfully set on the VPLS service, either the service name on the VPLS service must be removed and reapplied or the IP interface must be re-initialized using the shutdown or no shutdown commands. This will cause the system to reattempt the name resolution process between the IP interface and the VPLS service.

The no form of the command resets the allow-ip-int-binding flag on the VPLS service. If the VPLS service currently has an IP interface attached, the no allow-ip- int-binding command will fail. Once the allow-ip-int-binding flag is reset on the VPLS service, the configuration restrictions associated with setting the flag are removed.

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>service>vpls

Description

Platforms Supported: 7210 SAS-E.

This command creates a logical IP routing interface for a VPLS service. Once created, attributes such as IP address and service access points (SAP) can be associated with the IP interface.

The interface command, under the context of services, is used to create and maintain IP routing interfaces within the VPLS service IDs. The IP interface created is associated with the VPLS management routing instance.This instance does not support routing.

Interface names are case-sensitive and must be unique within the group of defined IP interfaces defined for the network core router instance. Interface names in the dotted decimal notation of an IP address are not allowed. For example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. Duplicate interface names can exist in different router instances.

Enter a new name to create a logical router interface. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, no default IP interface names are defined within the system. All VPLS IP interfaces must be explicitly defined in an enabled state.

The no form of this command removes the IP interface and the entire associated configuration. The interface must be administratively shutdown before issuing the no interface command.

For VPLS services, the IP interface must be shutdown before the SAP on that interface is removed.

For VPLS service, ping and traceroute are the only applications supported.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP.

An interface name:

- Should not be in the form of an IP address.
- Can be from 1 to 32 alphanumeric characters.
- If the string contains special characters (such as #,$,spaces), the entire string must be enclosed within double quotes.

If ip-int-name already exists within the service ID, the context changes to maintain that IP interface. If ip-int-name already exists within another service ID, an error occurs and the context does not change to that IP interface. If ip-int-name does not exist, the interface is created and the context is changed to that interface for further command processing.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*}
address *ip-address mask*

Context config>service>vpls>interface

Description

Platforms Supported: 7210 SAS-E.

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 within the 7210 SAS.

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 adminstratively 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.

Values 1 — 16383

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>vpls>interface

Description **Platforms Supported**: 7210 SAS-E.

This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

The default value for **arp-timeout** is 14400 seconds (4 hours).

The **no** form of this command restores **arp-timeout** to the default value.

Default 14400 seconds

Parameters *seconds* — The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 — 65535

mac

Syntax **mac** *ieee-address*
no mac

Context config>service>vpls>interface

Description **Platforms Supported**: 7210 SAS-E.

This command assigns a specific MAC address to a VPLS IP interface.

The **no** form of the command returns the MAC address of the IP interface to the default value.

Default The system chassis MAC address.

Parameters *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

static-arp

Syntax **static-arp** *ip-address ieee-address*
no static-arp *ip-address* [*ieee-address*]

Context config>service>vpls>interface

Description **Platforms Supported**: 7210 SAS-E.

This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of the command removes a static ARP entry.

Default None

Parameters *ip-address* — Specifies the IP address for the static ARP in dotted decimal notation.

ieee-mac-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

VPLS STP Commands

stp

Syntax **stp**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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 - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>pw-template - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures automatic detection of the edge port characteristics of the SAP.

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 - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the SAP 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 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

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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 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 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

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the Spanning Tree Protocol (STP) hello time for the Virtual Private LAN Service (VPLS) STP instance.

The hello time parameter defines the default timer value that controls the sending interval between BPDU configuration messages by this bridge, on ports where this bridge assumes the designated role.

The active hello time for the spanning tree is determined by the root bridge (except when the STP is running in RSTP mode, then the hello time is always taken from the locally configured parameter).

The configured hello-time can also be used to calculate the forward delay. See auto-edge on page 444.

The **no** form of this command returns the hello time to the default value.

Default 2 seconds

Parameters *hello-time* — The hello time for the STP instance in seconds.

hold-count

Syntax **hold-count** *BDPU tx hold count*
no hold-count

Context config>service>vpls>stp

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the peak number of BPDUs that can be transmitted in a period of one second.

The **no** form of this command returns the hold count to the default value

Default 6

Parameters *BDPU tx hold count* — The hold count for the STP instance in seconds.

Values 1 — 10

link-type

Syntax **link-type** {**pt-pt** | **shared**}
no link-type

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command 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 via a shared media, their SAP should all be configured as shared, and timer-based transitions are used.

The **no** form of this command returns the link type to the default value.

Default pt-pt

mst-instance

Syntax **mst-instance** *mst-inst-number*

Context config>service>vpls>sap>stp
config>service>vpls>stp

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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.

Default none

Parameters *mst-inst-number* — Specifies an existing Multiple Spanning Tree Instance number.

Values 1 — 4094

mst-path-cost

Syntax **mst-path-cost** *inst-path-cost*
no mst-path-cost

Context config>service>vpls>sap>stp>mst-instance
config>service>vpls>stp>mst-instance

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This commands specifies path-cost within a given instance. If a loop occurs, this parameter indicates the probability of a given port being assigned a forwarding state. (The highest value expresses lowest priority).

The **no** form of this command sets port-priority to its default value.

Default The path-cost is proportional to link speed.

Parameters *inst-path-cost* — Specifies the contribution of this port to the MSTI path cost.

Values 1 — 200000000

mst-port-priority

Syntax **mst-port-priority** *stp-priority*
no mst-port-priority

Context config>service>vpls>sap>stp>mst-instance
config>service>vpls>stp>mst-instance

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This commands specifies the port priority within a given instance. If a loop occurs, this parameter indicates the probability of a given port being assigned a forwarding state.

The **no** form of this command sets port-priority to its default value.

Default 128

Parameters *stp-priority* — Specifies the value of the port priority field.

max-age

Syntax **max-age** *seconds*
no max-age

Context config>service>vpls>stp

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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.

STP uses the max-age value configured in the root bridge. This value is propagated to the other bridges via the BPDUs.

The **no** form of this command returns the max ageto the default value.

Default 20 seconds

Parameters

seconds — The max info age for the STP instance in seconds. Allowed values are integers in the range 6 to 40. mode

Syntax **mode** {**rstp** | **comp-dot1w** | **dot1w** | **mstp**}
no mode

Context config>service>vpls>stp

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the version of Spanning Tree Protocol the bridge is currently running.

See section Spanning Tree Operating Modes on page 299 for details on these modes.

The **no** form of this command returns the STP variant to the default.

Default rstp

Parameters **rstp —** Corresponds to the Rapid Spanning Tree Protocol specified in IEEE 802.1D/D4-2003.

dot1w — Corresponds to the mode where the Rapid Spanning Tree is backward compatible with IEEE 802.1w.

compdot1w — Corresponds to the Rapid Spanning Tree Protocol fully conformant to IEEE 802.1w.

mstp — Sets MSTP as the STP mode of operation. Corresponds to the Multiple Spanning Tree Protocol specified in 802.1Q REV/D5.0-09/2005

mst-instance

Syntax [**no**] **mst-instance** *mst-inst-number*

Context config>service>vpls>stp

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command creates the context to configure Multiple Spanning Tree Instance (MSTI) related parameters. MSTP supports “16” instances. The instance “0” is mandatory (by protocol) and cannot be created by the CLI. The software automatically maintains this instance.

Default none

Parameters *mst-inst-number* — Specifies the Multiple Spanning Tree instance.

Values 1 — 4094

mst-priority

Syntax **mst-priority** *bridge-priority*
no mst-priority

Context config>service>vpls>stp>mst-instance

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the bridge priority for this specific Multiple Spanning Tree Instance for this service. The *bridge-priority* value reflects likelihood that the switch will be chosen as the regional root switch (65535 represents the least likely). It is used as the highest 4 bits of the Bridge ID included in the MSTP BPDU's generated by this bridge.

The values of the priority are only multiples of 4096 (4k). If a value is specified that is not a multiple of 4K, the value is replaced by the closest multiple of 4K(lower than the value entered).

The **no** form of this command sets the bridge-priority to its default value.

Default 32768 — All instances created by the **vlan-range** command do not have explicit definition of bridge-priority and will inherit the default value.

Parameters *bridge-priority* — Specifies the priority of this specific Multiple Spanning Tree Instance for this service.

Values 0 — 65535

vlan-range

Syntax [**no**] **vlan-range** [*vlan-range*]

Context config>service>vpls>stp>mst-instance

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies a range of VLANs associated with a certain MST-instance. This range applies to all SAPs of the mVPLS.

Every VLAN range that is not assigned within any of the created **mst-instance** is automatically assigned to mst-instance 0. This instance is automatically maintained by the software and cannot be modified. Changing the VLAN range value can be performed only when the given mst-instance is shutdown.

The **no** form of this command removes the **vlan-range** from given **mst-instance**.

Parameters *vlan-range* — The first VLAN range specifies the left-bound (i.e., minimum value) of a range of VLANs 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 associated with the mVPLS SAP.

Values 1— 4094

mst-max-hops

Syntax **mst-max-hops** *hops-count*
no mst-max-hops

Context config>service>vpls>stp

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the number of hops in the region before BPDU is discarded and the information held for the port is aged out. The root bridge of the instance sends a BPDU (or M-record) with remaining-hop-count set to configured <*max-hops*>. When a bridge receives the BPDU (or M-record), it decrements the received remaining-hop-count by 1 and propagates it in BPDU (or M-record) it generates.

The **no** form of this command sets the *hops-count* to its default value.

Default 20

Parameters *hops-count* — Specifies the maximum number of hops.

Values 1 — 40

mst-name

Syntax **mst-name** *region-name*
no mst-name

Context config>service>vpls>stp

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command defines an MST region name. Two bridges are considered as a part of the same MST region as soon as their configuration of the MST region name, the MST-revision and VLAN-to-instance assignment is identical.

The **no** form of this command removes *region-name* from the configuration.

Default no mst-name

Parameters *region-name* — Specifies an MST-region name up to 32 characters in length.

mst-revision

Syntax **mst-revision** *revision-number*

Context config>service>vpls>stp

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command defines the MST configuration revision number. Two bridges are considered as a part of the same MST region if their configured MST-region name, MST-revision, and VLAN-to-instance are identical.

The **no** form of this command returns MST configuration revision to its default value.

Default 0

Parameters *revision-number* — Specifies the MSTP region revision number to define the MSTP region.

Values 0 — 65535

path-cost

Syntax **path-cost** *sap-path-cost*
no path-cost

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the Spanning Tree Protocol (STP) path cost for the SAP.

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. When BPDUs are sent out other egress SAPs , the newly calculated root path cost is used.

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 7210 SAS the STP path cost is a purely static configuration.

The **no** form of this command returns the path cost to the default value.

path-cost — The path cost for the SAP.

Values 1 — 200000000 (1 is the lowest cost)

Default 10

port-num

Syntax [**no**] **port-num** *virtual-port-number*

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the virtual port number which uniquely identifies a SAP within configuration bridge protocol data units (BPDUs). The internal representation of a SAP is unique to a system and has a reference space much bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a SAP and identifies it with it's own virtual port number that is unique to every other SAP defined on the TLS. The virtual port number is assigned at the time that the SAP is added to the TLS. Since the order that the SAP was added to the TLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance.

The virtual port number cannot be administratively modifed.

priority

Syntax **priority** *stp-priority*
no priority

Context config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>sap>stp

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the Alcatel-Lucent Spanning Tree Protocol (STP) priority for the SAP.

STP priority is a configurable parameter associated with a SAP. When configuration BPDUs are received, the priority is used in some circumstances as a tie breaking mechanism to determine whether the SAPwill 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 SAPpriority. 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.

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 priority parameter. For instance, if a value of 0 is entered, masking out the lower 4 bits results in a parameter value of 0. If a value of 255 is entered, the result is 240.

The **no** form of this command returns the STP priority to the default value.

Default 128

Parameters *stp-priority* — The STP priority value for the SAP . 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

Syntax **sap** *sap-id* [**eth-ring** *ring-index*]
no sap *sap-id*

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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 7210 SAS. Each SAP must be unique.

A physical port can have only one SAP to be part of one service. Multiple SAPS can be defined over a physical port but each of these SAPs should belong to a different service.

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 (IES), the IP interface must be shutdown before the SAP on that interface may be removed.

This command is also used to create a Ring APS Control SAP or a Data SAP whose traffic is protected by a Ring APS Instance.

No SAPs are defined.

Special Cases 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).The 7210 SAS does not support explicit null encapsulation for VPLS service.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 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 — The keyword to create an instance of a Ring APS Control SAP or a Data SAP whose traffic is protected by a Ring APS Instance.

ring-index — Specifies the ring index of the Ethernet ring.

discard-unknown-source

Syntax [**no**] **discard-unknown-source**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

When this command is enabled, packets received on a SAP or with an unknown source MAC address will be dropped only if the maximum number of MAC addresses for that SAP (see max-nbr-mac-addr on page 464) has been reached. If max-nbr-mac-addr has not been set for the SAP, enabling discard-unknown-source has no effect.

When disabled, the packets are forwarded based on the destination MAC addresses.

The **no** form of this command causes packets with an unknown source MAC addresses to be forwarded by destination MAC addresses in VPLS.

Default **no discard-unknown-source**

ETH-CFM Service Commands

NOTES: The following contexts for the commands are supported as below:

- SAP context is supported for 7210 SAS-E, 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.
- Epipe spoke-sdp context is supported for 7210 SAS-K 2F4T6C only.
- VPLS mesh-sdp and spoke-sdp context is supported for 7210 SAS-K 2F4T6C only.

eth-cfm

Syntax **eth-cfm**

Context config>service>vpls>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables the context to configure ETH-CFM parameters.

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {up | **down**}]
no mep *mep-id* **domain** *md-index* **association** *ma-index*

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

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the ETH-CFM maintenance endpoint (MEP).

Parameters *mep-id* — Specifies the maintenance association end point identifier.

Values 1 — 8191

md-index — Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index — Specifies the MA index value.

Values 1 — 4294967295

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. This value cannot be used on 7210 SAS-E devices.

ais-enable

Syntax [**no**] **ais-enable**

Context config>service>vpls>mesh-sdp>eth-cfm>mep - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>epipe>spoke-sdp>eth-cfm>mep - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>sap>eth-cfm

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables the generation and the reception of AIS messages.

client-meg-level

Syntax **client-meg-level** [[*level* [*level* ...]]
no client-meg-level

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>sap>eth-cfm

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures the client maintenance entity group (MEG) level(s) to use for AIS message generation. Up to 7 levels can be provisioned with the restriction that the client MEG level must be higher than the local MEG level.

Parameters *level* — Specifies the client MEG level.

Values 1 — 7

Default 1

interval

Syntax **interval** {**1** | **60**}
no interval

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>sap>eth-cfm

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the transmission interval of AIS messages in seconds.

Parameters **1 | 60 —** The transmission interval of AIS messages in seconds.

Default 1

priority

Syntax **priority** *priority-value*
no priority

Context config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>sap>eth-cfm

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the priority of AIS messages originated by the node.

Parameters *priority-value* — 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>mep
config>service>vpls>sap>eth-cfm>mep
config>service>vpls>mesh-sdp>mep - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>epipe>spoke-sdp>eth-cfm>mep - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables the generation of CCM messages.

The **no** form of the command disables the generation of CCM messages.

ccm-ltm-priority

Syntax **ccm-ltm-priority** *priority*
no ccm-ltm-priority

Context config>service>vpls>sap>eth-cfm>mep
config>service>vpls>mesh-sdp>mep - Platforms Supported: 7210 SAS-K 2F4T6C

config>service>epipe>spoke-sdp>eth-cfm>mep - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the priority value for CCMs and LTMs transmitted by the MEP.

The **no** form of the command removes the priority value from the configuration.

Default The highest priority on the bridge-port.

Parameters *priority* — Specifies the priority of CCM and LTM messages.

Values 0 — 7

eth-test-enable

Syntax [**no**] **eth-test-enable**

Context config>service>vpls>sap>eth-cfm>mep
config>service>vpls>spoke-sdp>eth-cfm>mep - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

For ETH-test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands:

oam eth-cfm eth-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [data-length *data-length*]

A check is done for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP will indicate the problem.

test-pattern

Syntax **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
no test-pattern

Context config>service>vpls>sap>eth-cfm>mep>eth-test-enable
config>service>vpls>mesh-sdp>eth-cfm>mep>eth-test-enable - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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** <errors>
no bit-error-threshold

Context config>service>vpls>sap>eth-cfm>mep

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command is used to specify the threshold value of bit errors.

fault-propagation-enable

Syntax **fault-propagation-enable {use-if-tlv | suspend-ccm}**
no fault-propagation-enable

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

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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.

mac-address

Syntax **mac-address** *mac-address*
no mac-address

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

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the MAC address of the MEP.

The **no** form of this command reverts the MAC address of the MEP back to that of the port (if the MEP is on a SAP) or the bridge (if the MEP is on a spoke).

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

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MEP. Must be unicast. Using the all zeros address is equivalent to the no form of this command.

low-priority-defect

Syntax **low-priority-defect** {**allDef**|**macRemErrXcon**|**remErrXcon**|**errXcon**|**xcon**|**noXcon**}

Context config>service>epipe>spoke-sdp>eth-cfm>mep - Platforms Supported: 7210 SAS-K

2F4T6C
config>service>vpls>mesh-sdp>eth-cfm>mep - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>epipe>sap>eth-cfm>mep>eth-test-enable
config>service>epipe>sap>eth-cfm>mep

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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** *mac-address*
no mac-address

Context config>service>vpls>sap>eth-cfm>mep
config>service>vpls>mesh-sdp>eth-cfm>mep - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>epipe>spoke-sdp>eth-cfm>mep - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the MAC address of the MEP.

The **no** form of this command reverts the MAC address of the MEP back to that of the port (if the MEP is on a SAP) or the bridge (if the MEP is on a spoke).

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

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MEP. Must be unicast. Using the all zeros address is equivalent to the no form of this command.

one-way-delay-threshold

Syntax **one-way-delay-threshold** *seconds*

Context config>service>vpls>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables/disables eth-test functionality on MEP.

Parameters *seconds* — Specifies the one way delay threshold, in seconds.

Values 0..600

Default 3

mip

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

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

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command allows Maintenance Intermediate Points (MIPs) to be created if mhf-creation for the MA is configured using the default option.

For more information on the ETH-CFM matrix support for different platforms, see the tables below:

Table 9, “ETH-CFM Support Matrix for 7210 SAS-D,” on page 94

Table 10, “ETH-CFM Support Matrix for 7210 SAS-E,” on page 95

Table 11, “ETH-CFM Support Matrix for 7210 SAS-K 2F2T1C,” on page 95

Table 12, “ETH-CFM Support Matrix for 7210 SAS-K 2F4T6C,” on page 96

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

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MIP. The MAC must be unicast. Using the all zeros address is equivalent to the no form of this command.

default-mac — Using the no command deletes the MIP. If the operator wants to change the mac back to the default mac without having to delete the MIP and reconfiguring this command is useful.

Default no mip

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.

NOTE: This parameter is supported only with UP MEP on 7210 SAS-D with a Dot1q range SAP.

Values 0 — 4094

VPLS COMMANDS

limit-mac-move

Syntax **limit-mac-move** [**blockable** | **non-blockable**]
no limit-mac-move

Context config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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>sap
config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>endpoint
config>service>pw-template - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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 with mac-pinning enabled will remain in the FIB on this SAP forever. Every event that would otherwise result in re-learning will be logged (MAC address; original-SAP; new-SAP).

max-nbr-mac-addr

Syntax **max-nbr-mac-addr** *table-size*
no max-nbr-mac-addr

Context config>service>vpls>sap
config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>endpoint
config>service>pw-template - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this SAP.

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 455), packets with unknown source MAC addresses will be discarded.

The **no** form of the command restores the global MAC learning limitations for the SAP.

Default no max-nbr-mac-addr

Parameters *table-size* — Specifies the maximum number of learned and static entries allowed in the FDB of this service.

Values [1..30719] SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C
[1..16383]- SAS-E

statistics

Syntax **statistics**

Context config>service>vpls>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C.

This command enables the context to configure the counters associated with SAP ingress and egress.

egress

Syntax **egress**

Context config>service>vpls>sap>statistics

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables the context to configure the egress SAP statistics counter and set the mode of the counter.

This counter counts the number of packets forwarded through the SAP.

ingress

Syntax **ingress**

Context config>service>epipe>sap>statistics
config>service>vpls>sap>statistics

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C.

This command enables the context to configure the ingress SAP statistics counter.

For 7210 E, by default, SAP ingress counters are associated with a SAP and cannot be disabled.

In 7210 SAS-D devices, for access-uplink SAPs the ingress counters are not enabled by default. For access SAPs if the ingress counter is enabled by default, it can be disabled.

The two types of ingress SAP counters are:

- A counter that counts the total packets or octets received on the SAP
- A counter associated with meters defined in the QoS policy of the SAP. This counter counts the in-profile and out-of-profile packets or octets received on the SAP.

forwarded-count

Syntax [**no**] **forwarded-count**

Context config>service>vpls>sap>statistics>egress
config>service>ies>sap>statistics>egress

Platform supported: 7210 SAS-D

Description This command associates a counter with the SAP. The counter counts the number of packets forwarded through the SAP.

A limited amount of such counters are available for use with access SAPs and access-uplink SAPs.

Use this command before enabling applicable accounting record collection on the SAP to associate a counter with the SAP.

The **no** form of this command disables the packet count.

packets-forwarded-count

Syntax [**no**] **packets-forwarded-count**

Context config>service>vpls>sap>statistics>egress

Platform supported: 7210 SAS-E

Description This command associates a counter with the SAP. The counter counts the number of packets forwarded through the SAP.

A limited amount of such counters are available for use with access SAPs and access-uplink SAPs.

Use this command before enabling applicable accounting record collection on the SAP to associate a counter with the SAP.

The **no** form of this command disables the packet count.

counter-mode

Syntax **counter-mode {in-out-profile-count| forward-drop-count}**

Context config>service>epipe>sap>statistics>ingress
config>service>vpls>sap>statistics>ingress

Description **Platforms Supported:** 7210 SAS-D.

This command allows the user to set the counter mode for the counters associated with sap ingress meters (a.ka. policers). A pair of counters is available with each meter. These counters count different events based on the counter mode value.

Note: The counter mode can be changed if an accounting policy is associated with a SAP. If the counter mode is changed the counters associated with the meter are reset and the counts are cleared. If an accounting policy is in use when the counter-mode is changed a new record will be written into the current accounting file.

Execute the following sequence of commands to ensure a new accounting file is generated when the counter-mode is changed:

1. Execute the command **config>service>epipe/vpls>sap> no collect-stats**, to disable writing of accounting records.
2. Change the counter-mode to the desired value, execute the command **config>service>epipe/vpls>sap>counter-mode {in-out-profile-count| forward-drop-count}**.
3. Execute the command **config>service>epipe/vpls>sap> collect-stats**, to enable writing of accounting records.

The **no** form of the command restores the counter mode to the default value.

Default in-out-profile-count

Parameters **forward-drop-count** — If the counter mode is specified as "forward-drop-count", one counter counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

in-out-profile-count — If the counter mode is specified as "in-out-profile-count", one counter counts the total in-profile packets and octets received on ingress of a SAP and another counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

static-mac

Syntax [**no**] **static-mac** *ieee-mac-address* [**create**]

Context config>service>vpls>sap
config>service>vpls>mesh-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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 **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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 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 **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

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 **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command configures a range of VLANs on an access port 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 390.

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 **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables the context to configure egress filter policies.

If **no** egress filter is defined, no filtering is performed.

ingress

Syntax **ingress**

Context config>service>vpls>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

filter

Syntax **filter ip** *ip-filter-id*
filter ipv6 *ipv6-filter-id*
filter mac *mac-filter-id*

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress
config>service>ies>sap>egress
config>service>ies>sap>ingress

Description **Platforms Supported:** 7210 SAS-E (only ingress context), 7210 SAS-D (both ingress and egress context), 7210 SAS-K 2F2T1C (both ingress and egress context), 7210 SAS-K 2F4T6C (both ingress and egress context).

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

Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria. There are two types of filter policies: IP and MAC. Only one type may be applied to a SAP at a time.

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

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system.

Special Cases **VPLS —** Both MAC and IP filters are supported on a VPLS service SAP.

Parameters **ip** *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 — 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

qos

Syntax **qos** *policy-id*
no qos

Context config>service>vpls>sap>ingress
config>service>ies>sap>ingress

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP) or IP interface.

QoS ingress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate ingress policies. The **qos** command only allows ingress policies to be associated on SAP ingress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress 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.

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

policy-id — The ingress policy ID to associate with SAP on ingress. The policy ID must already exist.

Values 1 — 65535

agg-shaper-rate

Syntax **agg-shaper-rate** *agg-rate*
no agg-shaper-rate

Context config>service>vpls>sap>ingress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.
This command allows user to specify the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues. User can specify the CIR rate and the PIR rate. User must not oversubscribe the total bandwidth available for use by ingress queues.
The no form of the command disables the use of SAP aggregate rate shaper. In other words, the SAP can use up the maximum bandwidth available.

Default no agg-shaper-rate

Parameters ***agg-rate*** **—** Specifies the rate in kilobits per second.

Values 50 — 3000000 | max

Default max

agg-shaper-rate

Syntax **agg-shaper-rate** *agg-rate*
no agg-shaper-rate

Context config>service>vpls>sap>egress

Description **Platforms Supported:** 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.
This command allows user to specify the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues. User can specify the CIR rate and the PIR rate. User must not oversubscribe the total bandwidth available for use by ingress queues.
The no form of the command disables the use of SAP aggregate rate shaper. In other words, the SAP can use up the maximum bandwidth available.

Default no agg-shaper-rate

Parameters ***agg-rate*** **—** Specifies the rate in kilobits per second.

Values 50 — 1000000 | max

Default max

aggregate-meter-rate

Syntax **aggregate-meter-rate** *rate-in-kbps* [**burst** *burst-in-kbits*]
no aggregate-meter-rate

Context config>service>vpls>sap>ingress
config>service>ies>if>sap>ingress
config>service>epipe>sap>ingress

Description **Platforms Supported:** 7210 SAS-D.

This command allows the user to configure the SAP ingress aggregate policer. The rate of the SAP ingress aggregate policer must be specified by the user. The user can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the ingress traffic on different FCs and determines the final disposition of the packet. The packet is either forwarded to an identified profile or dropped.

Note: The sum of CIR of the individual FCs configured under the SAP cannot exceed the PIR rate configured for the SAP. Although the 7210 SAS software does not block this configuration, it is not recommended for use.

The table below provides information about the final disposition of the packet based on the operating rate of the per-FC policer and the per-SAP aggregate policer:

Per FC meter Operating Rate	Per FC Assigned Color	SAP aggregate meter Operating Rate	SAP aggregate meter color	Final Packet Color
Within CIR	Green	Within PIR	Green	Green or In-profile
Within CIR*	Green	Above PIR	Red	Green or In-profile
Above CIR, Within PIR	Yellow	Within PIR	Green	Yellow or Out-of-Profile
Above CIR, Within PIR	Yellow	Above PIR	Red	Red or Dropped
Above PIR	Red	Within PIR	Green	Red or Dropped
Above PIR	Red	Above PIR	Red	Red or Dropped

Table 21: Final Disposition of the packet based on per FC and per SAP policer or meter.

Note*: Row number 2 in the above table is not recommended for use. For more information on this, see the Note in the **aggregate-meter-rate** description.

When the SAP aggregate policer is configured, the per-FC policer can be only configured in “trtcm2” mode (RFC 4115).

Note: The meter modes “srtcm” and “trtcm1” are used in the absence of an aggregate meter.

The SAP ingress meter counters increment the packet or octet counts based on the final disposition of the packet.

If ingress frame-based accounting is used, the SAP aggregate meter rate accounts for the Ethernet frame overhead. The system accounts for 12 bytes of IFG and 8 bytes of start delimiter.

The **no** form of the command removes the aggregate policer from use.

Default no aggregate-meter-rate

Parameters *rate-in-kbps* — Specifies the rate in kilobits per second.

Values 0 — 20000000 | max

Default max

burst-in-kilobits — Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 —2146959

Default 512

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>sap

Note: This command is not applicable for access uplink SAPs.

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command creates the accounting policy context that can be applied to a SAP.
An accounting policy must be defined before it can be associated with a SAP.
If the *policy-id* does not exist, an error message is generated.
A maximum of one accounting policy can be associated with a SAP at one time. Accounting policies are configured in the **config>log** context.

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

Default Default accounting policy.

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 — 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>vpls>spoke-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>sap

Note: This command is not applicable for access uplink SAPs.

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C.

This command enables accounting and statistical data collection for either the SAP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default no collect-stats

VPLS SDP Commands

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}]
no mesh-sdp *sdp-id*[:*vc-id*]

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

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.

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.

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

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

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

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

Special Cases **VPLS —** Several SDPs can be bound to a VPLS. Each SDP must be destined to a different router. If two *sdp-id* bindings terminate on the same router, an error occurs and the second SDP is binding is rejected.

Parameters *sdp-id* — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit identifier. This value is used to validate the VC ID portion of each mesh SDP binding defined in the service. The default value of this object is equal to the service ID.

Values 1 — 4294967295

vc-type — This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled.
VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

- The VC type value for Ethernet is 0x0005.

- The VC type value for an Ethernet VLAN is 0x0004.

ether — Defines the VC type as Ethernet. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke SDP binding. (hex 5)

vlan — Defines the VC type as VLAN. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for mesh SDP bindings.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [create] [**split-horizon-group** *group-name*] **endpoint** *no-endpoint*
no spoke-sdp *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [create] **endpoint** *no-endpoint*

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command binds a service to an existing Service Distribution Point (SDP). The SDP defines the transport tunnel that this VPRN service is bound to.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

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

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

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

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

Special Cases **VPLS —** Several SDPs can be bound to a VPLS service. Each SDP must use unique *vc-ids*. An error message is generated if two SDP bindings with identical *vc-ids* terminate on the same router. Split horizon groups can only be created in the scope of a VPLS service.

Parameters *sdp-id* — The SDP identifier.

Values 1 — 17407

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

create — This keyword is mandatory while creating a spoke SDP.

endpoint — Specifies the service endpoint to which this SDP bind is attached. The service ID of the SDP binding must match the service ID of the service endpoint.

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

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)

split-horizon-group *group-name* **—** Specifies the name of the split horizon group to which the SDP belongs.

vc-type — This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled.
VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

- The VC type value for Ethernet is 0x0005.
- The VC type value for an Ethernet VLAN is 0x0004.

Values ether, vlan

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 requires at least one dot1Q tag within each encapsulated Ethernet packet transmitted to the far end.

control-word

Syntax [**no**] **control word**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

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.

Default no control word

egress

Syntax **egress**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures the egress SDP context.

ingress

Syntax **ingress**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures the ingress SDP context.

vc-label

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

Context config>service>vpls>mesh-sdp>egress
config>service>vpls>spoke-sdp>egress

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures the egress VC label.

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

Values 16 — 1048575

vc-label

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

Context config>service>vpls>mesh-sdp>ingress
config>service>vpls>spoke-sdp>ingress

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures the ingress VC label.

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

Values 2048 — 18431

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 **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command specifies an explicit Dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured Dot1q tag can be overridden by a received TLV specifying the Dot1q value expected by the far end. This signaled value must be stored as the remote signaled Dot1q value for the binding. The provisioned local Dot1q tag must be stored as the administrative Dot1q value for the binding.

When the Dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The **no** form of this command disables the command.

Default no vlan-vc-tag

Parameters *0..4094* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

IGMP Snooping Commands

disable-router-alert-check

Syntax [**no**] **disable-router-alert-check**

Context config>service>vpls>sap>igmp-snooping

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command enables the IGMP router alert check option.

The no form of the command disables the router alert check.

description

Syntax **description** *description-string*
no description

Context config>service>vpls>igmp-snooping>mvr

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the content in the configuration file.

The no form of this command removes the string from the configuration.

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

fast-leave

Syntax [**no**] **fast-leave**

Context config>service>vpls>sap>igmp-snooping

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command enables fast leave. When IGMP fast leave processing is enabled, the 7210 SAS E will immediately remove a SAP from the multicast group when it detects an IGMP “leave” on that SAP. Fast leave processing allows the switch to remove a SAP that sends a 'leave' from the forwarding table without first sending out group-specific queries to the SAP, and thus speeds up the process of changing channels ('zapping').

Fast leave should only be enabled when there is a single receiver present on the SAP.
When fast leave is enabled, the configured last-member-query-interval value is ignored.

Default no fast-leave

from-vpls

Syntax **from-vpls** *service-id*
no from-vpls

Context config>service>vpls>sap>igmp-snooping>mvr

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

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 *service-id* — Specifies the MVR VPLS from which multicast channels should be copied into this SAP.

Values *service-id*: 1 — 2147483648

group

Syntax [**no**] **group** *grp-address*

Context config>service>vpls>sap>igmp-snooping>static

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command adds a static multicast group as a (*, g). When a static IGMP group is added, multicast data for that (*,g) or (s,g) is forwarded to the specific SAP without receiving any membership report from a host.

Default none

Parameters *grp-address* — Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

group-policy

Syntax **group-policy** *policy-name*
no group-policy

Context config>service>vpls>igmp-snooping>mvr

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command identifies filter policy of multicast groups to be applied to this VPLS entity. The sources of the multicast traffic must be a member of the VPLS.

The **no** form of the command removes the policy association from the VPLS configuration.

Default No group policy is specified.

Parameters *policy-name* — 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.

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
config>service>pw-template

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command forces vc-vlan-type forwarding in the data path for spoke/mesh SDPs which have either vc-type. This command is not allowed on vlan-vc-type SDPs.
The no form of this command sets default behavior.

Default disabled

Description

igmp-snooping

Syntax **igmp-snooping**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp **-** 7210 SAS-K 2F4T6C.
config>service>vpls>mesh-sdp**-** 7210 SAS-K 2F4T6C.
config>service>pw-template**-** 7210 SAS-K 2F4T6C.

Description This command enables the Internet Group Management Protocol (IGMP) snooping context.

Default none

max-num-sources

Syntax **max-num-sources** *max-num-sources*
no max-num-sources

Context config>service>vpls>sap>igmp-snooping

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command configures the maximum number of multicast sources allowed per group.

The **no** form of the command removes the value from the configuration.

Parameters *max-num-sources* — Specifies the maximum number of multicast sources allowed per group.

Values 1 — 2047 - 7210 SAS-D and 7210 SAS-E
[1..900] - 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

import

Syntax **import** *policy-name*
no import

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config> service>vpls> mesh-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>pw-template>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP or SDP at any time.

The **no** form of the command removes the policy association from the SAP or SDP.

Default **no import** — No import policy is specified.

Parameters *policy-name* — The import policy name. Values can be string up to 32 characters long of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. These policies are configured in the **config>router> policy-options** context The router policy must be defined before it can be imported.

last-member-query-interval

Syntax **last-member-query-interval** *tenths-of-seconds*
no last-member-query-interval

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K

2F4T6C
config>service>vpls>mesh-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>pw-template>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command configures the maximum response time used in group-specific queries sent in response to 'leave'messages, and is also the amount of time between 2 consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.
The configured last-member-query-interval is ignored when fast-leave is enabled on the SAP.

Default 10

Parameters *seconds* — Specifies the frequency, in tenths of seconds, at which query messages are sent.

Values 1 — 50

max-num-groups

Syntax **max-num-groups** *count*
no max-num-groups

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>pw-template>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

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.

Default no max-num-groups

Parameters *count* — Specifies the maximum number of groups that can be joined on this SAP.

Values 1 — 1024
[1..900] - 7210 SAS-K

mrouter-port

Syntax [**no**] **mrouter-port**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C

config>service>vpls>mesh-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command specifies whether a multicast router is attached behind this SAP.

Configuring a SAP or SDP 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 (v1or 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 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>sap>igmp-snooping

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

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 - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>pw-template>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

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 **config>service>vpls>igmp-snooping, config>service>vpls>spoke-sdp>igmp-snooping, config>service>vpls>mesh-sdp>igmp-snooping:** 1 - 65535
config>service>vpls>sap>igmp-snooping, config>service>pw-template>igmp-snooping: 2 - 1024

query-src-ip

Syntax **query-src-ip** *ip-address*
no query-src-ip

Context config>service>vpls>igmp-snooping

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

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 - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>pw-template>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command configures the IGMP query response interval. If the **send-queries** command is enabled, this parameter specifies the maximum response time advertised in IGMP queries.

The configured query-response-interval must be smaller than the configured query-interval.

If send-queries is not enabled on this SAP or SDP, the configured query-response-interval value is ignored.

Default 10

Parameters *seconds* — Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 — 1023

report-src-ip

Syntax **report-src-ip** *ip-address*
no report-src-ip

Context config>service>vpls>igmp-snooping

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This parameter specifies the source IP address used when generating IGMP reports. According the IGMPv3 standard, a zero source address is allowed in sending IGMP reports. However, for interoperability with some multicast routers, the source IP address of IGMP group reports can be configured using this command.

Default a.b.c.d

Parameters *ip-address* — The source IP source address in transmitted IGMP reports.

report-src-ip

Syntax **report-src-ip**-*address*
no report-src-ip

Context config>service>vpls>igmp-snooping

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This parameter specifies the source IP address used when generating IGMP reports. According the IGMPv3 standard, a zero source address is allowed in sending IGMP reports. However, for interoperability with some multicast routers, the source IP address of IGMP group reports can be configured using this command.

Default 0.0.0.0

Parameters *ip-address* — The source IP source address in transmitted IGMP reports.

robust-count

Syntax **robust-count** *robust-count*
no robust-count

Context config>service>vpls>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>pw-template>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

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.

Values **config>service>vpls>sap>igmp-snooping**: 2— 7
config>service>vpls>igmp-snooping: 1 — 255

precedence

Syntax **precedence** *precedence-value*| **primary**
no precedence

Context config>service>vpls>spoke-sdp

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command configures the spoke SDP precendence.

Default 4

Parameters *precedence-value* — Specify the spoke SDP precedence.

Values 0 — 4

primary — Specifies that the prececence is primary.

propagate-mac-flush

Syntax [**no**] **propagate-mac-flush**

Context config>service>vpls

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

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 - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>pw-template>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

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.

Default no send-queries

group

Syntax **[no] group** *grp-address*

Context config>service>vpls>sap>igmp-snooping>static
config>service>vpls>spoke-sdp>snooping>static - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp>snooping>static - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command adds a static multicast group either as a (*, g) or as one or more (s,g) records. When a static 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.

NOTE: SDP is supported only on 7210 SAS-K 2F4T6C.

Default none

Parameters *grp-address* — Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

source

Syntax [**no**] **source** *ip-address*

Context config>service>vpls>sap>igmp-snooping>static>group

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

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 unicast address.

starg

Syntax [**no**] **starg**

Context config>service>vpls>sap>igmp-snooping>static>group
config>service>vpls>spoke-sdp>igmp-snooping>static>group - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp>igmp-snooping>static>group - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

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 - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

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) 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>sap>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>spoke-sdp>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>vpls>mesh-sdp>snooping>static - Platforms Supported: 7210 SAS-K 2F4T6C
config>service>pw-template>igmp-snooping - Platforms Supported: 7210 SAS-K 2F4T6C

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command specifies the version of IGMP which is running on this SAP. This object can be used to configure a router capable of running either value. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP 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 orwill be the version of the querier.

Parameters *version* — Specify the IGMP version.

Values 1, 2 or 3

to-sap

Syntax **to-sap** *sap-id*
no to-sap

Context config>service>vpls>sap>igmp-snooping>mvr

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, 7210 SAS-K 2F2T1C, and 7210 SAS-K 2F4T6C

This command configures the SAP to which the multicast data needs to be copied.

In some scenarios, 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.

Default no to-sap

Parameters *sap-id* — Specifies the SAP to which multicast channels should be copied.

Internet Enhanced Service on 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

In This Chapter

This chapter provides information about Internet Enhanced Services the process overview, and implementation notes.

Topics in this chapter include:

- IES Service Overview on page 494
- IES Features on page 496
  - → IP Interfaces on page 496
    - – Subscriber Interfaces on page 315
    - – Encapsulations on page 498
  - → Routing Protocols on page 511
    - – CPE Connectivity Check on page 498
  - → QoS Policies on page 499
  - → Filter Policies on page 500
- Configuring an IES Service with CLI on page 501
- Basic Configuration on page 502
- Common Configuration Tasks on page 503
- Service Management Tasks on page 507

IES Service Overview

NOTE: On 7210 SAS-D, 7210 SAS-E and 7210 SAS-K 2F2T1C, IES is designed for in-band management of the node. IES service (standalone, without being associated with Routed VPLS) is not supported on 7210 SAS-K 2F2T1C.

Internet Enhanced Service (IES) is a routed connectivity service where the subscriber communicates with an IP router interface to send and receive Internet traffic. An IES has one or more logical IP routing interfaces each with a SAP which acts as the access point to the subscriber's network.

IES allows IP interfaces to participate in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber be unique between other provider addressing schemes and potentially the entire Internet. While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be reserved for service IP provisioning, and be administered by a separate, but subordinate address authority.

IP interfaces defined within the context of an IES service must have a SAP associated as the uplink access point to the subscriber network. Multiple IES services are created to segregate subscriber owned IP interfaces.

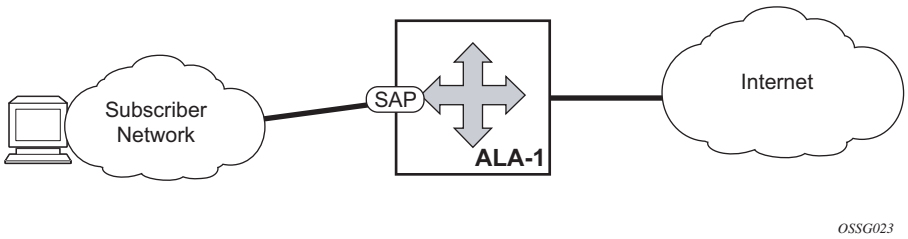

Figure 66: Internet Enhanced Service

The IES service provides in-band management connectivity. Other features include:

- Multiple IES services are created to separate IP interfaces.
- More than one IES service can be created for a single customer ID.
- More than one IP interface can be created within a single IES service ID. All IP interfaces created within an IES service ID belong to the same customer.

On 7210 SAS-D and 7210 SAS-E, the IES services provide IP connectivity to the node for in-band management of the node. Most of the management tasks supported with the out-of-band management port are supported with in-band management.

IES Features

This section describes various general service features and any special capabilities or considerations as they relate to IES services.

IP Interfaces

IES customer IP interfaces can be configured with most of the options found on the core IP interfaces. The advanced configuration options supported are:

- ICMP Options

IPv6 support for IES IP interfaces associated with Access-Uplink SAPs (on 7210 SAS-D)

In access-uplink mode, IES IP interfaces associated with access-uplink SAPs support IPv6 addressing. IPv6 can be used for in-band management of the node using the IES IP interface.

NOTE: IPv6 IES IP interfaces on access-uplink SAPs is supported only on 7210 SAS-D. It is not supported for 7210 SAS-E and 7210 SAS-K 2F2T1C.

In 7210 SAS-D, IPv4 and IPv6 route table lookup entries are shared. Before adding routes for IPv6 destinations, route entries in the routed lookup table needs to be allocated for IPv6 addresses. This can be done using the CLI command *config> system> resource-profile> router> max-ipv6-routes*. This command allocates route entries for /64 IPv6 prefix route lookups. The system does not allocate any IPv6 route entries by default and user needs to allocate some resources before using IPv6. For the command to take effect, the node must be rebooted after making the change. For more information, see the example below and the 7210 SAS Basic System Configuration Guide.A separate route table is used for IPv6 /128-bit prefix route lookup. A limited number of IPv6 /128 prefixes route lookup entries is supported on 7210 SAS-D. The software enables lookups in this table by default (in other words no user configuration is required to enable IPv6 / 128-bit route lookup).

NOTE: In 7210 SAS-D, IPv6 interfaces are created without allocating IPv6 route entries.

Following features and restrictions is applicable for IPv6 IES IP interfaces:

- IPv6 interfaces supports only static routing.
- Only port-based ingress QoS policies are supported.
- IPv6 filter policies can be used on SAP ingress and egress.
- Routing protocols, such as OSPFv3, and others are not supported.
- A limited number of IPv6 /128 prefixes route lookup entries is supported.

SAPs

See the following sections:

- Encapsulations on page 498
- CPE Connectivity Check on page 498
- QoS Policies on page 499
- CPU QoS for IES interfaces in access-uplink mode on page 499
- Filter Policies on page 500

Encapsulations

In 7210 SAS-D, the following Access SAP encapsulation is supported on IES services in both network mode and access-uplink mode:

- Ethernet null
- Ethernet dot1q
- Ethernet QinQ

NOTE: 7210 SAS-E does not support access SAP based IES interfaces.

In 7210 SAS-D and 7210 SAS-E, the following access-uplink SAP encapsulations are supported:

- Ethernet QinQ (access-uplink QinQ SAP)

CPE Connectivity Check

NOTE: This service is supported only in 7210 SAS-D.

Static routes are used within many IES services. Unlike dynamic routing protocols, there is no way to change the state of routes based on availability information for the associated CPE. CPE connectivity check adds flexibility so that unavailable destinations will be removed from the service provider's routing tables dynamically and minimize wasted bandwidth.

The availability of the far-end static route is monitored through periodic polling. The polling period is configured. If the poll fails a specified number of sequential polls, the static route is marked as inactive.

An ICMP ping mechanism is used to test the connectivity. If the connectivity check fails and the static route is deactivated, the router will continue to send polls and re-activate any restored routes.

QoS Policies

When applied to 7210 SAS IES services, service ingress QoS policies only create the unicast meters defined in the policy. The multi-point meters are not created on the service. With IES services, service egress QoS policies function as with other services where the class-based queues are created as defined in the policy.

On 7210 SAS ingress, only meters are supported on all the platforms.

Note:

- QoS policies only create the unicast meters defined in the policy if PIM is not configured on the associated IP interface; if PIM is configured, the multipoint meters are applied as well.
- On the 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C, both Unicast and Multicast Queues are configured.

In access-uplink mode, IES IP interface associated with an access SAP supports use of service ingress QoS policies. IES IP interface associated with an access-uplink SAP does not support use of service ingress QoS policies. IES IP interfaces associated with an access-uplink SAP share the port based ingress and egress QoS policies.

Note that both MAC and IPv4 criteria can be used in the QoS policies for traffic classification in an IES.

CPU QoS for IES interfaces in access-uplink mode

In access-uplink mode, IES IP interface bound to routed VPLS services, IES IP interface on access SAPs and IES IP interface on Access-Uplink SAPs are designed for use with inband management of the node. Consequently, they share a common set of queues for CPU bound management traffic. All CPU bound traffic is policed to pre-defined rates before being queued into CPU queues for application processing. The system uses meters per application or a set of applications. It does not allocate meters per IP interface. The possibility of CPU overloading has been reduced by use of these mechanisms. Users must use appropriate security policies either on the node or in the network to ensure that this does not happen.

Filter Policies

In access-uplink mode, only IP filter policies can be applied to IES service when either access SAP or access-uplink SAP is associated with the service.

Configuring an IES Service with CLI

This section provides information to configure IES services using the command line interface.

Topics in this section include:

- Basic Configuration on page 502
- Common Configuration Tasks on page 503
  - → Configuring IES Components on page 504
    - – Configuring an IES Service on page 504
    - – Configuring IES Interface Parameters on page 504
      - Configuring SAP Parameters on page 506
      - Configuring VRRP on page 517
- Service Management Tasks on page 507
  - → Modifying IES Service Parameters on page 507
  - → Deleting an IES Service on page 508
  - → Disabling an IES Service on page 508
  - → Re-Enabling an IES Service on page 508

Basic Configuration

The most basic IES service configuration has the following entities:

- Customer ID (refer to Configuring Customers on page 79)
- An interface to create and maintain IP routing interfaces within IES service ID.
- A SAP on the interface specifying the access port and encapsulation values.

The following example displays a sample configuration of an IES service on ALA-48 on an access-uplink SAP (applicable for access-uplink mode only).

```
*A:ALA-48>config>service# info
----------------------------------------------
     ies 1000 customer 50 create
            description "to internet"
            interface "to-web" create
                address 10.1.1.1/24
                sap 1/1/5:0.* create
                exit
            exit
            no shutdown
----------------------------------------------
*A:ALA-48>config>service#
```

The following example displays a sample configuration of an IES service on ALA-50.

```
*A:ALA-50>config>service# info
----------------------------------------------
ies 1000 customer 50 vpn 1000 create
     description "to internet"
     interface "to-web" create
          address 10.1.1.1/24
               sap 1/1/10:100 create
          exit
     exit
     no shutdown
----------------------------------------------
*A:ALA-50>config>service#
```

The following example displays a basic IES service configuration for IPv6, along with the use of max-ipv6-routes.

The following displays an example of allocation of IPv6 routes on the node.

```
*A:7210SAS>config>system>res-prof# info
----------------------------------------------
            max-ipv6-routes 1000
----------------------------------------------
```

NOTE: The node must be rebooted after the above change.

```
*A:7210SAS>config>service# info
----------------------------------------------
     ies 1000 customer 50 vpn 1000 create
          description "to inband-mgmt"
          interface "to-mgmt" create
               ipv6
                    address 10::1/24
                         sap 1/1/10:100 create
                    exit
          exit
     no shutdown
----------------------------------------------
*A:7210SAS>config>service#
```

Common Configuration Tasks

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

1. Associate an IES service with a customer ID.
2. Associate customer ID with the service.
3. Assign an IP address.
4. Create an interface.
5. Define SAP parameters on the interface
   → Select node(s) and port(s).
   → Optional — select filter policies (configured in the **config>filter** context).
6. Enable service.

Configuring IES Components

Use the CLI syntax to configure the following entities:

- Configuring an IES Service on page 504
- Configuring IES Interface Parameters on page 504
  - → Configuring SAP Parameters on page 506
  - → Configuring VRRP on page 517

Configuring an IES Service

Use the following CLI syntax to create an IES service:

The following example displays a basic IES service configuration.

```
A:ALA-48>config>service#
------------------------------------------------
...
     ies 1001 customer 1730 create
            description "to-internet"
            no shutdown
     exit
------------------------------------------------
A:ALA-48>config>service#
```

Configuring IES Interface Parameters

The following example displays an IES configuration with interface parameters in access-uplink mode:

```
*A:7210-SAS>config>service>ies>if# info
------------------------------------------------
     arp-timeout 10000
     allow-directed-broadcasts
     icmp
          ttl-expired 120 38
     exit
     ip-mtu 1000
------------------------------------------------
*A:7210-SAS>config>service>ies>if#
```

Configuring IES SAP Parameters

A SAP is a combination of a port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique within a router.

When configuring IES access SAP parameters, a default QoS policy is applied to each SAP ingress. Additional QoS policies must be configured in the config>qos context. Filter policies are configured in the config>filter context and must be explicitly applied to a SAP. There are no default filter policies.

This example displays an IES SAP configuration.

```
-----------------------------------------------
*A:ALA-A>config>service>ies>if# info
-----------------------------------------------
address 10.10.36.2/24
sap 1/1/3:100 create
ingress
qos 101
exit
exit
-----------------------------------------------
*A:ALA-A>config>service>ies>if#
```

Configuring SAP Parameters

A SAP is a combination of a port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique within a router.

When configuring IES access SAP parameters, a default QoS policy is applied to each SAP ingress . Additional QoS policies must be configured in the config>qos context.

Filter policies are configured in the config>filter context and must be explicitly applied to a SAP. There are no default filter policies.

Only in 7210 SAS-D, SAP ingress Qos policy is supported only for access SAPs. It is not supported for access-uplink SAP.

Access-uplink SAPs (on both 7210 SAS-E and 7210 SAS-D) use the port based ingress and egress QoS policies.

This example displays an IES SAP configuration.

```
-----------------------------------------------
*A:ALA-A>config>service>ies>if# info
-----------------------------------------------
     address 10.10.36.2/24
     sap 1/1/3:100 create
          ingress
               qos 101
          exit
     exit
-----------------------------------------------
*A:ALA-A>config>service>ies>if#
```

Service Management Tasks

This section discusses the following service management tasks:

- Modifying IES Service Parameters on page 507
- Deleting an IES Service on page 508
- Disabling an IES Service on page 508
- Re-Enabling an IES Service on page 508

Modifying IES Service Parameters

Existing IES service parameters in the CLI or NMS can be modified, added, removed, enabled or disabled. The changes are applied immediately to all services when the charges are applied.

To display a list of customer IDs, use the **show service customer** command.
Enter the parameter(s) (such as description and SAP information) and then enter the new information.

The following displays the modified service:

```
*A:ALA-A>config>service>ies# info
----------------------------------------------
     ies 1000 customer 50 create
           description "This is a new description"
           interface "to-web" create
               address 10.1.1.1/24
               mac 00:dc:98:1d:00:00
               sap 1/1/5:0.* create
               exit
           exit
           no shutdown
     exit
----------------------------------------------
*A:ALA-A>config>service#
```

Deleting an IES Service

An IES service cannot be deleted until SAPs and interfaces are shut down *and* deleted and the service is shutdown on the service level.

Use the following CLI syntax to delete an IES service:

CLI Syntax:
```
config>service#
        [no] ies service-id
        shutdown
           [no] interface ip-int-name
           shutdown
              [no] sap sap-id
              shutdown
```

Disabling an IES Service

An IES service can be shut down without deleting the service parameters.

CLI Syntax:
```
config>service> ies service-id
        shutdown
```

Re-Enabling an IES Service

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

CLI Syntax:
```
config>service> ies service-id
        [no] shutdown
```

Example:
```
config>service# ies 2000
config>service>ies# no shutdown
config>service>ies# exit
```

IES Services Command Reference

Command Hierarchies

- Global Commands for 7210 SAS-D and 7210 SAS-E on page 509
- Interface Commands on page 509
- Routed VPLS Commands for 7210 SAS-D and 7210 SAS-K 2F2T1C on page 510
- Interface SAP Commands on page 511
- IES SAP Configuration - QoS and Filter Commands for 7210 SAS-D and 7210 SAS-E on page 512
- Interface IPv6 commands (applicable only to 7210 SAS-D and supported only for access-uplink SAPs) on page 513
- Show Commands on page 514

Global Commands for 7210 SAS-D and 7210 SAS-E

```
config
    — service
        — ies service-id [customer customer-id] [create] [vpn vpn-id]
        — no ies service-id
            — description description-string
            — no description
            — interface
            — no interface
            — service-name service-name
            — no service-name
            — [no] shutdown
```

Interface Commands

```
config
    — service
        — ies service-id [customer customer-id] [create] [vpn vpn-id]
            — [no] interface ip-int-name [create]
                — address {ip-address/mask | ip-address netmask}
                — no address
                — [no] arp-timeout
                — arp-timeout seconds
                — no arp-timeout
                — dhcp
                    — description description-string
```

```
                    — no description
                    — gi-address ip-address [src-ip-addr]
                    — no gi-address
                    — [no] option
                        — action {replace|drop|keep}
                        — no action
                        — [no] circuit-id [ascii-tuple|ifindex|sap-id|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
                    — no snoop
                    — snoop
                    — no server
                    — server server1 [server2...(upto 8 max)]
                    — [no] shutdown
                    — [no] trusted
                — delayed-enable seconds (supported only on 7210 SAS-D)
                — no delayed-enable
                — description description-string
                — no description
                — icmp
                    — mask-reply (supported only on 7210 SAS-D)
                    — no mask-reply
                    — redirects [number seconds]
                    — no redirects
                    — ttl-expired [number seconds]
                    — no ttl-expired
                    — unreachables [number seconds]
                    — no unreachables
                — ip-mtu octets (Supported only on 7210 SAS-D)
                — no ip-mtu
                — [no] loopback
                — [no] shutdown
                — [no] static-arp ip-address
```

Routed VPLS Commands for 7210 SAS-D and 7210 SAS-K 2F2T1C

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
            — interface ip-interface-name [create]
            — no interface-interface-name
                — vpls service-name
                — no vpls
                    — ingress
                        — v4-routed-override-filter ip-filter-id
                        — no v4-routed-override-filter
```

Interface SAP Commands

```
config
    — service
        — ies service-id [customer customer-id] [create]
            — [no] interface ip-int-name
                — [no] sap sap-id [create]
                    — accounting-policy acct-policy-id
                    — no accounting-policy
                    — collect-stats
                    — no collect-stats
                    — description description-string
                    — no description
                    — ingress
                        — meter-override
                        — no meter-override
                            — meter meter-id [create]
                            — no meter meter-id
                                — adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
                                — cbs size-in-kbytes (supported on 7210 SAS-E)
                                — no cbs
                                — cbs size [kbits|bytes|kbytes] (supported on 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C)
                                — no cbs
                                — mbs size-in-kbits (supported on 7210 SAS-E)
                                — no mbs
                                — mbs size [kbits|bytes|kbytes] (supported on 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C)
                                — no mbs
                                — no mode
                                — no mode
                                — rate cir cir-rate [pir pir-rate]
                    — statistics
                        — egress
                            — [no] forwarded-count
                    — [no] packets-forwarded-count
                        — ingress
                            — counter-mode {in-out-profile-count|forward-drop-count}
                            — [no] received-count
                    — [no] shutdown
```

IES SAP Configuration - QoS and Filter Commands for 7210 SAS-D and 7210 SAS-E

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id] [create]
            — [no] interface ip-int-name
                — [no] sap sap-id [create]
                    — egress
                        — aggregate-meter-rate rate-in-kbps [burst burst-in-kbits] [enable-stats]
                        — no aggregate-meter-rate
                        — filter ip ip-filter-id
                        — filter ipv6 ipv6 -filter-id
                        — no filter [ip ip-filter-id] [ipv6 ipv6 -filter-id]
                    — ingress
                        — aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
                        — no aggregate-meter-rate
                        — filter ip ip-filter-id
                        — filter [ipv6 ipv6-filter-id]
                        — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                        — qos policy-id
                        — no qos
```

Interface IPv6 commands (applicable only to 7210 SAS-D and supported only for access-uplink SAPs)

```
config
    — service
        — ies service-id [customer customer-id] [create]
            — [no] interface ip-int-name [create]
                — ipv6
                — no ipv6
                    — [no] address ipv6-address/prefix-length [eui-64] [preferred]
                    — icmp6
                        — [no] packet-too-big number seconds
                        — [no] param-problem number seconds
                        — [no] redirects number seconds
                        — [no] time-exceeded number seconds
                        — [no] unreachables number seconds
                    — [no] link-local-address ipv6-address [preferred]
                    — [no] local-proxy-nd
                    — [no] neighbor ipv6-address mac-address
                    — [no] proxy-nd-policy policy-name [policy-name...(upto 5 max)]
```

```
show
    — service
        — customer [customer-id] [site customer-site-name]
        — sap-using [sap sap-id]
        — sap-using interface [ip-address | ip-int-name]
        — sap-using [ingress | egress] filter filter-id
        — sap-using [ingress] qos-policy qos-policy-id
        — service-using [ies] [customer customer-id]
        — id service-id
            — all
            — arp [ip-address]|[mac ieee-address]|[sap sap-id]|[interface ip-int-name]
            — base
            — dhcp
                — statistics [sap sap-id] |[sdp sdp-id:vc-id]|[interface interface-name]
                — summary [interface interface-name | saps]
            — interface [ip-address | ip-int-name] [detail |summary]
```

IES Service Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>service>ies
config>service>ies>if

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Special Cases **IES —** The default administrative status of an IES service is down. While the service is down, all its associated virtual router interfaces will be operationally down. The administrative state of the service is not reflected in the administrative state of the virtual router interface.

For example if: 1) An IES service is operational and an associated interface is shut down.
2) The IES service is administratively shutdown and brought back up.
3) The interface shutdown will remain in administrative shutdown state.

A service is regarded as operational if one IP Interface is operational.

IES IP Interfaces — When the IP interface is shutdown, it enters the administratively and operationally down states. For a SAP bound to the IP interface, no packets are transmitted out the SAP and all packets received on the SAP will be dropped while incrementing the packet discard counter.

description

Syntax **description** *long description-string*
no description

Context config>service>ies

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

IES Global Commands

ies

Syntax **ies** *service-id* **customer** *customer-id* [**create**] [**vpn** *vpn-id*]
no ies *service-id*

Context config>service

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command creates or edits an IES service instance.

The **ies** command is used to create or maintain an (IES). If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

IP interfaces defined within the context of an IES service ID must have a SAP created.

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.

More than one IP interface may be created within a single IES service ID.

By default, no IES service instances exist until they are explicitly created.

The **no** form of this command deletes the IES service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces defined within the service ID have been shutdown and deleted.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR, 7450 ESS and 7710 SR on which this service is defined.

Values *service-id*: 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 assigned to the service.

Values 1 — 2147483647

service-name

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

Context config>service>epipe
config>service>ies
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 7210 SAS 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 o 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).

IES Interface IPv6 Commands

ipv6

Syntax [**no**] **ipv6**

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-D

This command enables the context to configure IPv6 for an IES interface.

address

Syntax **address** *ipv6-address/prefix-length* [**eui-64**]
no address *ipv6-address/prefix-length*

Context config>service>ies>if>ipv6

Description **Platforms Supported**: 7210 SAS-D

This command assigns an IPv6 address to the IES interface.

Parameters *ipv6-address/prefix-length* — Specify the IPv6 address on the interface.

Values

ipv6-address/prefix: 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
prefix-length	1 — 128

eui-64 — When the **eui-64** keyword is specified, a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from MAC address on Ethernet interfaces. For interfaces without a MAC address, for example ATM interfaces, the Base MAC address of the chassis is used.

icmp6

Syntax **icmp6**

Context config>service>ies>if>ipv6

Description **Platforms Supported**: 7210 SAS-D

This command configures ICMPv6 parameters for the IES interface.

packet-too-big

Syntax **packet-too-big** [*number seconds*]
no packet-too-big

Context config>service>ies>if>ipv6>icmp6

Description **Platforms Supported**: 7210 SAS-D.

This command specifies whether "packet-too-big" ICMPv6 messages should be sent. When enabled, ICMPv6 "packet-too-big" messages are generated by this interface.

The **no** form of the command disables the sending of ICMPv6 "packet-too-big" messages.

Default 100 10

Parameters *number* — Specifies the number of "packet-too-big" ICMPv6 messages to send in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of "packet-too-big" ICMPv6 messages issued.

Values 1 — 60

Default 10

param-problem

Syntax **param-problem** [*number seconds*]
no packet-too-big

Context config>service>ies>if>ipv6>icmp6

Description **Platforms Supported**: 7210 SAS-D

This command specifies whether "parameter-problem" ICMPv6 messages should be sent. When enabled', "parameter-problem" ICMPv6 messages are generated by this interface.

The **no** form of the command disables the sending of "parameter-problem" ICMPv6 messages.

Default 100 10

number — Specifies the number of "parameter-problem" ICMPv6 messages to send in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of "parameter-problem" ICMPv6 messages issued.

Values 1 — 60

Default 10

redirects

Syntax **redirects** [*number seconds*]
no redirects

Context config>service>ies>if>ipv6>icmp6

Description **Platforms Supported**: 7210 SAS-D

This command configures ICMPv6 redirect messages. When enabled, ICMPv6 redirects are generated when routes are not optimal on this router and another router on the same subnetwork has a better route in order to alert that node that a better route is available.

When disabled, ICMPv6 redirects are not generated.

Default 100 10

number — Specifies the number of version 6 redirects are to be issued in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of version 6 redirects issued.

Values 1 — 60

Default 10

time-exceeded

Syntax **time-exceeded** [*number seconds*]
no time-exceeded

Context config>service>ies>if>ipv6>icmp6

Description **Platforms Supported**: 7210 SAS-D

This command specifies whether "time-exceeded" ICMPv6 messages should be sent. When enabled, ICMPv6 "time-exceeded" messages are generated by this interface.

When disabled, ICMPv6 "time-exceeded" messages are not sent.

Default 100 10

number — Specifies the number of "time-exceeded" ICMPv6 messages are to be issued in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of "time-exceeded" ICMPv6 message to be issued.

Values 1 — 60

Default 10

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>ies>if>ipv6>icmp6

Description **Platforms Supported**: 7210 SAS-D

This command specifies that ICMPv6 host and network unreachable messages are generated by this interface.

When disabled, ICMPv6 host and network unreachable messages are not sent.

Default 100 10

number — Specifies the number of destination unreachable ICMPv6 messages are issued in the time frame specified by the *seconds* parameter.

Values 10 — 1000

Default 100

seconds — Specifies the time frame in seconds that is used to limit the number of destination unreachable ICMPv6 messages to be issued.

Values 1 — 60

Default 10

link-local-address

Syntax **link-local-address** *ipv6-address* [preferred]
no link-local-address

Context config>service>ies>if>ipv6

Description **Platforms Supported**: 7210 SAS-D

This command configures the IPv6 link local address.

local-proxy-nd

Syntax [**no**] **local-proxy-nd**

Context config>service>ies>if>ipv6

Description **Platforms Supported**: 7210 SAS-D

This command enables local proxy neighbor discovery on the interface.

The **no** form of the command disables local proxy neighbor discovery.

proxy-nd-policy

Syntax **proxy-nd-policy** *policy-name* [*policy-name*...(up to 5 max)]
no proxy-nd-policy

Context config>service>ies>if>ipv6

Description **Platforms Supported**: 7210 SAS-D

This command applies a proxy neighbor discovery policy for the interface.

Parameters *policy-name* — Specifies an existing neighbor discovery 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. The specified policy name(s) must already be defined.

neighbor

Syntax **neighbor** *ipv6-address mac-address*
no neighbor *ipv6-address*

Context config>service>ies>if>ipv6

Description **Platforms Supported**: 7210 SAS-D

This command configures IPv6-to-MAC address mapping on the IES interface.

Default none

Parameters *ipv6-address* — The IPv6 address of the interface for which to display information.

Values 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
prefix-length [1..128]

mac-address — Specifies the 48-bit MAC address for the IPv6-to-MAC address mapping 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.

IES Interface Commands

interface

Syntax **interface** *ip-int-name* [**create**]
no interface *ip-int-name*

Context config>service>ies

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

On 7210 SAS-K 2F2T1C, an IES interface must always be associated with RVPLS. It cannot be used "standalone".

This command creates a logical IP routing interface for an (IES). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within IES service IDs. The **interface** command can be executed in the context of an IES service ID. The IP interface created is associated with the service core network routing instance and default routing.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config service ies interface** (that is, the network core router instance). Interface names must not be in the dotted decimal notation of an IP address. 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. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All IES IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The **no** form of this command removes IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For IES services, the IP interface must be shutdown before the SAP on that interface may be removed.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

If *ip-int-name* already exists within the service ID, the context will be changed to maintain that IP interface. If *ip-int-name* already exists within another service ID, an error will occur and context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*} [**broadcast** [**all-ones** | **host-ones**}]
address *ip-address mask*
no address

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command assigns an IP address and IP subnet, 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 7210 SAS.

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 adminstratively up state and an address is assigned, it becomes operationally up.

ip-address — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

/ — The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "/" and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.

mask-length — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values 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.

broadcast - — Specifies the broadcast format.

Values all-ones, host-ones

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

The **no** form of this command restores **arp-timeout** to the default value.

Default 14400 seconds

Parameters *seconds* — The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 — 65535

Values

allow-directed-broadcasts

Syntax **[no] allow-directed-broadcasts**

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command enables the forwarding of directed broadcasts out of the IP interface. A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The allow-directed-broadcasts command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface.

When enabled, a frame destined to the local subnet on this IP interface will be sent as a subnet broadcast out this interface. Care should be exercised when allowing directed broadcasts as it is a well-known mechanism used for denial-of-service attacks.

When disabled, directed broadcast packets discarded at this egress IP interface will be counted in the normal discard counters for the egress SAP.

By default, directed broadcasts are not allowed and will be discarded at this egress IP interface.

The no form of this command disables the forwarding of directed broadcasts out of the IP interface.

Default no allow-directed-broadcasts — Directed broadcasts are dropped.

delayed-enable

Syntax **delayed-enable seconds [init-only]**
no delayed-enable

Context config>service>ies>if

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command delays making interface operational by the specified number of seconds. In environments with many subscribers, it can take time to synchronize the subscriber state between peers when the subscriber-interface is enabled (perhaps, after a reboot). To ensure that the state has time to be synchronized, the delayed-enable timer can be specified. The optional parameter init-only can be added to use this timer only after a reboot.

Default no delayed-enable

Parameters *seconds* — Specifies the number of seconds to delay before the interface is operational.

Values 1 — 1200

ip-mtu

Note: This command is supported only on 7210 SAS-D devices. The ip-mtu command for 7210 SAS-E devices can be enabled using the CLI "**config>port>ethernet>ip-mtu** *mtu-bytes* ".

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-D and 7210 SAS-K 2F2T1C

This command configures the maximum IP transmit unit (packet) for the interface.

The MTU that is advertized from the IES size is:

MINIMUM((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))

By default (for Ethernet network interface) if no ip-mtu is configured, the packet size is (1568 - 14) = 1554.

The **no** form of the command returns the default value.

Default no ip-mtu

Parameters *octets* — pecifies the number of octets in the IP-MTU.

Values 512 — 9000

loopback

Syntax [**no**] **loopback**

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated IES interface cannot be bound to a SAP.

Note that you can configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap** command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

Default none

static-arp

Syntax **static-arp** *ip-address ieee-address*
no static-arp *ip-address*

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of the command removes a static ARP entry.

Default None

Parameters *ip-address* — Specifies the IP address for the static ARP in IP address dotted decimal notation.

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.

vpls

Syntax vpls *service-name*

Context config>service
config>service>ies>if

Description **Platforms Supported**: 7210 SAS-D and 7210 SAS-K 2F2T1C

The vpls command, within the IP interface context, is used to bind the IP interface to the specified service name.

The system does not attempt to resolve the service name provided until the IP interface is placed into the administratively up state (no shutdown). Once the IP interface is administratively up, the system scans the available VPLS services that have the allow-ip-int-binding flag set for a VPLS service associated with the name. If the service name is bound to the service name when the IP interface is already in the administratively up state, the system will immediately attempt to resolve the given name.

If a VPLS service is found associated with the name and with the allow-ip-int-binding flag set, the IP interface will be attached to the VPLS service allowing routing to and from the service virtual ports once the IP interface is operational.

A VPLS service associated with the specified name that does not have the allow-ip-int-binding flag set or a non-VPLS service associated with the name will be ignored and will not be attached to the IP interface.

If the service name is applied to a VPLS service after the service name is bound to an IP interface and the VPLS service allow-ip-int-binding flag is set at the time the name is applied, the VPLS service is automatically resolved to the IP interface if the interface is administratively up or when the interface is placed in the administratively up state.

If the service name is applied to a VPLS service without the allow-ip-int-binding flag set, the system does not attempt to resolve the applied service name to an existing IP interface bound to the name. To rectify this condition, the flag must first be set and then the IP interface must enter or reenter the administratively up state.

While the specified service name may be assigned to only one service context in the system, it is possible to bind the same service name to more than one IP interface. If two or more IP interfaces are bound to the same service name, the first IP interface to enter the administratively up state (if currently administratively down) or to reenter the administratively up state (if currently administratively up) when a VPLS service is configured with the name and has the allow-ip-intbinding flag set will be attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

Once an IP interface is attached to a VPLS service, the name associated with the service cannot be removed or changed until the IP interface name binding is removed. Also, the allow-ip-int-binding flag cannot be removed until the attached IP interface is unbound from the service name. Unbinding the service name from the IP interface causes the IP interface to detach from the VPLS service context. The IP interface may then be bound to another service name or a SAP or SDP binding may be created for the interface using the sap or spoke-sdp commands on the interface.

Default none

Parameters *service-name* — The service-name parameter is required when using the IP interface vpls command and specifies the service name that the system will attempt to resolve to an allow-ip-int-binding enabled VPLS service associated with the name. The specified name is expressed as an ASCII string comprised of up to 32 characters. It does not need to already be associated with a service and the system does not check to ensure that multiple IP interfaces are not bound to the same name.

ingress

Syntax ingress

Context config>service>ies>if>vpls

Description **Platforms Supported**: 7210 SAS-D and 7210 SAS-K 2F2T1C

The ingress node in this context under the vpls binding is used to define the routed ip-filter-id optional filter overrides.

v4-routed-override-filter

Syntax v4-routed-override-filter *ip-filter-id*
no v4-routed-override-filter

Context config>service>ies>if>vpls>ingress

Description **Platforms Supported**: 7210 SAS-D and 7210 SAS-K 2F2T1C

The v4-routed-override-filter command is used to specify an IP filter ID that is applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IP filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and when it is not defined or it is removed, the IP routed packets uses the any existing ingress IP filter on the VPLS virtual port.

The no form of the command is used to remove the IP routed override filter from the ingress IP interface. When removed, the IP ingress routed packets within a VPLS service attached to the IP interface uses the IP ingress filter applied to the packets virtual port when defined.

Default none

Parameters *ip-filter-id* — Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the **configure>filter>ip-filter** context.

Values 1 — 65535

IES Interface ICMP Commands

icmp

Syntax **icmp**

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command enables the context to configure Internet Control Message Protocol (ICMP) parameters on an IES service

mask-reply

Syntax **[no] mask-reply**

Context config>service>ies>if>icmp

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance will reply to mask requests.

The no form of this command disables replies to ICMP mask requests on the router interface.

Default mask-reply — Reply to ICMP mask requests.

redirects

Syntax **redirects** [*number seconds*]
no redirects

Context config>service>ies>if>icmp

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.

When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

The **redirects** command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval. (*Default: redirects 100 10*)

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default **redirects 100 10** — Maximum of 100 redirect messages in 10 seconds

Parameters *number* — The maximum number of ICMP redirect messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP redirect messages that can be issued.

Values 1 — 60

ttl-expired

Syntax **ttl-expired** *number seconds*
no ttl-expired

Context config>service>ies>if>icmp

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default ttl-expired 100 10

Parameters *number* — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 — 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>ies>if>icmp

Description **Platforms Supported**: 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C

This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The **unreachables** command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *time* parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.

By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 10 per 60 second time interval.

The **no** form of this command disables the generation of icmp destination unreachable messages on the router interface.

Default **unreachables 100 10**

Parameters *number* — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 — 60

sap

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

Context config>service>ies>if

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C.

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 router. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access uplink port using the **configure port** *port number* **ethernet mode access uplink** 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.

Default No SAPs are defined.

Special Cases **IES —** A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Attempts to create a second SAP on an IP interface will fail and generate an error; the original SAP will not be affected.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 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 1/1/1 specifies port 1 on MDA 1 in slot 1.

The *port-id* must reference a valid port type. The port must be configured as an uplink access port.

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

IES Filter Commands

filter

Syntax **filter ipv6** *ipv6-filter-id*
filter ip *ip-filter-id*
no filter

Context config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress

Description **Platforms Supported:** 7210 SAS-D (both ingress and egress context), 7210 SAS-E (only egress context), and 7210 SAS-K 2F2T1C (both ingress and egress context).

This command associates a filter policy with an ingress or egress Service Access Point (SAP). Filter policies control the forwarding and dropping of packets based on the matching criteria.

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

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to the match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP. The filter ID itself is not removed from the system.

Special Cases **IES —** Only IP filters are supported on an IES IP interface, and the filters only apply to routed traffic.

Parameters **ip —** Keyword indicating the filter policy is an IP filter.

ipv6 — Keyword indicating the filter policy is an IPv6 filter.

ip-filter-id — Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the **configure>filter>ip-filter** context.

Values 1 — 65535

ipv6 ipv6-filter-id — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters (applicable only for 7210 SAS-D).

Values 1 — 65535

egress

Syntax **egress**

Context config>service>ies>if>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C.

This command enables the context to apply egress policies.

ingress

Syntax **ingress**

Context config>service>ies>if>sap

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-E, and 7210 SAS-K 2F2T1C.

This command enables the context to apply ingress policies

statistics

Syntax **statistics**

Context config>service>>ies>sap

Description **Platforms Supported**: 7210 SAS-D and 7210 SAS-E

This command enables the context to configure the counters associated with SAP ingress and egress.

egress

Syntax **egress**

Context config>service>>ies>sap>statistics

Description **Platforms Supported**: 7210 SAS-D and 7210 SAS-E

This command enables the context to configure the egress SAP statistics counter and set the mode of the counter.

This counter counts the number of packets forwarded through the SAP.

ingress

Syntax **ingress**

Description config>service>>ies>sap>statistics

Description **Platforms Supported**: 7210 SAS-D and 7210 SAS-E

This command enables the context to configure the ingress SAP statistics counter.

By default, SAP ingress counters are associated with a SAP and cannot be disabled.

The IES service supports a counter that counts the total packets or octets received on the SAP.

packets-forwarded-count

Syntax [**no**] **packets-forwarded-count**

Context config>service>>ies>sap>statistics>egress

Description **Platforms Supported**: 7210 SAS-E

This command associates a counter with the SAP. The counter counts the number of packets forwarded through the SAP.

The **no** form of this command disables the packet count.

counter-mode

Syntax **counter-mode** {in-out-profile-count|forward-drop-count}

Context config>service>ies>sap>statistics>ingress

Description **Platforms Supported**: 7210 SAS-D

This command sets the mode of ingress counters associated with the SAP to either octets or packets. On IES SAPs, collect stats cannot be enabled so the mode of the counter can be changed without any reference. Changing the mode of the counter results in loss of previously collected counts and resets the counter. The **no** form of this command is not supported.

Default in-out-profile-count

Parameters **in-out-profile-count —** If the counter mode is specified as "in-out-profile-count", one counter counts the total in-profile packets and octets received on ingress of a SAP and another counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

forward-drop-count — If the counter mode is specified as "forward-drop-count", one counter counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

Internet Enhanced Service on 7210 SAS-K 2F4T6C

In This Chapter

This chapter provides information about Internet Enhanced Services support on 7210 SAS-K 2F4T6C.

Topics in this chapter include:

- IES Service Overview on page 540
- IES Features on page 541
  - IP Interfaces on page 541
- SAPs on page 542
  - Encapsulations on page 542
  - CPE Connectivity Check on page 542
  - QoS Policies on page 543
  - CPU QoS for IES interfaces on page 543
  - Filter Policies on page 543
- Configuring an IES Service with CLI on page 501
  - Basic Configuration on page 502
  - Common Configuration Tasks on page 503
  - Service Management Tasks on page 507

IES Service Overview

Internet Enhanced Service (IES) is a routed connectivity service where the subscriber communicates with an IP router interface to send and receive Internet traffic. An IES has one or more logical IP routing interfaces each with a SAP which acts as the access point to the subscriber's network.

IES allows IP interfaces to participate in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber be unique between other provider addressing schemes and potentially the entire Internet. While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be reserved for service IP provisioning, and be administered by a separate, but subordinate address authority.

IP interfaces defined within the context of an IES service must have a SAP associated as the access point to the subscriber network. Multiple IES services are created to segregate subscriber owned IP interfaces.

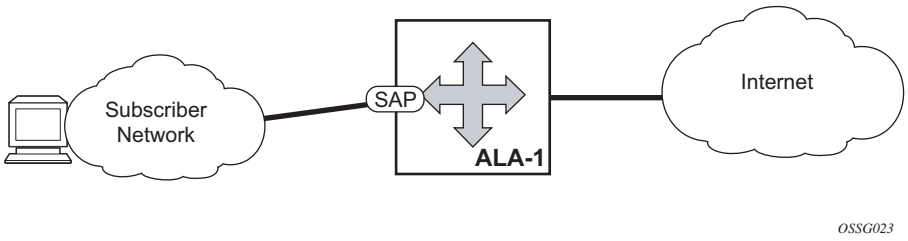

Figure 67: Internet Enhanced Service

The IES service provides in-band management connectivity. Other features include:

- Multiple IES services are created to separate IP interfaces.
- More than one IES service can be created for a single customer ID.
- More than one IP interface can be created within a single IES service ID. All IP interfaces created within an IES service ID belong to the same customer.

IES Features

This section describes various general service features and any special capabilities or considerations as they relate to IES services.

IP Interfaces

IES customer IP interfaces can be configured with most of the options found on the core IP interfaces. The advanced configuration options supported are:

- ICMP Options

In network mode, configuration options found on core IP interfaces not supported on IES IP interfaces are:

- NTP broadcast receipt

SAPs

See the following sections:

- Encapsulations on page 542
- CPE Connectivity Check on page 542
- QoS Policies on page 543
- CPU QoS for IES interfaces on page 543
- Filter Policies on page 543

Encapsulations

The following SAP encapsulation is supported on IES services:

- Ethernet null
- Ethernet dot1q
- Ethernet QinQ
- Ethernet QinQ (access-uplink QinQ SAP)

CPE Connectivity Check

Static routes are used within many IES services. Unlike dynamic routing protocols, there is no way to change the state of routes based on availability information for the associated CPE. CPE connectivity check adds flexibility so that unavailable destinations will be removed from the service provider's routing tables dynamically and minimize wasted bandwidth.

The availability of the far-end static route is monitored through periodic polling. The polling period is configured. If the poll fails a specified number of sequential polls, the static route is marked as inactive.

An ICMP ping mechanism is used to test the connectivity. If the connectivity check fails and the static route is deactivated, the router will continue to send polls and re-activate any routes that are restored.

QoS Policies

When applied to 7210 SAS IES services, service ingress QoS policies only create the unicast queues defined in the policy. The multi-point queues are not created on the service. With IES services, service egress QoS policies function as with other services where the class-based queues are created as defined in the policy.

Note that both MAC and IPv4 criteria can be used in the QoS policies for traffic classification in an IES.

CPU QoS for IES interfaces

Traffic bound to CPU received on IES access interfaces are policed/rate-limited and queued into CPU queues. The software allocates a policer per IP application or a set of IP applications, for rate-limiting CPU bound IP traffic from all IES SAPs. The policers CIR/PIR values are set to appropriate values based on feature scaling and these values are not user configurable. The software allocates a set of queues for CPU bound IP traffic from all IES SAPs. The queues are either shared by a set of IP applications or in some cases allocated to an IP application. The queues are shaped to appropriate rate based on feature scaling. The shaper rate is not user configurable.

NOTE: The instance of queues and policers used for traffic received on network port IP interfaces is different for traffic received from access port IP interfaces. Additionally the network CPU queues are accorded higher priority than the access CPU queues. This is done to provide better security and mitigate the risk of access traffic affecting network side.

Filter Policies

IP filter policies with IPv4 match entries can be applied to IES services.

Configuring an IES Service with CLI

This section provides information to configure IES services using the command line interface.

Topics in this section include:

- Basic Configuration on page 545
- Common Configuration Tasks on page 545
- Service Management Tasks on page 548

Basic Configuration

The most basic IES service configuration has the following entities:

- Customer ID (refer to Configuring Customers on page 79)
- An interface to create and maintain IP routing interfaces within IES service ID.
- A SAP on the interface specifying the access port and encapsulation values.

The following example displays a sample configuration of an IES service on ALA-50.

```
*A:ALA-50>config>service# info
----------------------------------------------
     ies 1000 customer 50 create
            description "to internet"
            interface "to-web" create
                address 10.1.1.1/24
                sap 1/1/10:100.* create
                exit
            exit
            no shutdown
----------------------------------------------
*A:ALA-50>config>service#
```

Common Configuration Tasks

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

1. Associate an IES service with a customer ID.
2. Associate customer ID with the service.

3. Assign an IP address.
4. Create an interface.
5. Define SAP parameters on the interface
6. Select node(s) and port(s).
7. Optional - select filter policies (configured in the config>filter context).
8. Enable service.

Configuring IES Components

Use the CLI syntax to configure the following entities:

- Configuring an IES Service on page 546
- Configuring IES Interface Parameters on page 546
- Configuring IES SAP Parameters on page 547

Configuring an IES Service

Use the following CLI syntax to create an IES service:

The following example displays a basic IES service configuration.

```
A:ALA-48>config>service#
-----------------------------------------------
...
ies 1001 customer 1730 create
               description "to-internet"
               no shutdown
exit
-----------------------------------------------
A:ALA-48>config>service#
```

Configuring IES Interface Parameters

The following is an example of configuring some IP interface parameters under IES service.

```
*A:K-SASK12>config>service>ies>if# info detail
----------------------------------------------
                no description
                no enable-ingress-stats
                no enable-mac-accounting
                no tcp-mss
                cpu-protection 254
                no address
                no mac
                arp-timeout 14400
                arp-retry-timer 50
                no arp-limit
                no allow-directed-broadcasts
                icmp
                    mask-reply
                    redirects 100 10
                    unreachables 100 10
                    ttl-expired 100 10
                exit
                dhcp
                    shutdown
                    no description
                    no option
                    no server
                    no trusted
                    no relay-proxy
                    no gi-address
                    no relay-plain-bootp
                exit
                no authentication-policy
                no ip-mtu
                no delayed-enable
                no multicast-network-domain
                no bfd
                no local-dhcp-server
                no proxy-arp-policy
                no local-proxy-arp
                no remote-proxy-arp
                no ptp-hw-assist
                no qos-route-lookup
                load-balancing
                    no teid-load-balancing
                    no egr-ip-load-balancing
                    no spi-load-balancing
                exit
                no vas-if-type
                no shutdown
----------------------------------------------
*A:K-SASK12>config>service>ies>if#
```

Configuring IES SAP Parameters

A SAP is a combination of a port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique within a router.

When configuring IES access SAP parameters, a default QoS policy is applied to each SAP ingress. Additional QoS policies must be configured in the config>qos context. Filter policies are configured in the config>filter context and must be explicitly applied to a SAP. There are no default filter policies.

This example displays an IES SAP configuration.

```
-----------------------------------------------
*A:ALA-A>config>service>ies>if# info
-----------------------------------------------
address 10.10.36.2/24
               sap 1/1/3:100 create
                       ingress
                       qos 101
                       exit
               exit
-----------------------------------------------
*A:ALA-A>config>service>ies>if#
```

Service Management Tasks

This section discusses the following service management tasks:

- Modifying IES Service Parameters on page 549
- Deleting an IES Service on page 549
- Disabling an IES Service on page 550
- Re-Enabling an IES Service on page 550

Modifying IES Service Parameters

Existing IES service parameters in the CLI or NMS can be modified, added, removed, enabled or disabled. The changes are applied immediately to all services when the charges are applied.

To display a list of customer IDs, use the show service customer command.

Enter the parameter(s) (such as description SAP information) and then enter the new information.

The following displays the modified service:

```
*A:ALA-A>config>service>ies# info
----------------------------------------------
     ies 1000 customer 50 create
           description "This is a new description"
           interface "to-web" create
               address 10.1.1.1/24
               mac 00:dc:98:1d:00:00
               sap 1/1/5:0.* create
               exit
           exit
           no shutdown
     exit
----------------------------------------------

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

Deleting an IES Service

An IES service cannot be deleted until SAPs and interfaces are shut down and deleted and the service is shutdown on the service level.

Use the following CLI syntax to delete an IES service:

CLI Syntax: `config>service#`

```
[no] ies service-id
shutdown
[no] interface ip-int-name
shutdown
[no] sap sap-id
shutdown
```

Disabling an IES Service

An IES service can be shut down without deleting the service parameters.

CLI Syntax: `config>service> ies service-id`

`shutdown`

Re-Enabling an IES Service

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

CLI Syntax: `config>service> ies service-id`

`[no] shutdown`

Example: `config>service# ies 2000`

```
config>service>ies# no shutdown
config>service>ies# exit
```

IES Services Command Reference

Command Hierarchies

Global Commands 7210 SAS-K 2F4T6C

```
config
    — service
        — ies service-id [customer customer-id] [create] [vpn vpn-id]
        — no ies service-id
            — description description-string
            — no description
            — interface
            — no interface
            — service-name service-name
            — no service-name
            — [no] shutdown
```

Interface Commands

```
config
    — service
        — ies service-id [customer customer-id] [create] [vpn vpn-id]
            — [no] interface ip-int-name [create]
                — address {ip-address/mask | ip-address netmask}
                — no address
                — [no] delayed-enable
                — arp-timeout seconds
                — no arp-timeout
                — dhcp
                    — description description-string
                    — no description
                    — gi-address ip-address [src-ip-addr]
                    — no gi-address
                    — [no] option
                        — action {replace|drop|keep}
                        — no action
                        — [no] circuit-id [ascii-tuple|ifindex|sap-id|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
```

```
                    — no relay-plain-bootp
                    — relay-plain-bootp
                    — no server
                    — server server1 [server2...(upto 8 max)]
                    — [no] shutdown
                    — [no] trusted
            — description description-string
            — no description
            — icmp
                    — mask-reply
                    — no mask-reply
                    — redirects [number seconds]
                    — no redirects
                    — ttl-expired [number seconds]
                    — no ttl-expired
                    — ttl-expired [number seconds]
                    — no unreachables
                    — unreachables octets
            — no ip-mtu
            — [no] loopback
            — [no] local-proxy-arp
            — [no] local-dhcp-server local-server-name
            — [no] shutdown
            — [no] static-arp ip-address [ieee-address]
```

Routed VPLS Commands

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
            — interface ip-interface-name [create]
            — no interface ip-interface-name
                — vpls service-name
                — no vpls
                    — ingress
                        — v4-routed-override-filter ip-filter-id
                        — no v4-routed-override-filter
```

IES SAP Configuration - QoS and Filter Commands for 7210 SAS-K 2F4T6C

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id] [create]
            — [no] interface ip-int-name
                — [no] sap sap-id [create]
                    — egress
                        — agg-shaper-rate agg-rate
                        — no agg-shaper-rate
                        — filter ip ip-filter-id
                        — no filter [ip ip-filter-id]
                        — qos policy-id
                        — no qos
                    — ingress
                        — agg-shaper-rate agg-rate
                        — no agg-shaper-rate
                        — aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
                        — no aggregate-meter-rate
                        — filter ip ip-filter-id
                        — no filter [ip ip-filter-id]
                        — qos policy-id
                        — no qos
```

```
show
  — service
      — customer [customer-id] [site customer-site-name]
      — sap-using [sap sap-id]
      — sap-using interface [ip-address | ip-int-name]
      — sap-using [ingress | egress] filter filter-id
      — sap-using [ingress] qos-policy qos-policy-id
      — service-using [ies] [customer customer-id]
      — id service-id
          — all
          — arp [ip-address]|[mac ieee-address]|[sap sap-id]|[interface ip-int-name]
          — base
          — dhcp
              — statistics [sap sap-id] |[sdp sdp-id:vc-id]|[interface interface-name]
              — summary [interface interface-name | saps]
          — interface [ip-address | ip-int-name] [detail |summary]
```

IES Service Configuration Commands for 7210 SAS-K 2F4T6C

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>service>ies
config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Special Cases **IES —** The default administrative status of an IES service is down. While the service is down, all its associated virtual router interfaces will be operationally down. The administrative state of the service is not reflected in the administrative state of the virtual router interface.

For example if: 1) An IES service is operational and an associated interface is shut down.
2) The IES service is administratively shutdown and brought back up.
3) The interface shutdown will remain in administrative shutdown state.

A service is regarded as operational if one IP Interface is operational.

IES IP Interfaces — When the IP interface is shutdown, it enters the administratively and operationally down states. For a SAP bound to the IP interface, no packets are transmitted out the SAP and all packets received on the SAP will be dropped while incrementing the packet discard counter.

description

Syntax **description** *long description-string*
no description

Context config>service>ies

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

IES Global Commands

ies

Syntax **ies** *service-id* **customer** *customer-id* [**create**] [**vpn** *vpn-id*]
no ies *service-id*

Context config>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command creates or edits an IES service instance.

The **ies** command is used to create or maintain an (IES). If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

IP interfaces defined within the context of an IES service ID must have a SAP created.

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.

More than one IP interface may be created within a single IES service ID.

By default, no IES service instances exist until they are explicitly created.

The **no** form of this command deletes the IES service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces defined within the service ID have been shutdown and deleted.

Parameters *service-id* — The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR, 7450 ESS and 7710 SR on which this service is defined.

Values *service-id*: 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 assigned to the service.

Values 1 — 2147483647

service-name

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

Context config>service>ies

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

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 7210 SAS 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 o 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).

IES Interface Commands

interface

Syntax **interface** *ip-int-name* [**create**]
no interface *ip-int-name*

Context config>service>ies

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command creates a logical IP routing interface for an (IES). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within IES service IDs. The **interface** command can be executed in the context of an IES service ID. The IP interface created is associated with the service core network routing instance and default routing.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config service ies interface** (that is, the network core router instance). Interface names must not be in the dotted decimal notation of an IP address. 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. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All IES IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The **no** form of this command removes IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For IES services, the IP interface must be shutdown before the SAP on that interface may be removed.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

If *ip-int-name* already exists within the service ID, the context will be changed to maintain that IP interface. If *ip-int-name* already exists within another service ID, an error will occur and context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*}
address *ip-address mask*
no address

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command assigns an IP address and IP subnet, 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 7210 SAS.

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 adminstratively up state and an address is assigned, it becomes operationally up.

ip-address — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

/ — The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "/" and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.

mask-length — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values 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.

broadcast - — Specifies the broadcast format.

Values all-ones, host-ones

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

The **no** form of this command restores **arp-timeout** to the default value.

Default 14400 seconds

Parameters *seconds* — The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 — 65535

Values

delayed-enable

Syntax **delayed-enable seconds [init-only]**
no delayed-enable

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command delays making interface operational by the specified number of seconds. In environments with many subscribers, it can take time to synchronize the subscriber state between peers when the subscriber-interface is enabled (perhaps, after a reboot). To ensure that the state has time to be synchronized, the delayed-enable timer can be specified. The optional parameter init-only can be added to use this timer only after a reboot.

Default no delayed-enable

Parameters *seconds* — Specifies the number of seconds to delay before the interface is operational.

Values 1 — 1200

loopback

Syntax **[no] loopback**

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated IES interface cannot be bound to a SAP.

Note that you can configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap** command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

Default none

static-arp

Syntax **static-arp** *ip-address ieee-mac-address*
no static-arp *ip-address*

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of the command removes a static ARP entry.

Default None

Parameters *ip-address* — Specifies the IP address for the static ARP in IP address dotted decimal notation.

ieee-mac-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

vpls

Syntax vpls *service-name*

Context config>service
config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

The vpls command, within the IP interface context, is used to bind the IP interface to the specified service name.

The system does not attempt to resolve the service name provided until the IP interface is placed into the administratively up state (no shutdown). Once the IP interface is administratively up, the system scans the available VPLS services that have the allow-ip-int-binding flag set for a VPLS service associated with the name. If the service name is bound to the service name when the IP interface is already in the administratively up state, the system will immediately attempt to resolve the given name.

If a VPLS service is found associated with the name and with the allow-ip-int-binding flag set, the IP interface will be attached to the VPLS service allowing routing to and from the service virtual ports once the IP interface is operational.

A VPLS service associated with the specified name that does not have the allow-ip-int-binding flag set or a non-VPLS service associated with the name will be ignored and will not be attached to the IP interface.

If the service name is applied to a VPLS service after the service name is bound to an IP interface and the VPLS service allow-ip-int-binding flag is set at the time the name is applied, the VPLS service is automatically resolved to the IP interface if the interface is administratively up or when the interface is placed in the administratively up state.

If the service name is applied to a VPLS service without the allow-ip-int-binding flag set, the system does not attempt to resolve the applied service name to an existing IP interface bound to the name. To rectify this condition, the flag must first be set and then the IP interface must enter or reenter the administratively up state.

While the specified service name may be assigned to only one service context in the system, it is possible to bind the same service name to more than one IP interface. If two or more IP interfaces are bound to the same service name, the first IP interface to enter the administratively up state (if currently administratively down) or to reenter the administratively up state (if currently administratively up) when a VPLS service is configured with the name and has the allow-ip-intbinding flag set will be attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

Once an IP interface is attached to a VPLS service, the name associated with the service cannot be removed or changed until the IP interface name binding is removed. Also, the allow-ip-int-binding flag cannot be removed until the attached IP interface is unbound from the service name. Unbinding the service name from the IP interface causes the IP interface to detach from the VPLS service context. The IP interface may then be bound to another service name or a SAP or SDP binding may be created for the interface using the sap or spoke-sdp commands on the interface.

Default none

Parameters *service-name* — The service-name parameter is required when using the IP interface vpls command and specifies the service name that the system will attempt to resolve to an allow-ip-int-binding enabled VPLS service associated with the name. The specified name is expressed as an ASCII string comprised of up to 32 characters. It does not need to already be associated with a service and the system does not check to ensure that multiple IP interfaces are not bound to the same name.

ingress

Syntax ingress

Context config>service>ies>if>vpls

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

The ingress node in this context under the vpls binding is used to define the routed ip-filter-id optional filter overrides.

v4-routed-override-filter

Syntax v4-routed-override-filter *ip-filter-id*
no v4-routed-override-filter

Context config>service>ies>if>vpls>ingress

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

The v4-routed-override-filter command is used to specify an IP filter ID that is applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IP filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and when it is not defined or it is removed, the IP routed packets uses the any existing ingress IP filter on the VPLS virtual port.

The no form of the command is used to remove the IP routed override filter from the ingress IP interface. When removed, the IP ingress routed packets within a VPLS service attached to the IP interface uses the IP ingress filter applied to the packets virtual port when defined.

Default none

Parameters *ip-filter-id* — Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the **configure>filter>ip-filter** context.

Values 1 — 65535

IES Interface ICMP Commands

icmp

Syntax **icmp**

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command enables the context to configure Internet Control Message Protocol (ICMP) parameters on an IES service.

mask-reply

Syntax **[no] mask-reply**

Context config>service>ies>if>icmp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance will reply to mask requests.

The no form of this command disables replies to ICMP mask requests on the router interface.

Default mask-reply — Reply to ICMP mask requests.

redirects

Syntax **redirects** [*number seconds*]
no redirects

Context config>service>ies>if>icmp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.

When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

The **redirects** command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval. (*Default: redirects 100 10*)

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default **redirects 100 10** — Maximum of 100 redirect messages in 10 seconds

Parameters *number* — The maximum number of ICMP redirect messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP redirect messages that can be issued.

Values 1 — 60

ttl-expired

Syntax **ttl-expired** *number seconds*
no ttl-expired

Context config>service>ies>if>icmp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default ttl-expired 100 10

Parameters *number* — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 — 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>ies>if>icmp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The **unreachables** command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *time* parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.

By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 10 per 60 second time interval.

The **no** form of this command disables the generation of icmp destination unreachable messages on the router interface.

Default **unreachables 100 10**

Parameters *number* — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 — 60

sap

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

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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 router. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port and access uplink port using the **configure port** *port number* **ethernet mode access uplink** 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.

Default No SAPs are defined.

Special Cases **IES —** A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Attempts to create a second SAP on an IP interface will fail and generate an error; the original SAP will not be affected.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 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 1/1/1 specifies port 1 on MDA 1 in slot 1.

The *port-id* must reference a valid port type. The port must be configured as an uplink access port.

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

IES Filter Commands

filter

Syntax **filter ip** *ip-filter-id*
no filter

Context config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command associates a filter policy with an ingress or egress Service Access Point (SAP). Filter policies control the forwarding and dropping of packets based on the matching criteria.

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

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to the match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP. The filter ID itself is not removed from the system.

Special Cases **IES —** Only IP filters are supported on an IES IP interface, and the filters only apply to routed traffic.

Parameters **ip —** Keyword indicating the filter policy is an IP filter.

ip-filter-id — Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the **configure>filter>ip-filter** context.

Values 1 — 65535

egress

Syntax **egress**

Context config>service>ies>if>sap

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to apply egress policies.

ingress

Syntax **ingress**

Context config>service>ies>if>sap

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to apply ingress policies

DHCP Server Commands

dhcp

Syntax **dhcp**

Context config>service>ies>sap

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure DHCP parameters.

option

Syntax [**no**] **option**

Context config>service>ies>sap>dhcp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.

The **no** form of this command returns the system to the default.

Default no option

action

Syntax **action** [*dhcp-action*] {***replace*** | ***drop*** | ***keep***}
no action

Context config>service>ies>sap>dhcp>option

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the Relay Agent Information Option (Option 82) processing.

The **no** form of this command returns the system to the default value.

Default The default is to keep the existing information intact.

Parameters *dhcp-action* — Specifies the DHCP option action.

replace — In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.

keep — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on towards the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

gi-address

Syntax **gi-address** *ip-address* [*src-ip-addr*]
no gi-address

Context config>service>ies>if>dhcp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the gateway interface address for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. The GI address is needed, when the router functions as a DHCP relay, to distinguish between the different subscriber interfaces and potentially between the group interfaces defined.

By default, the GI address used in the relayed DHCP packet is the primary IP address of a normal IES interface. Specifying the GI address allows the user to choose a secondary address.

Default no gi-address

Parameters *ip-address* — Specifies the host IP address to be used for DHCP relay packets.

src-ip-address — Specifies the source IP address to be used for DHCP relay packets.

circuit-id

Syntax **circuit-id** [**ascii-tuple** | **ifindex** | **sap-id** | **vlan-ascii-tuple**]
no circuit-id

Context config>service>ies>if>dhcp>option

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by "|".

In order to send a tuple in the circuit ID, the **action replace** command must be configured in the same context.

If disabled, the **circuit-id** sub-option of the DHCP packet will be left empty.

The **no** form of this command returns the system to the default.

Default circuit-id

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 "|".

ifindex — Specifies that the interface index will be used. (The If Index of a router interface can be displayed using the command `show>router>interface>detail`)

sap-id — Specifies that the SAP identifier will be 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-encapsulated ports only. Thus, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

remote-id

Syntax [**no**] **remote-id** [**mac** | **string** *string*]

Context config>service>ies>sap>dhcp>option

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies what information goes into the remote-id sub-option in the DHCP Relay packet.

If disabled, the **remote-id** sub-option of the DHCP packet will be left empty.

The **no** form of this command returns the system to the default.

Default no remote-id

Parameters **mac —** This keyword specifies the MAC address of the remote end is encoded in the sub-option.

string *string* **—** Specifies the remote-id.

vendor-specific-option

Syntax [**no**] **vendor-specific-option**

Context config>service>ies>if>dhcp>option

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the vendor specific sub-option of the DHCP relay packet.

client-mac-address

Syntax [**no**] **client-mac-address**

Context config>service>ies>sap>dhcp>option>vendor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the sending of the MAC address in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the vendor specific sub-option of the DHCP relay packet.

pool-name

Syntax [**no**] **pool-name**

Context config>service>ies>if>dhcp>option>vendor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command sends the pool name in the Nokia vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending.

sap-id

Syntax [**no**] **sap-id**

Context config>service>ies>sap>dhcp>option>vendor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the sending of the SAP ID in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the vendor specific sub-option of the DHCP relay packet.

service-id

Syntax [**no**] **service-id**

Context config>service>ies>sap>dhcp>option>vendor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the sending of the service ID in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the vendor specific sub-option of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

Context config>service>ies>sap>dhcp>option>vendor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the string in the vendor specific sub-option of the DHCP relay packet.

The **no** form of the command returns the default value.

Parameters *text* — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

system-id

Syntax [**no**] **system-id**

Context config>service>ies>sap>dhcp>option>vendor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies whether the system-id is encoded in the vendor specific sub-option of Option 82.

relay-plain-bootp

Syntax **relay-plain-bootp**
no relay-plain-bootp

Context config>service>ies>if>dhcp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the relaying of plain BOOTP packets.

The no form of the command disables the relaying of plain BOOTP packets.

server

Syntax **server** *server1* [*server2*...(up to 8 max)]

Context config>service>ies>if>dhcp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies a list of servers where requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list.

There can be a maximum of 8 DHCP servers configured.

Default no server

Parameters *server* — Specify the DHCP server IP address.

trusted

Syntax [**no**] **trusted**

Context config>service>ies>if>dhcp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables relaying of untrusted packets.

The **no** form of this command disables the relay.

Default not enabled

static-arp

Syntax **static-arp** *ip-address ieee-address*
no static-arp *ip-address*

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of the command removes a static ARP entry.

Default None

qos

Syntax **qos** *policy-id*
no qos

Context config>service>ies>sap>egress

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP) or IP interface.

QoS egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate ingress policies. The **qos** command only allows egress policies to be associated on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only 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.

By default, if no specific QoS policy is associated with the SAP for 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.

policy-id — The egress policy ID to associate with SAP on egress. The policy ID must already exist.

Values 1 — 65535

agg-shaper-rate

Syntax **agg-shaper-rate** *agg-rate*
no agg-shaper-rate

Context config>service>ies>sap>ingress

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command allows user to specify the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues. User can specify the CIR rate and the PIR rate. User must not oversubscribe the total bandwidth available for use by ingress queues.
The no form of the command disables the use of SAP aggregate rate shaper. In other words, the SAP can use up the maximum bandwidth available.

Default no agg-shaper-rate

Parameters ***agg-rate*** **—** Specifies the rate in kilobits per second.

Values 50 — 3000000 | max

Default max

agg-shaper-rate

Syntax **agg-shaper-rate** *agg-rate*
no agg-shaper-rate

Context config>service>ies>sap>egress

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.
This command allows user to specify the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues. User can specify the CIR rate and the PIR rate. User must not oversubscribe the total bandwidth available for use by ingress queues.
The no form of the command disables the use of SAP aggregate rate shaper. In other words, the SAP can use up the maximum bandwidth available.

Default no agg-shaper-rate

Parameters ***agg-rate*** **—** Specifies the rate in kilobits per second.

Values 50 — 1000000 | max

Default max

qos

Syntax **qos** *policy-id*
no qos

Context config>service>ies>sap>ingress

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP) or IP interface.

QoS ingress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate ingress policies. The **qos** command only allows ingress policies to be associated on SAP ingress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress 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.

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

policy-id — The ingress policy ID to associate with SAP on ingress. The policy ID must already exist.

Values 1 — 65535

local-proxy-arp

Syntax [**no**] **local-proxy-arp**

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus will become the forwarding point for all traffic between hosts in that subnet. When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

Default ies>if: no local-proxy-arp

local-dhcp-server

Syntax **local-dhcp-server *server-name* [create]**
no local-dhcp-server *server-name*

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command instantiates a local DHCP server. A local DHCP server can serve multiple interfaces but is limited to the routing context it was which it was created.

Default none

Parameters *server-name* — Specifies the name of local DHCP server.

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

loopback

Syntax [**no**] **loopback**

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated IES interface cannot be bound to a SAP.

You can configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap** command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

Default n/a

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>ies>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the IP maximum transmit unit (packet) for this interface.

Because this connects a Layer 2 to a Layer 3 service, this parameter can be adjusted under the IES interface.

The MTU that is advertised from the IES size is:

MINIMUM ((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))

By default (for Ethernet network interface) if no ip-mtu is configured it is (1568 - 14) = 1554.

The **no** form of the command returns the default value.

Default no ip-mtu

Virtual Private Network Service on 7210 SAS-K 2F4T6C

In This Chapter

This chapter provides information about the Virtual Private Routed Network (VPN) service and implementation notes on 7210 SAS-K 2F4T6C.

Topics in this chapter include:

- VPRN Service Overview on page 583
  - → Routing Prerequisites on page 584
  - → BGP Support on page 584
  - → Route Distinguisher on page 585
  - → Route Reflector on page 586
  - → CE to PE Route Exchange on page 586
  - → on page 588
- VPRN Features on page 589
  - → IP Interfaces on page 589
  - → SAPs on page 589
  - → CE to PE Routing Protocols on page 590
  - → Spoke SDP termination on page 591
  - → Service Label Mode of a VPRN on page 592

- Configuring a VPRN Service with CLI on page 595
  - → Basic Configuration on page 595
  - → Common Configuration Tasks on page 596
  - → Configuring VPRN Protocols - OSPF on page 603
  - → Service Management Tasks on page 604

VPRN Service Overview

RFC2547bis is an extension to the original RFC 2547, which details a method of distributing routing information and forwarding data to provide a Layer 3 Virtual Private Network (VPN) service to end customers.

Each Virtual Private Routed Network (VPRN) consists of a set of customer sites connected to one or more PE routers. Each associated PE router maintains a separate IP forwarding table for each VPRN. Additionally, the PE routers exchange the routing information configured or learned from all customer sites via MP-BGP peering. Each route exchanged via the MP-BGP protocol includes a Route Distinguisher (RD), which identifies the VPRN association.

The service provider uses BGP to exchange the routes of a VPN among the PE routers that are attached to that VPN. This is done in a way which ensures that routes from different VPNs remain distinct and separate, even if two VPNs have an overlapping address space. The PE routers distribute routes from other CE routers in that VPN to the CE routers in a VPN. Since the CE routers do not peer with each other there is no overlay visible to the VPN's routing algorithm.

When BGP distributes a VPN route, it also distributes an MPLS label for that route. The label distributed with a VPN route depends on the configured label-mode of the VPRN that is originating the route. On 7210 SAS routers, only one method of label allocation mode is supported: label per VRF. In other words, when 7210 allocates a label for a VPN route, it allocates a single label for all the VPN routes belonging to a single VRF. This does not restrict the label distribution method on the remote PE. The label distribution method configured on the remote PE must consider the scale of the PEs participating in the VPRN service.

Before a customer data packet travels across the service provider's backbone, it is encapsulated with the MPLS label that corresponds, in the customer's VPN, to the route which best matches the packet's destination address. The MPLS packet is further encapsulated with either another MPLS label header, so that it gets tunneled across the backbone to the proper PE router. Each route exchanged by the MP-BGP protocol includes a route distinguisher (RD), which identifies the VPRN association. Thus, the backbone core routers do not need to know the VPN routes. Figure 68 displays a VPRN network diagram example.

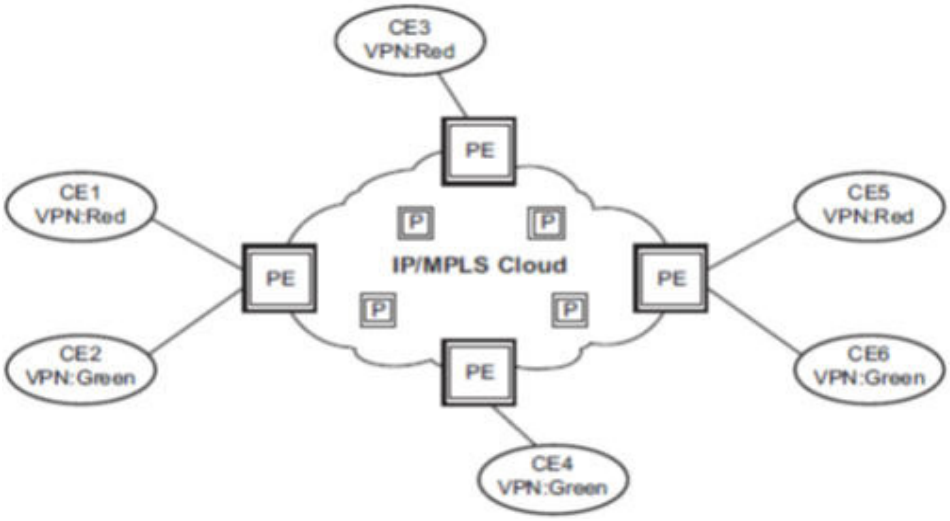

Figure 68: Virtual Private Routed Network

Routing Prerequisites

RFC2547bis requires the following features:

- Multi-protocol extensions
- Extended BGP community support
- BGP capability negotiation
- Parameters defined in RFC 2918

Tunneling protocol options are as follows:

- Label Distribution Protocol (LDP)
- MPLS RSVP-TE tunnels

BGP Support

BGP is used with BGP extensions mentioned in Routing Prerequisites on page 643 to distribute VPRN routing information across the service provider's network. BGP was initially designed to distribute IPv4 routing information. Therefore, multi-protocol extensions and the use of a VPN-IPv4 address were created to extend BGP's ability to carry overlapping routing information. A VPN-IPv4 address is a 12-byte value consisting of the 8-byte route distinguisher (RD) and the 4-byte IPv4 IP address prefix. The RD must be unique within the scope of the VPRN. This allows the IP address prefixes within different VRFs to overlap.

Route Distinguisher

The route distinguisher (RD) is an 8-byte value consisting of 2 major fields, the Type field and value field. The type field determines how the value field should be interpreted. The 7210 SAS implementation supports the three (3) type values as defined in the internet draft.

Figure 69: Route Distinguisher

The three Type values are:

- **Type 0: Value Field - Administrator subfield (2 bytes)**

  Assigned number subfield (4 bytes)

  The administrator field must contain an AS number (using private AS numbers is discouraged). The Assigned field contains a number assigned by the service provider.

- **Type 1: Value Field - Administrator subfield (4 bytes)**

  Assigned number subfield (2 bytes)

  The administrator field must contain an IP address (using private IP address space is discouraged). The Assigned field contains a number assigned by the service provider.

- **Type 2: Value Field - Administrator subfield (4 bytes)**

  Assigned number subfield (2 bytes)

  The administrator field must contain a 4-byte AS number (using private AS numbers is discouraged). The Assigned field contains a number assigned by the service provider.

Route Reflector

Per RFC2547bis the use of Route Reflectors is supported in the service provider core. Multiple sets of route reflectors can be used for different types of BGP routes, including VPN-IPv4, etc. 7210 SAS-K 2F4T6C can only be configured a route reflector client. It cannot be used as a route reflector "server" in the network.

CE to PE Route Exchange

Routing information between the Customer Edge (CE) and Provider Edge (PE) can be exchanged by the following methods:

- Static Routes (IPv4)
- E-BGP (IPv4)
- OSPFv2 (IPv4)

Each protocol provides controls to limit the number of routes learned from each CE router.

Route Redistribution

Routing information learned from the CE-to-PE routing protocols and configured static routes should be injected in the associated local VPN routing/forwarding (VRF). In the case of dynamic routing protocols, there may be protocol specific route policies that modify or reject certain routes before they are injected into the local VRF.

Route redistribution from the local VRF to CE-to-PE routing protocols is to be controlled via the route policies in each routing protocol instance, in the same manner that is used by the base router instance.

The advertisement or redistribution of routing information from the local VRF to or from the MPBGP instance is specified per VRF and is controlled by VRF route target associations or by VRF route policies.

VPN-IP routes imported into a VPRN, have the protocol type bgp-vpn to denote that it is an VPRN route. This can be used within the route policy match criteria.

CPE Connectivity Check

Static routes are used within many IES and VPRN services. Unlike dynamic routing protocols, there is no way to change the state of routes based on availability information for the associated CPE. CPE connectivity check adds flexibility so that unavailable destinations will be removed from the VPRN routing tables dynamically and minimize wasted bandwidth.

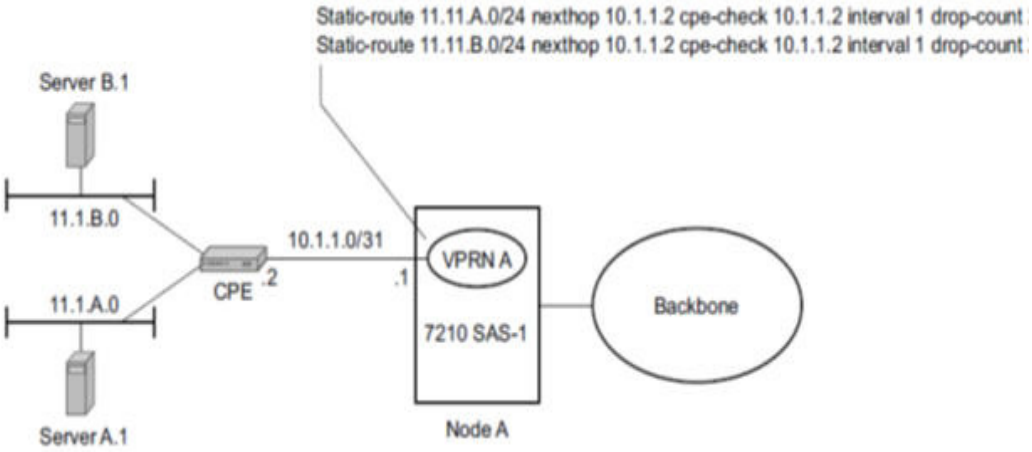

Figure 70: Directly connected to IP target

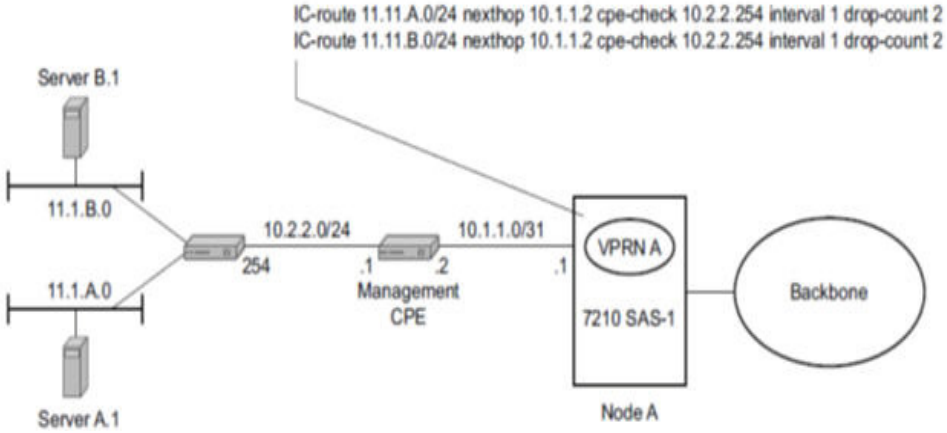

Figure 71: Multiple Hops to IP target

The availability of the far-end static route is monitored through periodic polling. The polling period is configured. If the poll fails a specified number of sequential polls, the static route is marked as inactive.

Either ICMP ping or unicast ARP mechanism can be used to test the connectivity. ICMP ping is preferred.

If the connectivity check fails and the static route is deactivated, the 7210 SAS router will continue to send polls and re-activate any routes that are restored.

VPRN Features

This section describes various VPRN features and any special capabilities or considerations as they relate to VPRN services.

- IP Interfaces on page 589
- SAPs on page 589
  - → Encapsulations on page 589
  - → QoS Policies for 7210 SAS-K 2F4T6C on page 590
  - → Filter Policies on page 590
  - → CPU QoS and DSCP Marking for Self-Generated Traffic on page 590
- CE to PE Routing Protocols on page 590
  - → PE to PE Tunneling Mechanisms on page 591
  - → Per VRF Route Limiting on page 591

IP Interfaces

VPRN customer IP interfaces can be configured with most of the same options found on the core IP interfaces.

The advanced configuration options supported are:

- ICMP Options

SAPs

Encapsulations

The following SAP encapsulations are supported on the 7210 SAS VPRN service:

- Ethernet null
- Ethernet dot1q
- QinQ
- LAG

QoS Policies for 7210 SAS-K 2F4T6C

When applied to a VPRN SAP, service ingress and service egress QoS policies creates the unicast and multicast queues defined in the policy (as multicast is not supported in VPRN service). Note that both Layer-2 (dot1p only) or Layer-3 criteria can be used in the QoS policies for traffic classification in an VPRN.

Filter Policies

Ingress and egress IPv4 filter policies can be applied to VPRN SAPs.

CPU QoS and DSCP Marking for Self-Generated Traffic

DSCP values, dot1p values and forwarding class for all applications is assigned by the system. On ingress, the system uses meters with default values to rate-limit all applications to system defined values. A queue and policer is used, one each for all access ports and for all network ports.

CE to PE Routing Protocols

The 7210 SAS VPRN supports the following PE to CE routing protocols:

- eBGP (IPv4)
- Static (IPv4)
- OSPFv2 (IPv4)

PE to PE Tunneling Mechanisms

The 7210 SAS supports multiple mechanisms to provide transport tunnels for the forwarding of traffic between PE routers within the 2547bis network. The 7210 SAS VPRN implementation supports the use of:

- RSVP-TE protocol to create tunnel LSP's between PE routers
- LDP protocol to create tunnel LSP's between PE routers

These transport tunnel mechanisms provide the flexibility of using dynamically created LSPs where the service tunnels are automatically bound (the “auto-bind” feature) and the ability to provide certain VPN services with their own transport tunnels by explicitly binding SDPs if desired. When the auto-bind is used, all services traverse the same LSPs and do not allow alternate tunneling mechanisms or the ability to craft sets of LSP's with bandwidth reservations for specific customers as is available with explicit SDPs for the service.

Per VRF Route Limiting

The 7210 SAS allows setting the maximum number of routes that can be accepted in the VRF for a VPRN service. There are options to specify a percentage threshold at which to generate an event that the VRF table is near full and an option to disable additional route learning when full or only generate an event.

Spoke SDP termination

Spoke-SDP termination into a Layer-3 service is not supported on 7210 SAS platforms.

Using OSPF in IP-VPNs

NOTE: OSPFv2 as PE-CE routing protocol is supported for IPv4 VPNs.

Using OSPF as a CE to PE routing protocol allows OSPF that is currently running as the IGP routing protocol to migrate to an IP-VPN backbone without changing the IGP routing protocol, introducing BGP as the CE-PE or relying on static routes for the distribution of routes into the service providers IP-VPN. The following features are supported:

- Advertisement/redistribution of BGP-VPN routes as summary (type 3) LSAs flooded to CE neighbors of the VPRN OSPF instance. This occurs if the OSPF route type (in the OSPF route type BGP extended community attribute carried with the VPN route) is not external (or NSSA) and the locally configured domain-id matches

the domain-id carried in the OSPF domain ID BGP extended community attribute carried with the VPN route.

- OSPF sham links. A sham link is a logical PE-to-PE unnumbered point-to-point interface that essentially rides over the PE-to-PE transport tunnel. A sham link can be associated with any area and can therefore appear as an intra-area link to CE routers attached to different PEs in the VPN.

Service Label Mode of a VPRN

The 7210 SAS allocates one unique (platform-wide) service label per VRF. All VPN-IP routes exported by the PE from a particular VPRN service with that configuration have the same service label. When the PE receives a terminating MPLS packet, the service label value determines the VRF to which the packet belongs. A lookup of the IP packet DA in the forwarding table of the selected VRF determines the next-hop interface.

Inter-AS VPRNs

Inter-AS IP-VPN services have been driven by the popularity of IP services and service provider expansion beyond the borders of a single Autonomous System (AS) or the requirement for IP VPN services to cross the AS boundaries of multiple providers. Three options for supporting inter-AS IP-VPNs are described in RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs).

NOTE: 7210 SAS-K 2F4T6C platforms support only option-A and option-C. It does not support option-B. It described below only for the sake of completeness.

The first option, referred to as Option-A (Figure 72), is considered inherent in any implementation. This method uses a back-to-back connection between separate VPRN instances in each AS. Thus, each VPRN instance views the inter-AS connection as an external interface to a remote VPRN customer site. The back-to-back VRF connections between the ASBR nodes require individual sub-interfaces, one per VRF.

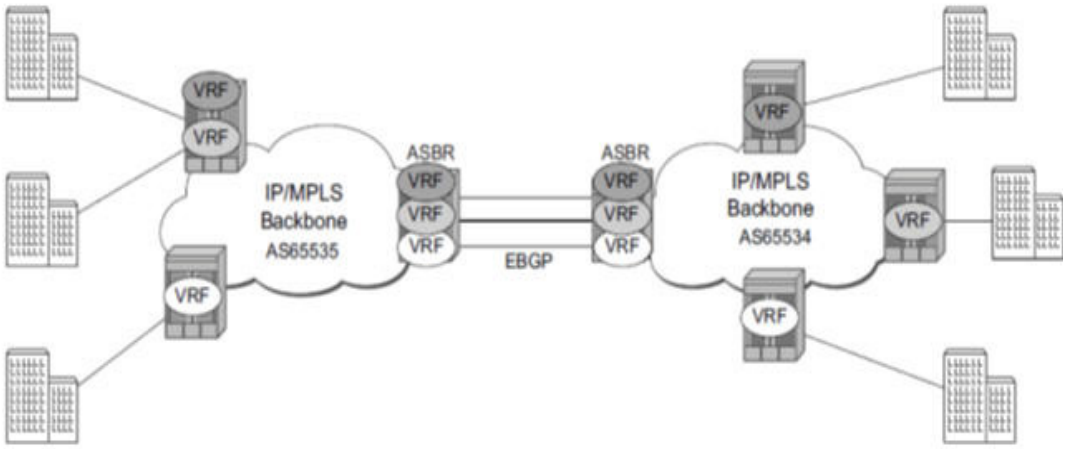

Figure 72: Inter-AS Option-A: VRF-to-VRF model

The second option, referred to as Option-B (Figure 73), relies heavily on the AS Boundary Routers (ASBRs) as the interface between the autonomous systems. This approach enhances the scalability of the eBGP VRF-to-VRF solution by eliminating the need for per-VPRN configuration on the ASBR(s). However, it requires that the ASBR(s) provide a control plan and forwarding plane connection between the autonomous systems. The ASBR(s) are connected to the PE nodes in its local autonomous system using iBGP either directly or through route reflectors. This means the ASBR(s) receive all the VPRN information and will forward these VPRN updates,

VPN-IPV4, to all its EBGP peers, ASBR(s), using itself as the next-hop. It also changes the label associated with the route. This means the ASBR(s) must maintain an associate mapping of labels received and labels issued for those routes. The peer ASBR(s) will in turn forward those updates to all local IBGP peers.

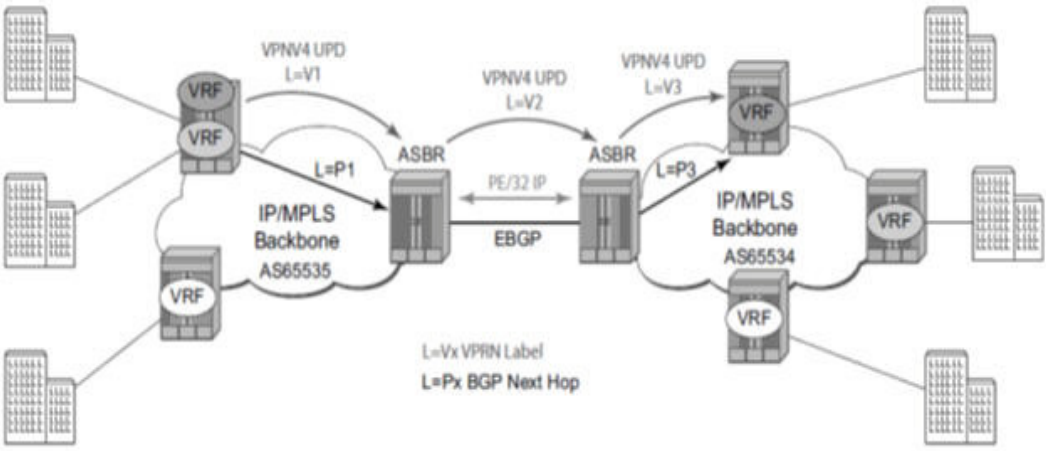

Figure 73: Inter-AS Option-B

This form of inter-AS VPRNs does not require instances of the VPRN to be created on the ASBR, as in option-A, thus there is less management overhead. This is also the most common form of Inter-AS VPRNs used between different service providers as all routes advertised between autonomous systems can be controlled by route policies on the ASBRs.

The third option, referred to as Option-C (Figure 74), allows for a higher scale of VPRNs across AS boundaries but also expands the trust model between ASNs. Thus, this model is typically used within a single company that may have multiple ASNs for various reasons.

This model differs from Option-B, in that in Option-B all direct knowledge of the remote AS is contained and limited to the ASBR. Thus, in option-B the ASBR performs all necessary mapping functions and the PE routers do not need perform any additional functions then in a non-Inter-AS VPRN.

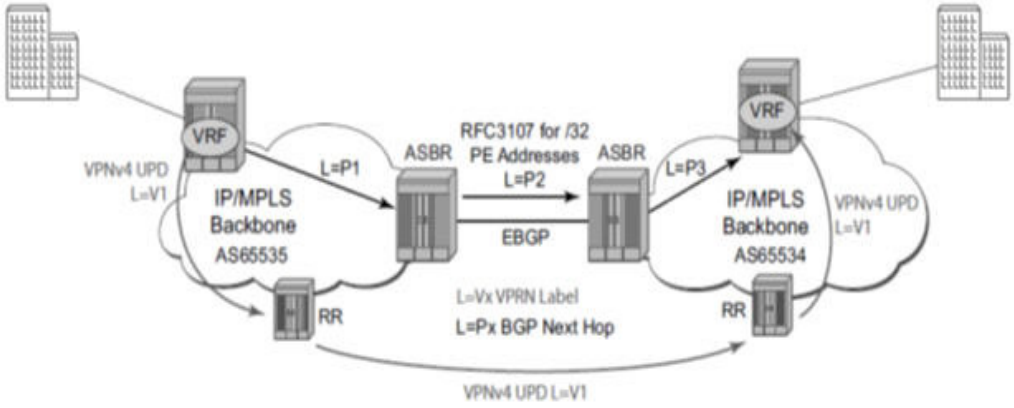

Figure 74: Option C Example

With Option-C, knowledge from the remote AS is distributed throughout the local AS. This distribution allows for higher scalability but also requires all PEs and ASBRs involved in the Inter-AS VPRNs to participate in the exchange of inter-AS routing information. In Option-C, the ASBRs distribute reachability information for remote PE's system IP addresses only. This is done between the ASBRs by exchanging MP-eBGP labeled routes, using RFC 3107, *Carrying Label Information in BGP-4*.

NOTE: 7210 SAS-K 2F4T6C does not play the role of option C ASBR.

Distribution of VPRN routing information is handled by either direct MP-BGP peering between PEs in the different ASNs or more likely by one or more route reflectors in ASN.

Configuring a VPRN Service with CLI

This section provides information to configure Virtual Private Routed Network (VPRN) services using the command line interface.

Topics in this section include:

- Basic Configuration on page 595
- Common Configuration Tasks on page 596
- Configuring VPRN Protocols - OSPF on page 603
- Service Management Tasks on page 604

Basic Configuration

The following fields require specific input (there are no defaults) to configure a basic VPRN service:

- Customer ID (refer to Configuring Customers on page 79
- Specify interface parameters

The following example displays a sample configuration of a VPRN service.

```
*A:K-SASK12>config>service>vprn$ info detail
----------------------------------------------
            shutdown
            no description
            dhcp
            exit
            dhcp6
            exit
            no router-id
            no maximum-routes
            no maximum-ipv6-routes
            snmp
                no access
            exit
            log
            exit
            no mc-maximum-routes
            no autonomous-system
            no route-distinguisher
            auto-bind-tunnel
                resolution-filter
                    no ldp
                    no rsvp
                exit
                resolution disabled
            exit
            no enable-bgp-vpn-backup
            no vrf-target
            no static-route-hold-down
            no ptp
```

```
            mvpn
                provider-tunnel
                exit
                no vrf-target
            exit
            radius-proxy
            exit
            radius-server
            exit
            no allow-export-bgp-vpn
            twamp-light
            exit
            network
                ingress
                    no filter
                    no qos
                exit
            exit
            no vsd-domain
-------------------------------------------
*A:K-SASK12>config>service>vprn$
```

Common Configuration Tasks

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

1. Associate a VPRN service with a customer ID.
2. Define an autonomous system (optional).
3. Define a route distinguisher (mandatory).
4. Define VRF route-target associations or VRF import/export policies.
5. Create an interface.
6. Define SAP parameters on the interface.
   - 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).
7. Define BGP parameters (optional).
   - BGP must be enabled in the config>router>bgp context.
8. Enable the service.

Configuring VPRN Components

This section provides VPRN configuration examples for the following entities:

- Creating a VPRN Service on page 597
- Configuring Global VPRN Parameters on page 598
- Configuring Router Interfaces on page 598
- Configuring VPRN Protocols - BGP on page 598

Creating a VPRN Service

Use the following CLI syntax to create a VRPN service. A route distinguisher must be defined for VPRN to be operationally active.

CLI Syntax:
```
config>service# vprn service-id [customer customer-id]
          route-distinguisher [ip-address:number1 | asn:number2]
          description description-string
          no shutdown
```

The following example displays a VPRN service configuration.

```
*A:ALA-1>config>service# info
----------------------------------------------
...
vprn 1 customer 1 create
route-distinguisher 10001:0
no shutdown
exit
...
----------------------------------------------
*A:ALA-1>config>service>vprn#
```

Configuring Global VPRN Parameters

Refer to VPRN Services Command Reference on page 681 for CLI syntax to configure VPRN parameters.

The following example displays a VPRN service with configured parameters.

```
*A:ALA-1>config>service# info
----------------------------------------------
...
vprn 1 customer 1 create
exit
            no enable-bgp-vpn-backup
            no vrf-target
            no static-route-hold-down
            no ptp
            mvpn
                provider-tunnel
                exit
                no vrf-target
            exit
exit
...
----------------------------------------------
*A:ALA-1>config>service#
```

Configuring Router Interfaces

Refer to the 7210 SAS OS Router Configuration Guide for command descriptions and syntax information to configure router interfaces.

The following example displays a router interface configurations:

```
*A:K-SASK12>config>router# info
----------------------------------------------
#------------------------------------------------
echo "IP Configuration"
#------------------------------------------------
        interface "system"
            no shutdown
        exit
#------------------------------------------------
#
```

Configuring VPRN Protocols - BGP

The autonomous system number and router ID configured in the VPRN context only applies to that particular service.

The minimal parameters that should be configured for a VPRN BGP instance are:

- Specify an autonomous system number for the router. See Configuring Global VPRN Parameters on page 598.
- Specify a router ID - Note that if a new or different router ID value is entered in the BGP context, then the new values takes precedence and overwrites the VPRN-level router ID.

See Configuring Global VPRN Parameters on page 598 on page 669.

- Specify a VPRN BGP peer group.
- Specify a VPRN BGP neighbor with which to peer.
- Specify a VPRN BGP peer-AS that is associated with the above peer.

VPRN BGP is administratively enabled upon creation. Minimally, to enable VPRN BGP in a VPRN instance, you must associate an autonomous system number and router ID for the VPRN service, create a peer group, neighbor, and associate a peer AS number. There are no default VPRN BGP groups or neighbors. Each VPRN BGP group and neighbor must be explicitly configured.

All parameters configured for VPRN BGP are applied to the group and are inherited by each peer, but a group parameter can be overridden on a specific basis. VPRN BGP command hierarchy consists of three levels:

- The global level
- The group level
- The neighbor level

For example:

CLI Syntax: `config>service>vprn>bgp# (global level)`

```
group (group level)

neighbor (neighbor level)
```

Note that the local-address must be explicitly configured if two systems have multiple BGP peer sessions between them for the session to be established.

For more information about the BGP protocol, refer to the 7210 SAS OS Router configuration Guide.

Configuring VPRN BGP Group and Neighbor Parameters

A group is a collection of related VPRN BGP peers. The group name should be a descriptive name for the group. Follow your group, name, and ID naming conventions for consistency and to help when troubleshooting faults.

All parameters configured for a peer group are applied to the group and are inherited by each peer (neighbor), but a group parameter can be overridden on a specific neighbor-level basis.

After a group name is created and options are configured, neighbors can be added within the same autonomous system to create IBGP connections and/or neighbors in different autonomous systems to create EBGP peers. All parameters configured for the peer group level are applied to each neighbor, but a group parameter can be overridden on a specific neighbor basis.

VPRN BGP CLI Syntax

Use the CLI syntax to configure VPRN BGP parameters (BGP Configuration Commands on page 691).

The following example displays a VPRN BGP configuration:

```
*A:K-SASK12>config>service>vprn>bgp$ info detail
----------------------------------------------
                no description
                no authentication-key
                no connect-retry
                no keepalive
                no hold-time
                no damping
                no local-preference
                no loop-detect
                no min-route-advertisement
                no aggregator-id-zero
                no preference
                no remove-private
                no multihop
                no med-out
                no cluster
                no disable-4byte-asn
                no import
                no export
                no local-as
                no path-mtu-discovery
                no router-id
                no disable-fast-external-failover
                no disable-communities
                no advertise-inactive
                no enable-peer-tracking
                no auth-keychain
                no rapid-withdrawal
                no split-horizon
                no backup-path
                best-path-selection
                    no always-compare-med
```

```
                    no deterministic-med
                    no as-path-ignore
                    no ignore-nh-metric
                    no ignore-router-id
                exit
                next-hop-resolution
                    no policy
                exit
                no peer-tracking-policy
                error-handling
                    no update-fault-tolerance
                exit
                no damp-peer-oscillations
                rib-management
                    ipv4
                        no leak-import
                    exit
                    ipv6
                        no leak-import
                    exit
                exit
                no shutdown
-----------------------------------------------
*A:K-SASK12>config>service>vprn>bgp$
```

Configuring a VPRN Interface

Interface names associate an IP address to the interface, and then associate the IP interface with a physical port. The logical interface can associate attributes like an IP address, port, Link Aggregation Group (LAG) or the system.

There are no default interfaces.

Note that you can configure a VPRN interface as a loop-back interface by issuing the loop-back command instead of the sap sap-id command. The loop-back flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loop-back interface.

The following example displays a VPRN interface configuration:

```
*A:K-SASK12>config>service>vprn>if$ info detail
-----------------------------------------------
                no description
                no enable-ingress-stats
                no enable-mac-accounting
                no tcp-mss
                cpu-protection 254
                no address
                no mac
                arp-timeout 14400
                arp-retry-timer 50
                no arp-limit
                no allow-directed-broadcasts
                icmp
```

```
                    mask-reply
                    redirects 100 10
                    unreachables 100 10
                    ttl-expired 100 10
                exit
                dhcp
                    shutdown
                    no description
                    no option
                    no server
                    no trusted
                    no relay-proxy
                    no gi-address
                    no relay-plain-bootp
                exit
                no authentication-policy
                no ip-mtu
                no delayed-enable
                no bfd
                no local-dhcp-server
                no proxy-arp-policy
                no local-proxy-arp
                no remote-proxy-arp
                no ptp-hw-assist
                no qos-route-lookup
                load-balancing
                    no teid-load-balancing
                    no egr-ip-load-balancing
                    no spi-load-balancing
                exit
                no vas-if-type
                no shutdown
-------------------------------------------
*A:K-SASK12>config>service>vprn>if$
```

Configuring a VPRN Interface SAP

A SAP is a combination of a port and encapsulation parameters which identifies the service access point on the interface and within the 7210 SAS. Each SAP must be unique within a router. A SAP cannot be defined if the interface loop-back command is enabled.

When configuring VPRN interface SAP parameters, a default QoS policy is applied to each ingress and egress SAP. Additional QoS policies and scheduler policies must be configured in the config>qos context. Filter policies are configured in the config>filter context and must be explicitly applied to a SAP. There are no default filter policies.

The following example displays a VPRN interface SAP configuration:

```
exit
        vprn 1435 customer 1 create
            shutdown
            interface "test" create
            exit
```

```
            bgp
                no shutdown
            exit
            ospf
                no shutdown
            exit
        exit
-------------------------------------------
*A:K-SASK12>config>service#
```

Configuring VPRN Protocols - OSPF

In a VPRN interface, each VPN routing instance is isolated from any other VPN routing instance, and from the routing used across the backbone. OSPF can be run with any VPRN, independently of the routing protocols used in other VPRNs, or in the backbone itself. For more information about the OSPF protocol, see the *7210 SAS OS Routing Protocols Guide*.

CLI Syntax: `config>service>vprn>ospf#`

VPRN OSPF CLI Syntax

Refer to Configuring VPRN Protocols - OSPF on page 685 for CLI syntax to configure VPRN parameters.

The following example displays the VPRN OSPF configuration shown above:

```
*A:K-SASK12>config>service>vprn>ospf$ info detail
-----------------------------------------------
                router-id 10.10.10.1
                no rtr-adv-lsa-limit
                no super-backbone
                no vpn-domain
                vpn-tag 0
                reference-bandwidth 100000000
                compatible-rfc1583
                no external-db-overflow
                no overload-include-ext-2
                no overload-include-stub
                no overload
                no overload-on-boot
                timers
                    no spf-wait
```

```
                    no lsa-generate
                    no lsa-arrival
                    no incremental-spf-wait
                    no lsa-accumulate
                    no redistribute-delay
                exit
                preference 10
                external-preference 150
                no rib-priority high
                no export-limit
                no export
                no graceful-restart
                no ignore-dn-bit
                no suppress-dn-bit
                no import
                no advertise-router-capability
                no loopfree-alternate
                no loopfree-alternate-exclude
                no shutdown
-------------------------------------------
*A:K-SASK12>config>service>vprn>ospf$
```

For more information about the OSPF protocol, refer to the *7210 SAS OS Routing Protocols Guide*.

Service Management Tasks

This section discusses the following service management tasks:

- Modifying VPRN Service Parameters on page 604
- Deleting a VPRN Service on page 605
- Disabling a VPRN Service on page 606
- Re-enabling a VPRN Service on page 607

Modifying VPRN Service Parameters

Use the CLI syntax to modify VPRN parameters (VPRN Services Command Reference on page 681).

The following example displays the VPRN service creation output.

```
vprn 2 customer 1 create
            shutdown
            no description
            dhcp
            exit
            dhcp6
            exit
            no router-id
            no maximum-routes
            no maximum-ipv6-routes
```

```
snmp
    no access
exit
log
exit
no mc-maximum-routes
no autonomous-system
no route-distinguisher
auto-bind-tunnel
    resolution-filter
        no ldp
        no rsvp
    exit
    resolution disabled
exit
no enable-bgp-vpn-backup
no vrf-target
interface "test" create
    no description
    no enable-ingress-stats
    no enable-mac-accounting
    no tcp-mss
    cpu-protection 254
    no address
    no mac
    arp-timeout 14400
    arp-retry-timer 50
    no arp-limit
    no allow-directed-broadcasts
    icmp
        mask-reply
        redirects 100 10
        unreachables 100 10
        ttl-expired 100 10
    exit
    dhcp
        shutdown
        no description
        no option
        no server
        no trusted
        no relay-proxy
        no gi-address
        no relay-plain-bootp
    exit
```

Deleting a VPRN Service

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

Use the following CLI syntax to delete a VPRN service:

CLI Syntax: config>service#

[no] vprn service-id [customer customer-id]

shutdown

[no] interface ip-int-name

shutdown

[no] sap sap-id]

[no] bgp

shutdown

[no] spoke-sdp sdp-id

[no] shutdown

Disabling a VPRN Service

A VPRN service can be shut down without deleting any service parameters.

CLI Syntax: config>service#

vprn service-id [customer customer-id]

shutdown

Example: config>service# vprn 1

config>service>vprn# shutdown

config>service>vprn# exit

```
vprn 2 customer 1 create
            shutdown
            no description
            dhcp
            exit
            dhcp6
            exit
            no router-id
            no maximum-routes
            no maximum-ipv6-routes
            snmp
                no access
            exit
            log
```

```
            exit
            no mc-maximum-routes
            no autonomous-system
            no route-distinguisher
            auto-bind-tunnel
                resolution-filter
                    no ldp
                    no rsvp
                exit
                resolution disabled
            exit
            no enable-bgp-vpn-backup
            no vrf-target
            interface "test" create
                no description
                no enable-ingress-stats
                no enable-mac-accounting
                no tcp-mss
                cpu-protection 254
                no address
                no mac
                arp-timeout 14400
                arp-retry-timer 50
                no arp-limit
                no allow-directed-broadcasts
                icmp
                    mask-reply
                    redirects 100 10
                    unreachables 100 10
                    ttl-expired 100 10
                exit
                dhcp
                    shutdown
                    no description
                    no option
                    no server
                    no trusted
                    no relay-proxy
                    no gi-address
                    no relay-plain-bootp
                exit

#
```

Re-enabling a VPRN Service

To re-enable a VPRN service that was shut down.

CLI Syntax: `config>service#`

```
vprn service-id [customer customer-id]

no shutdown
```

VPRN Services Command Reference for 7210 SAS-K 2F4T6C

Command Hierarchies

- VPRN Service Configuration Commands on page 610
- DHCP Server Commands on page 611
- IPoE Commands on page 613
- Interface Commands on page 615
- Routed VPLS Commands for VPRN Service on page 616
- Interface SAP Commands on page 617
- BGP Configuration Commands on page 619
- OSPF Configuration Commands on page 623
- Show Commands on page 626
- Clear Commands on page 628
- Debug Commands on page 629

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — aggregate ip-prefix /ip-prefix-length [summary-only] [as-set] [aggregator as-number: ip-address]
            — no aggregate  ip-prefix/ ip-prefix-length
            — auto-bind-tunnel
                — resolution {any | filter | disabled}
                — [no] resolution-filter ldp
                — [no] resolution-filter rsvp
            — description description-string
            — no description
            — [no] ignore-nh-metric
            — maximum-routes number [log-only] [threshold percent]
            — no maximum-routes
            — route-distinguisher [ip-address:number1 | asn:number2]
            — no route-distinguisher
            — router-id ip-address
            — no router-id
            — [no] shutdown
            — snmp
                — community community community-name  [hash|hash2] [access-permissions] [version SNMP-version] [src-access-list list-name]
                — no community community-name
            — source-address
                — application app [ip-int-name | ip-address]
                — no application app
            — [no] spoke-sdp sdp-id
                — [no] description description string
            — [no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag tag] [enable | disable] {next-hop ip-int-name|ip-address | {cpe-check cpe-ip-address [interval seconds] [drop-count count] [log]}} {prefix-list prefix-list-name [all|none]}]
            — [no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag tag] [enable | disable] indirect ip-address [cpe-check cpe-ip-address [interval seconds][drop-count count] [log]] {prefix-list prefix-list-name [all|none]}]
            — [no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag tag] [enable | disable] black-hole {prefix-list prefix-list-name [all|none]}]
            — [no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference][metric metric] [tag tag] [enable | disable] next-hop ip-int-name|ip-address [bfd-enable|{cpe-check cpe-ip-address [interval seconds] [drop-count count] [log]}|{prefix-list prefix-list-name [all|none]}]][description description]
            — vrf-export policy-name [policy-name...(upto 5 max)]
            — no vrf-export
            — vrf-import policy-name [policy-name...(upto 5 max)]
            — no vrf-import
            — vrf-target {ext-community|export ext-community|import ext-community}
            — no vrf-target
            — [no] shutdown
```

DHCP Server Commands

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — dhcp
                — local-dhcp-server server-name [create]
                — no local-dhcp-server server-name
                    — description description-string
                    — no description
                    — [no] force-renews
                    — pool pool-name [create]
                    — no pool pool-name
                        — description description-string
                        — no description
                        — max-lease-time [days days] [hrs hours] [min minutes] [sec seconds]
                        — no max-lease-time
                        — min-lease-time [days days] [hrs hours] [min minutes] [sec seconds]
                        — no min-lease-time
                        — minimum-free minimum-free [percent] [event-when-depleted]
                        — no minimum-free
                        — [no] nak-non-matching-subnet
                        — offer-time [min minutes] [sec seconds]
                        — no offer-time
                        — options
                            — custom-option option-number address [ip-address...(up to 4 max)]
                            — custom-option option-number hex hex-string
                            — custom-option option-number string ascii-string
                            — no custom-option option-number
                            — dns-server [ip-address...(up to 4 max)]
                            — domain-name domain-name
                            — no domain-name
                            — lease-rebind-time [days days] [hrs hours] [min minutes] [sec seconds]
                            — no lease-rebind-time
                            — lease-renew-time [days days] [hrs hours] [min minutes] [sec seconds]
                            — no lease-renew-time
                            — lease-time [days days] [hrs hours] [min minutes] [sec seconds]
                            — no lease-time
                        — subnet {ip-address/mask|ip-address netmask} [create]
                        — no subnet {ip-address/mask|ip-address netmask}
                            — [no] address-range start-ip-address end-ip-address
                            — [no] exclude-addresses start-ip-address [end-ip-address]
                            — maximum-declined maximum-declined
```

- **no maximum-declined**
- **minimum-free** *minimum-free*
- **no minimum-free**
- **options**
  - **custom-option** *option-number* **address** [*ip-address...*(upto 4 max)]
  - **custom-option** *option-number* **hex** *hex-string*
  - **custom-option** *option-number* **string** *ascii-string*
  - **no custom-option** *option-number*
  - **default-router** *ip-address* [*ip-address...*(up to 4 max)]
  - **no default-router**
  - **subnet-mask** *ip-address*
  - **no subnet-mask**
- [**no**] **use-gi-address**
- **user-db** *local-user-db-name*
- **no user-db**

IPoE Commands

```
config
    — subscriber-mgmt
        — local-user-db local-user-db-name [create]
        — no local-user-db local-user-db-name
            — description description-string
            — no description
            — ipoe
                — host host-name [create]
                — no host host-name
                    — address gi-address [scope scope]
                    — address ip-address
                    — address pool pool-name [secondary-pool sec-pool-name] [delimiter delimiter]
                    — address use-pool-from-client [delimiter delimiter]
                    — no address
                    — gi-address ip-address
                    — no gi-address
                    — host-identification
                        — circuit-id string ascii-string
                        — circuit-id hex hex-string
                        — no circuit-id
                        — encap-tag-range start-tag start-tag end-tag end-tag
                        — no encap-tag-range
                        — mac ieee-address
                        — no mac
                        — option60 hex-string
                        — no option60
                        — remote-id hex hex-string
                        — remote-id string ascii-string
                        — no remote-id
                        — sap-id sap-id
                        — no sap-id
                        — service-id service-id
                        — no service-id
                        — string string
                        — no string
                        — system-id system-id
                        — no system-id
                        — shutdown
                        — no shutdown
                    — options
                        — custom-option option-number address [ip-address...(up to 4 max)]
                        — custom-option option-number hex hex-string
                        — custom-option option-number string ascii-string
                        — no custom-option option-number
                        — default-router ip-address [ip-address (up to 4 max)]
                        — no default-router
                        — dns-server ip-address [ip-address...(up to 4 max)]
                        — no dns-server
                        — domain-name domain-name
```

- — **no domain-name**
- — **lease-rebind-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- — **no lease-rebind-time**
- — **lease-renew-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- — **no lease-renew-time**
- — **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- — **no lease-time**
- — **subnet-mask** *ip-address*
- — **no subnet-mask**

— **mask type** *ipoe-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}
— **no mask type** *ipoe-match-type*
— **match-list** *ipoe-match-type-1* [*ipoe-match-type-2* (up to 4 max)]
— **no match-list**
— **shutdown**
— **no shutdown**

Interface Commands

```
config
    — service
        — vprn service-id [customer customer-id] [create]
        — no vprn service-id
            — [no] interface ip-int-name
                — address ip-address[/mask] [netmask] [broadcast {all-ones | host-ones}]
                — no address
                — [no] allow-directed-broadcasts
                — arp-timeout [seconds]
                — no arp-timeout
                — delayed-enable seconds
                — no delayed-enable
                — description description-string
                — no description [description-string]
                — dhcp
                    — description description-string
                    — no description
                    — gi-address ip-address [src-ip-addr]
                    — no gi-address
                    — [no] option
                        — action {replace | drop | keep}
                        — no action
                        — circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tuple]
                        — no circuit-id
                        — 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
                    — server server1 [server2...(up to 8 max)]
                    — no server
                    — [no] shutdown
                    — [no] trusted
                — icmp
                    — [no] mask-reply
                    — redirects number seconds
                    — no redirects [number seconds]
                    — ttl-expired number seconds
                    — no ttl-expired [number seconds]
                    — unreachables number seconds
                    — no unreachables [number seconds]
                — [no] local-proxy-arp
                — [no] loopback
                — [no] remote-proxy-arp <policy-name> [<policy-name>...(upto 5 max)]
                — [no] remote-proxy-arp
                — [no] shutdown
                — [no] static-arp ip-address [ieee-address]
                    —
```

Routed VPLS Commands for VPRN Service

```
config
    — service
        — vprn
            — interface ip-interface-name [create]
            — no interface ip-interface-name
                — vpls service-name
                — no vpls
                    — ingress
                        — v4-routed-override-filter ip-filter-id
                        — no v4-routed-override-filter
```

Interface SAP Commands

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — [no] interface ip-int-name [create] [tunnel]
                — [no] sap sap-id
                    — accounting-policy acct-policy-id
                    — no accounting-policy [acct-policy-id]
                    — [no] collect-stats
                    — description description-string
                    — no description [description-string]
                    — egress
                        — agg-shaper-rate agg-rate
                        — no agg-shaper-rate
                        — filter ip ip-filter-id
                        — no filter [ip ip-filter-id]
                        — qos policy-id
                        — no qos [policy-id]
                        —
                    — ingress
                        — meter-override
                        — no meter-override
                            — meter meter-id [create]
                            — no meter meter-id
                                — adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
                                — cbs size-in-kbytes (supported on 7210 SAS-E)
                                — no cbs
                                — cbs size [kbits|bytes|kbytes] (supported on 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C)
                                — no cbs
                                — mbs size-in-kbits (supported on 7210 SAS-E)
                                — no mbs
                                — mbs size [kbits|bytes|kbytes] (supported on 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C)
                                — no mbs
                                — no mode
                                — no mode
                                — rate cir cir-rate [pir pir-rate]
                    — ingress
                        — agg-shaper-rate agg-rate
                        — no agg-shaper-rate
                        — aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
                        — no aggregate-meter-rate
                        — filter ip ip-filter-id
                        — no filter [ip ip-filter-id]
                        — qos policy-id [shared-queuing | multipoint-shared]
```

— **no qos** [*policy-id*]
— [**no**] **shutdown**

BGP Configuration Commands

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — [no] bgp
                — [no] advertise-inactive
                — [no] aggregator-id-zero
                — always-compare-med {zero | infinity}
                — no always-compare-med
                — [no] as-path-ignore
                — auth-keychain name
                — authentication-key [authentication-key | hash-key] [hash | hash2]
                — no authentication-key
                — bfd-enable
                — no bfd-enable
                — best-path-selection
                    —always-compare-med {zero | infinity}
                    —always-compare-med strict-as {zero | infinity}
                    —no always-compare-med
                    —as-path-ignore [ipv4] [ipv6]
                    — no as-path-ignore
                    —as-path-ignore
                    —no ignore-nh-metric
                    —ignore-nh-metric
                    —no ignore-router-id
                    —ignore-nh-metric
                — [no] connect-retry seconds
                — [no] damping
                — description description-string
                — no description
                — [no] disable-4byte-asn
                — [no] disable-client-reflect
                — disable-communities [standard] [extended]
                — no disable-communities
                — [no] disable-fast-external-failover
                — [no] enable-peer-tracking
                — export policy-name [policy-name...(upto 5 max)]
                — no export
                — family [ipv4]
                — no family
                — hold-time seconds [strict]
                — no hold-time
                — import policy-name [policy-name...(up to 5 max)]
                — no import
                — keepalive seconds
                — no keepalive
                — local-as as-number [private]
                — no local-as
                — local-preference local-preference
                — no local-preference
                — loop-detect {drop-peer | discard-route | ignore-loop| off}
                — no loop-detect
                — med-out {number | igp-cost}
```

— **no med-out**
— **min-as-origination** *seconds*
— **no min-as-origination**
— **min-route-advertisement** *seconds*
— **no min-route-advertisement**
— **multihop** *ttl-value*
— **no multihop**
— **preference** *preference*
— **no preference**
— [**no**] **rapid-withdrawal**
— [**no**] **remove-private**
— **router-id** *ip-address*
— **no router-id**
— [**no**] **shutdown**
— [**no**] **group** *name* [**dynamic-peer**]
    — [**no**] **advertise-inactive**
    — [**no**] **aggregator-id-zero**
    — [**no**] **as-override**
    — **auth-keychain** *name*
    — **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
    — **no authentication-key**
    — [**no**] **bfd-enable**
    — **cluster** *cluster-id*
    — **no cluster**
    — **connect-retry** *seconds*
    — **no connect-retry**
    — [**no**] **damping**
    — **description** *description-string*
    — **no description**
    — [**no**] **disable-4byte-asn**
    — [**no**] **disable-client-reflect**
    — **disable-communities** [**standard**] [**extended**]
    — **no disable-communities**
    — [**no**] **disable-fast-external-failover**
    — [**no**] **enable-peer-tracking**
    — **export** *policy-nam*e [*policy-name*...(upto 5 max)]
    — **no export**
    — **family** [**ipv4**]
    — **no family**
    — **hold-time** *seconds* [**strict**]
    — **no hold-time**
    — **import** *policy-name* [*policy-name*...(upto 5 max)]
    — **no import**
    — **keepalive** *seconds*
    — **no keepalive**
    — **local-address** *ip-address*
    — **no local-address**
    — **local-as** *as-number* [**private**]
    — **no local-as**
    — **local-preference** *local-preference*
    — **no local-preference**
    — **loop-detect** {**drop-peer**|**discard-route**|**ignore-loop**|**off**}
    — **no loop-detect**
    — **med-out** {**number** | **igp-cost**}

- **no med-out**
- **min-as-origination** *seconds*
- **no min-as-origination**
- **min-route-advertisement** *seconds*
- **no min-route-advertisement**
- **multihop** *ttl-value*
- **no multihop**
- [**no**] **passive**
- **peer-as** *as-number*
- **no peer-as**
- **preference** *preference*
- **no preference**
- **prefix-limit** *limit* [**log-only**] [**threshold** *percent*]
- **no prefix-limit**
- [**no**] **remove-private**
- [**no**] **shutdown**
- **type** {**internal** | **external**}
- **no type**
- [**no**] **neighbor** *ip-address*
  - [**no**] **advertise-inactive**
  - [**no**] **aggregator-id-zero**
  - [**no**] **as-override**
  - **auth-keychain** *name*
  - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
  - **no authentication-key**
  - [**no**] **bfd-enable**
  - **cluster** ***cluster-id***
  - **no cluster**
  - **connect-retry** *seconds*
  - **no connect-retry**
  - [**no**] **damping**
  - **description** *description-string*
  - **no description**
  - [**no**] **disable-4byte-asn**
  - [**no**] **disable-client-reflect**
  - **disable-communities** [**standard**] [**extended**]
  - **no disable-communities**
  - [**no**] **disable-fast-external-failover**
  - [**no**] **enable-peer-tracking**
  - **export** *policy-name* [*policy-name*...(upto 5 max)]
  - **no export**
  - **family** [**ipv4**]
  - **no family**
  - **hold-time** *seconds* [**strict**]
  - **no hold-time**
  - **import** *policy-name* [*policy-name*...(upto 5 max)]
  - **no import**
  - **keepalive** *seconds*
  - **no keepalive**
  - **local-address** *ip-address*
  - **no local-address**
  - **local-as** *as-number* [**private**]
  - **no local-as**
  - **local-preference** *local-preference*

— **no local-preference**
— **loop-detect** {**drop-peer** | **discard-route** | **ignore-loop** | **off**}
— **no loop-detect**
— **med-out** {**number** | **igp-cost**}
— **no med-out**
— **min-as-origination** *seconds*
— **no min-as-origination**
— **min-route-advertisement** *seconds*
— **no min-route-advertisement**
— **multihop** *ttl-value*
— **no multihop**
— [**no**] **passive**
— **peer-as** *as-number*
— **no peer-as**
— **preference** *preference*
— **no preference**
— **prefix-limit** *limit* [**log-only**] [**threshold** *percent*]
— **no prefix-limit**
— [**no**] **remove-private**
— [**no**] **shutdown**
— **type** {**internal** | **external**}
— **no type**

OSPF Configuration Commands

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — [no] ospf
                — [no] area area-id
                    — area-range ip-prefix/mask [advertise | not-advertise]
                    — no area-range ip-prefix/mask
                    — [no] blackhole-aggregate
                    — [no] interface ip-int-name [secondary]
                        — [no] advertise-subnet
                        — authentication-key [authentication-key | hash-key] [hash | hash2]
                        — no authentication-key
                        — authentication-type {password | message-digest}
                        — no authentication-type
                        — bfd-enable [remain-down-on-failure]
                        — no bfd-enable
                        — dead-interval seconds
                        — no dead-interval
                        — hello-interval seconds
                        — no hello-interval
                        — interface-type {broadcast | point-to-point}
                        — no interface-type
                        — message-digest-key key-id md5 [key | hash-key] [hash | hash2]
                        — no message-digest-key key-id
                        — metric metric
                        — no metric
                        — mtu bytes
                        — no mtu
                        — [no] passive
                        — priority number
                        — no priority
                        — retransmit-interval seconds
                        — no retransmit-interval
                        — [no] shutdown
                        — transit-delay seconds
                        — no transit-delay
                    — [no] nssa
                        — area-range ip-prefix/mask [advertise | not-advertise]
                        — no area-range ip-prefix/mask
                        — originate-default-route [type-7]
                        — no originate-default-route
                        — [no] redistribute-external
                        — [no] summaries
                    — [no] sham-link ip-int-name ip-address
                        — authentication-key [authentication-key | hash-key] [hash | hash2]
                        — no authentication-key
                        — authentication-type {password | message-digest}
                        — no authentication-type
```

— **dead-interval** *seconds*
— **no dead-interval**
— **hello-interval** *seconds*
— **no hello-interval**
— **message-digest-key** *key-id* **md5** [*key* | *hash-key*] [**hash** | **hash2**]
— **no message-digest-key** *key-id*
— **metric** *metric*
— **no metric**
— **retransmit-interval** *seconds*
— **no retransmit-interval**
— [**no**] **shutdown**
— **transit-delay** *seconds*
— **no transit-delay**

— [**no**] **stub**
— **default-metric** *metric*
— **no default-metric**
— [**no**] **summaries**

— [**no**] **virtual-link** *router-id* **transit-area** *area-id*
— **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
— **no authentication-key**
— **authentication-type** {**password** | **message-digest**}
— **no authentication-type**
— **dead-interval** *seconds*
— **no dead-interval**
— **hello-interval** *seconds*
— **no hello-interval**
— **message-digest-key** *key-id* **md5** [*key* | *hash-key*] [**hash** | **hash2**]
— **no message-digest-key** *key-id*
— **retransmit-interval** *seconds*
— **no retransmit-interval**
— [**no**] **shutdown**
— **transit-delay** *seconds*
— **no transit-delay**

— [**no**] **compatible-rfc1583**
— **export** *policy-name* [*policy-name*...(up to 5 max)]
— **no export**
— **external-db-overflow** *limit seconds*
— **no external-db-overflow**
— **external-preference** *preference*
— **no external-preference**
— [**no**] **ignore-dn-bit**
— **overload** [**timeout** *seconds*]
— **no overload**
— [**no**] **overload-include-stub**
— **overload-on-boot** [**timeout** *seconds*]
— **no overload-on-boot**
— **preference** *preference*
— **no preference**
— **reference-bandwidth** *bandwidth-in-kbps*
— **no reference-bandwidth**
— **router-id** *ip-address*
— **no router-id**

- [**no**] **shutdown**
- [**no**] **super-backbone**
- [**no**] **suppress-dn-bit**
- **timers**
  - [**no**] **lsa-arrival** ***lsa-arrival-time***
  - [**no**] **lsa-generate** *max-lsa-wait* [*lsa-initial-wait* [*lsa-second-wait*]]
  - [**no**] **spf-wait** ***max-spf-wait*** [*spf-initial-wait* [*spf-second-wait*]]
- **vpn-domain** *id* {**0005** | **0105** | **0205** | **8005**}
- **no vpn-domain**
- **vpn-tag** *vpn-tag*
- **no vpn-tag**

```
show
    — service
        — egress-label start-label [end-label]
        — ingress-label start-label [[end-label]
        — id service-id
            — all
            — sap [sap-id [detail]]
            — sdp [sdp-id | far-end ip-address] [detail]
        — sap-using [sap sap-id]
        — sap-using interface [ip-address | ip-int-name]
        — sap-using [ingress | egress] filter filter-id
        — sap-using [ingress | egress] qos-policy qos-policy-id
        — sdp-using [sdp-id | far-end ip-address] [detail | keep-alive-history]
        — sdp-using [sdp-id[:vc-id]
        — service-using [vprn] [sdp sdp-id] [customer customer-id]
```

show
- **router** [*vprn-service-id*]
  - **aggregate** [<family>] [**active**]
  - **arp** [<ip-int-name|ip-address[/mask]>|**mac**<ieee-mac-address>|sumary]**[local|dynamic|static|managed]**
  - **bgp**
    - **damping** [*ip-prefix*[/*prefix-length*]] [**decayed**|**history**|**suppressed**] [detail] [ipv4]
    - **damping** [*ip-prefix*[/*prefix-length*]] [**decayed**|**history**|**suppressed**] [detail] vpn-ipv4
    - **group** [*name*] [**detail**]
    - **neighbor** [*ip-address* [[**family** *family*] *filter1*]]
    - **neighbor** [*as-number* [[**family** *family*] *filter2*]]
    - **paths**
    - **routes** [**family** *family*] [*prefix* [**detail** | **longer**]]
    - **routes** [**family** *family*] [*prefix* [**hunt** | **brief**]]
    - **routes** [**family** *family*] [**community** *comm-id*]
    - **routes** [**family** *family*] [**aspath-regex** *reg-ex*1]
    - **routes** [**family**] [*ipv6-prefix*[/*prefix-length*] [**detail** | **longer**]|[**hunt** [**brief**]]]
    - **summary** [**all**]
  - **interface** [{[*ip-address* | *ip-int-name*] [**detail**]} | **summary**
  - **route-table** [*ip-address*[/*mask*] [**longer | best**]] | [**protocol** *protocol*] | [**summary**]
  - **static-arp** [*ip-address* | *ip-int-name* | **mac** *ieee-mac-address*]
  - **static-route** [*ip-prefix* /*mask*] | [**preference** *preference*] | [**next-hop** *ip-address*] [**detail**]
  - **tunnel-table** [*ip-address*[/*mask*] [**protocol** *protocol* | **sdp** *sdp-id*]
  - **tunnel-table** [**summary**]

```
clear
    — router
        — bgp
            — damping [{prefix/mask [neighbor ip-address]} | {group name}]
            — flap-statistics [[ip-prefix/mask] [neighbor ip-address]] | [group group-name] |
              [regex reg-exp] | [policy policy-name]
            — neighbor {ip-address | as as-number | external | all} [soft | soft-inbound | statis-
              tics]
            — protocol
        — interface [ip-int-name | ip-address] [icmp]
clear
    — service
        — id service-id
            — spoke-sdp sdp-id:vc-id ingress-vc-label
        — statistics
            — sap sap-id {all | counters | stp}
            — sdp sdp-id keep-alive
            — id service-id
                — counters
                — spoke-sdp sdp-id:vc-id {all | counters | stp}
                — spoke-sdp
```

Debug Commands

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

VPRN Service Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>service>vprn
config>service>vprn>if
config>service>vprn>if>sap
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
config>service>vprn>spoke-sdp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

If the AS number was previously changed, the BGP AS number inherits the new value.

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.

A service is regarded as operational providing that one IP Interface SAP and one SDP is operational.

VPRN BGP — This command disables the BGP instance on the given IP interface. Routes learned from a neighbor that is shutdown are immediately removed from the BGP database and RTM. If BGP is globally shutdown, then all group and neighbor interfaces are shutdown operationally. If a BGP group is shutdown, all member neighbor interfaces are shutdown operationally. If a BGP neighbor is shutdown, just that neighbor interface is operationally shutdown.

description

Syntax **description** *description-string*
no description

Context config>service>vprn>bgp
config>service>vprn
config>service>vprn>if
config>service>vprn>if>sap
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Default No description associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Global Commands

vprn

Syntax **vprn** *service-id* [**customer** *customer-id*] [**create**]
no vprn *service-id*

Context config>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C

This command creates or edits a Virtual Private Routed Network (VPRN) service instance.

If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

VPRN services allow the creation of customer-facing IP interfaces in the same routing instance used for service network core routing connectivity. VPRN services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.

IP interfaces defined within the context of an VPRN service ID must have a SAP created as the access point to the subscriber network.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the customer command in the service context. When a service is created with a customer association, it is not possible to edit the customer association. The service must be deleted and re-created with a new customer association.

When a service is created, the use of the **customer** *customer-id* is optional to navigate into the service configuration context. If attempting to edit a service with the incorrect *customer-id* results in an error.

Multiple VPRN services are created to separate customer-owned IP interfaces. More than one VPRN service can be created for a single customer ID. More than one IP interface can be created within a single VPRN service ID. All IP interfaces created within an VPRN service ID belongs to the same customer.

The **no** form of the command deletes the VPRN service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces and all routing protocol configurations defined within the service ID have been shutdown and deleted.

Default None — No VPRN service instances exist until they are explicitly created.

Parameters *service-id* — The unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values

service-id:	1 — 2147483648
svc-name:	64 characters maximum

customer *customer-id* — Specifies an existing customer identification number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 — 2147483647

aggregate

Syntax **aggregate** *ip-prefix/ip-prefix-length* [**summary-only**] [**as-set**] [**aggregator** *as-number:ip-address*]
aggregate *ip-prefix/ip-prefix-length*

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command creates an aggregate route.

Use this command to automatically install an aggregate in the routing table when there are one or more component routes. A component route is any route used for forwarding that is a more-specific match of the aggregate.

The use of aggregate routes can reduce the number of routes that need to be advertised to neighbor routers, leading to smaller routing table sizes.

Overlapping aggregate routes may be configured; in this case a route becomes a component of only the one aggregate route with the longest prefix match. For example if one aggregate is configured as 10.0.0.0/16 and another as 10.0.0.0/24, then route 10.0.128/17 would be aggregated into 10.0.0.0/16, and route 10.0.0.128/25 would be aggregated into 10.0.0.0/24. If multiple entries are made with the same prefix and the same mask the previous entry is overwritten.

A standard 4-byte BGP community may be associated with an aggregate route in order to facilitate route policy matching.

The **no** form of the command removes the aggregate.

Default no aggregate

Parameters *ip-prefix* — the destination address of the aggregate route in dotted decimal notation

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32

the mask associated with the network address expressed as a mask length

Values: 0 to 32

summary-only — This optional parameter suppresses advertisement of more specific component routes for the aggregate.

To remove the **summary-only** option, enter the same aggregate command without the **summary-only** parameter.

as-set — This optional parameter is only applicable to BGP and creates an aggregate where the path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized. Use this feature carefully as it can increase the amount of route churn due to best path changes.

aggregator *as-number***:***ip-address* **—** This optional parameter specifies the BGP aggregator path attribute to the aggregate route. When configuring the aggregator, a two-octet AS number used to form the aggregate route must be entered, followed by the IP address of the BGP system that created the aggregate route.

allow-ip-int-bind

Syntax **[no] allow-ip-int-bind**

Context config>service>vpls

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables IP interface binding.

The no form of the command disables IP interface binding.

Default no allow-ip-int-bind

auto-bind-tunnel

Syntax **auto-bind-tunnel**

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure automatic binding of a VPRN service using tunnels to MP-BGP peers.

The **auto-bind-tunnel** node is simply a context to configure the binding of VPRN routes to tunnels. The user must configure the **resolution** option to enable auto-bind resolution to tunnels in TTM. If the **resolution** option is explicitly set to **disabled**, the auto-binding to tunnel is removed.

If **resolution** is set to **any**, any supported tunnel type in VPRN context will be selected following TTM preference. If one or more explicit tunnel types are specified using the **resolution-filter** option, then only these tunnel types will be selected again following the TTM preference.

The following tunnel types are supported in a VPRN context in order of preference: RSVP, LDP, Segment Routing (SR), and GRE. The BGP tunnel type is not explicitly configured and is thus implicit. It is always preferred over any other tunnel type enabled in the **auto-bind-tunnel** context.

The **ldp** value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.

The **rsvp** value instructs BGP to search for the best metric RSVP LSP to the address of the BGP next-hop. This address can correspond to the system interface or to another loopback used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id.

When the **sr-isis** (**sr-ospf**) value is enabled, a SR tunnel to the BGP next-hop is selected in the TTM from the lowest numbered ISIS (OSPF) instance.

The **sr-te** value instructs BGP to search for the best metric SR-TE LSP to the address of the BGP next-hop. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple SR-TE LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id.

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under **resolution-filter**.

When an explicit SDP to a BGP next-hop is configured in a VPRN service (**configure>service>vprn>spoke-sdp**), it overrides the **auto-bind-tunnel** selection for that BGP next-hop only. There is no support for reverting automatically to the **auto-bind-tunnel** selection if the explicit SDP goes down. The user must delete the explicit spoke-sdp in the VPRN service context to resume using the **auto-bind-tunnel** selection for the BGP next-hop.

resolution

Syntax **resolution** {**any** | **filter** | **disabled**}

Context config>service>vprn>auto-bind-tunnel

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the resolution mode in the automatic binding of a VPRN service to tunnels to MP-BGP peers.

Parameters *any* — enables the binding to any supported tunnel type in VPRN context following TTM preference

filter — enables the binding to the subset of tunnel types configured under resolution-filter

disabled — disables the automatic binding of a VPRN service to tunnels to MP-BGP peers

resolution-filter

Syntax **resolution-filter**

Context config>service>vprn>auto-bind-tunnel

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the subset of tunnel types which can be used in the resolution of VPRN prefixes within the automatic binding of VPRN service to tunnels to MP-BGP peers.

The following tunnel types are supported in a VPRN context in order of preference: RSVP, LDP, Segment Routing (SR), and GRE. The BGP tunnel type is not explicitly configured and is thus implicit. It is always preferred over any other tunnel type enabled in the **auto-bind-tunnel** context.

Parameters **ldp —** selects the LDP tunnel type

rsvp — selects the RSVP-TE tunnel type

autonomous-system

Syntax **autonomous-system** *as-number*
no autonomous-system

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command defines the autonomous system (AS) to be used by this VPN routing/forwarding (VRF). This command defines the autonomous system to be used by this VPN routing

The **no** form of the command removes the defined AS from this VPRN context.

Default no autonomous-system

Parameters *as-number* — Specifies the AS number for the VPRN service.

Values 1 — 4294967295

Values

maximum-routes

Syntax **maximum-routes** *number* [**log-only**] [**threshold** *percentage*]
no maximum-routes

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the maximum number of remote routes that can be held within a VPN routing/ forwarding (VRF) context. Note that **local**, **host**, **static** and **aggregate** routes are not counted.

Note that the VPRN service ID must be in a shutdown state in order to modify maximum-routes command parameters.

If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then the offending RIP peer (if applicable) is brought down (but the VPRN instance remains up). BGP peering will remain up but the exceeding BGP routes will not be added to the VRF.

The maximum route threshold can dynamically change to increase the number of supported routes even when the maximum has already been reached. Protocols will resubmit their routes which were initially rejected.

The **no** form of the command disables any limit on the number of routes within a VRF context. Issue the **no** form of the command only when the VPRN instance is shutdown.

Default 0 or disabled — The threshold will not be raised.

Parameters *number* — An integer that specifies the maximum number of routes to be held in a VRF context.

Values 1 — 2147483647

log-only — This parameter specifies that if the maximum limit is reached, only log the event. **log-only** does not disable the learning of new routes.

threshold *percentage* — The percentage at which a warning log message and SNMP trap should be set. There are two warnings, the first is a mid-level warning at the threshold value set and the second is a high-level warning at level between the maximum number of routes and the mid-level rate ([mid+max] / 2).

Values 0 — 100

route-distinguisher

Syntax **route-distinguisher** [*ip-address*:*number* | *asn*:*number*]
no route-distinguisher

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command sets the identifier attached to routes the VPN belongs to. Each routing instance must have a unique (within the carrier's domain) route distinguisher associated with it. A route distinguisher must be defined for a VPRN to be operationally active.

Default no route-distinguisher

Parameters The route distinguisher is a 6-byte value that can be specified in one of the following formats:

ip-address:number — Specifies the IP address in dotted decimal notation. The assigned number must not be greater than 65535.

asn:number — The ASN is a 2-byte value less than or equal to 65535. The assigned number can be any 32-bit unsigned integer value.

router-id

Syntax **router-id** *ip-address*
no router-id

Context config>service>vprn
config>service>vprn>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command sets the router ID for a specific VPRN context.

If neither the router ID nor system interface are defined, the router ID from the base router context is inherited.

The **no** form of the command removes the router ID definition from the given VPRN context.

Default no router-id

Parameters *ip-address* — The IP address must be given in dotted decimal notation.

service-name

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

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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 7210 SAS 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).

Platforms Supported: 7210 SAS-K 2F4T6C.

snmp

Syntax **snmp**

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure SNMP parameters for this VPRN.

community

Syntax **community** *community-name* [**hash** | **hash2**] [**access-permissions**] [**version** *SNMP-version*] [**src-access-list list-name**]
no community [*community-name*]

Context config>service>vprn>snmp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command sets the SNMP community name(s) to be used with the associated VPRN instance. These VPRN community names are used to associate SNMP v1/v2c requests with a particular vprn context and to return a reply that contains VPRN-specific data or limit SNMP access to data in a specific VPRN instance.

VPRN snmp communities configured with an access permission of 'r' are automatically associated with the default access group "snmp-vprn-ro" and the "vprn-view" view (read only). VPRN snmp communities configured with an access permission of 'rw' are automatically associated with the default access group "snmp-vprn" and the "vprn-view" view (read/write).

The community in an SNMP v1/v2 request determines the SNMP context (i.e., the vprn# for accessing SNMP tables) and not the VPRN of the incoming interface on which the request was received. When an SNMP request arrives on VPRN 5 interface "ringo" with a destination IP address equal to the "ringo" interface, but the community in the SNMP request is the community configured against VPRN 101, then the SNMP request will be processed using the VPRN 101 context. (the response will contain information about VPRN 101). It is recommended to avoid using a simple series of vprn snmp-community values that are similar to each other (for example, avoid my-vprncomm-1, my-vprn-comm-2, etc).

The **no** form of the command removes the SNMP community name from the given VPRN context.

Default n/a — The SNMP community must be explicitly specified.

Parameters *community-name* — Specifies the SNMP v1/v2c community name. This is a secret/confidential key used to access SNMP and specify a context (base vs vprn1 vs vprn2).

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

version *SNMP-version* **—** specifies the SNMP version

Values v1, v2c, both

access-permissions — specifies the access rights to MIB objects

Values **r** — Grants only read access to MIB objects. Creates an association of the community-name with the **snmp-vprn-ro** access group.
rw — Grants read and write access to MIB objects. Creates an association of the community-name with the **snmp-vprn** access group.

list-name — Configures the **community** to reference a specific **src-access-list** (created under **configure system security snmp**), which will be used to validate the source IP address of all received SNMP requests that use this **community**. Multiple **community** (vprn or base router) and **usm-community** instances can reference the same **src-access-list**.

source-address

Syntax **source-address**

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to specify the source address and application that should be used in all unsolicited packets.

application

Syntax **application** *app* [*ip-int-name*|*ip-address*]
no application *app*

Context config>service>vprn>source-address

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the source address and application.

Parameters *app* — Specify the application name.

Values telnet, ssh, traceroute, ping

ip-int-name | *ip-address* — Specifies the name of the IP interface or IP address. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

static-route

Syntax [**no**] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] {**next-hop** *ip-int-name* | *ip-address* | **ipsec-**

tunnel *ipsec-tunnel-name*} [**bfd-enable** | {**cpe-check** *cpe-ip-address* [**interval** *seconds*] [**drop-count** *count*] [**log**]}]
[**no**] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **indirect** *ip-address* [**cpe-check** *cpe-ip-address* [**interval** *seconds*][**drop-count** *count*] [**log**]]
[**no**] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **black-hole**

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command creates static route entries within the associated router instance. When configuring a static route, either **next-hop**, **indirect** or **black-hole** must be configured.

The **no** form of the command deletes the static route entry. If a static route needs to be removed when multiple static routes exist to the same destination, then as many parameters to uniquely identify the static route must be entered.

If a CPE connectivity check target address is already being used as the target address in a different static route, then cpe-check parameters must match. If they do not, the new configuration command will be rejected.

If a static-route command is issued with no cpe-check target but the destination prefix/netmask and next-hop matches a static route that did have an associated cpe-check, the cpe-check test will be removed from the associated static route.

Default No static routes are defined**.**

Parameters *ip-prefix* — The destination address of the aggregate route in dotted decimal notation.

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 — 32

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 — 255.255.255.255 (network bits all 1 and host bits all 0)

ip-int-name — The name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed with

ip-address — The IP address of the IP interface. The *ip-addr* 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 ipv4-address a.b.c.d (host bits must be 0)

enable — Static routes can be administratively enabled or disabled. Use the **enable** parameter to re-enable a disabled static route. In order to enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route.

The administrative state is maintained in the configuration file.

Default enable

disable — Static routes can be administratively enabled or disabled. Use the **disable** parameter to disable a static route while maintaining the static route in the configuration. In order to enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route.

The administrative state is maintained in the configuration file.

Default enable

interval *seconds* — This optional parameter specifies the interval between ICMP pings to the target IP address.

Values 1 —255 seconds

Default 1 seconds

drop-count *count* — This optional parameter specifies the number of consecutive ping-replies that must be missed to declare the CPE down and to de-active the associated static route.

Values Value range: 1 —255

Default 3

log — This optional parameter enables the ability to log transitions between active and in-active based on the CPE connectivity check. Events should be sent to the system log, syslog and SNMP traps.

next-hop [*ip-address* | *ip-int-name*] **—** Specifies the directly connected next hop IP address used to reach the destination. If the next hop is over an unnumbered interface, the *ip-int-name* of the unnumbered interface (on this node) can be configured.

The **next-hop** keyword and the **indirect** or **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of either the **indirect** or **black-hole** parameters), then this static route will be replaced with the newly entered command, and unless specified, the respective defaults for preference and metric will be applied.

The *ip-addr* configured here can be either on the network side or the access side on this node. This address must be associated with a network directly connected to a network configured on this node.

indirect *ip-address* **—** Specifies that the route is indirect and specifies the next hop IP address used to reach the destination.

The configured *ip-addr* is not directly connected to a network configured on this node. The destination can be reachable via multiple paths. The static route remains valid as long as the address configured as the indirect address remains a valid entry in the routing table. Indirect static routes cannot use an ip-prefix/mask to another indirect static route.

The **indirect** keyword and the **next-hop** or **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of either the **next-hop** or **black-hole** parameters), then this static route will be replaced with the newly entered command and unless specified the respective defaults for preference and metric will be applied.

The *ip-addr* configured can be either on the network or the access side and is normally at least one hop away from this node.

black-hole — Specifies a black hole route meaning that if the destination address on a packet matches this static route it will be silently discarded.

The **black-hole** keyword is mutually exclusive with either the **next-hop** or **indirect** keywords. If an identical command is entered, with exception of either the **next-hop** or **indirect** parameters, then the static route is replaced with the new command, and unless specified, the respective defaults for **preference** and **metric** are applied.

preference *preference* — The preference of this static route (as opposed to the routes from different sources such as BGP or OSPF), expressed as a decimal integer. When modifying the **preference** value of an existing static route, unless specified, the metric will not change.

If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of which route to use is determined by the configuration of the ECMP command.

Default 5

Values 1 — 255

metric *metric* — The cost metric for the static route, expressed as a decimal integer. This value is used when importing this static route into other protocols such as OSPF. This value is also used to determine the static route to install in the forwarding table: When modifying the metrices of an existing static route, unless specified, the preference will not change.

If there are multiple static routes with the same preference but unequal metrices, the lower cost (metric) route is installed. If there are multiple static routes with equal preference and metrics then ECMP rules apply. If there are multiple routes with unequal preferences, then the lower preference route is installed.

Default 1

Values 0 — 65535

tag — Adds a 32-bit integer tag to the static route. The tag is used in route policies to control distribution of the route into other protocols.

Values 1..4294967295

bfd-enable — Associates the state of the static route to a BFD session between the local system and the configured nexthop. This keyword cannot be configured if the nexthop is **indirect** or a **blackhole** keywords are specified.

cpe-check *target-ip-address* **—** This parameter specifies the IP address of the target CPE device. ICMP pings will be sent to this target IP address. This parameter must be configured to enable the CPE connectivity feature for the associated static route. The target-ip-address cannot be in the same subnet as the static route subnet itself to avoid possible circular references. This option is mutually exclusive with BFD support on a given static route.

Default no cpe-check enabled

vrf-export

Syntax **vrf-export** *policy* [*policy...*]
no vrf-export

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the export policies to control routes exported from the local VPN routing/forwarding (VRF) to other VRFs on the same or remote PE routers (via MP-BGP).

The **no** form of the command removes all route policy names from the export list.

Default None — No routes are exported from the VRF by default.

Parameters *policy* — The route policy statement name.

vrf-import

Syntax **vrf-import** *policy* [*policy...*]
no vrf-import

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command sets the import policies to control routes imported to the local VPN routing/forwarding (VRF) from other VRFs on the same or remote PE routers (via MP-BGP). BGP-VPN routes imported with a vrf-import policy will use the BGP preference value of 170 when imported from remote PE routers, or retain the protocol preference value of the exported route when imported from other VRFs on the same router, unless the preference is changed by the policy.

The **no** form of the command removes all route policy names from the import list

Default None — No routes are accepted into the VRF by default.

Parameters *policy* — The route policy statement name.

Router DHCP Configuration Commands

dhcp

Syntax **dhcp**

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure DHCP parameters.

local-dhcp-server

Syntax **local-dhcp-server** *server-name* [**create**]
no local-dhcp-server *server-name*

Context config>service>vprn>dhcp
config>service>vprn>if
config>service>vprn>nw-if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command instantiates a local DHCP server. A local DHCP server can serve multiple interfaces but is limited to the routing context it was which it was created.

NOTE: On 7210 SAS-K 2F4T6C DHCP server is supported only for craft access for local management of the node. It must not be used as a general purpose DHCP server.

Default none

Parameters *server-name* — Specifies the name of local DHCP server.

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

force-renews

Syntax [**no**] **force-renews**

Context config>service>vprn>dhcp>server

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the sending of sending forcerenew messages.

The no form of the command disables the sending of forcerenew messages.

Default no force-renews

pool

Syntax **pool** *pool-name* [**create**]
no pool *pool-name*

Context config>service>vprn>dhcp>local-dhcp-server

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures a DHCP address pool on the router.

Default none

Parameters *pool name* — Specifies the name of this IP address pool. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters.

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

max-lease-time

Syntax **max-lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no max-lease-time

Context config>service>vprn>dhcp>server>pool

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the maximum lease time.

The **no** form of the command returns the value to the default.

Default 10 days

Parameters *time* — Specifies the maximum lease time.

Values	days :	0 — 3650
	hours	0 — 23
	minutes:	0 — 59
	seconds	0 — 59

min-lease-time

Syntax **min-lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no min-lease-time

Context config>service>vprn>dhcp>server>pool

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the minimum lease time.

The **no** form of the command returns the value to the default.

Default 10 minutes

Parameters *time* — Specifies the minimum lease time.

Values

days :	0 — 3650
hours	0 — 23
minutes:	0 — 59
seconds	0 — 5

minimum-free

Syntax **minimum-free** *minimum-free* [**percent**] [**event-when-depleted**]
no minimum-free

Context config>router>dhcp>server>pool

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the desired minimum number of free addresses in this pool.

The **no** form of the command reverts to the default.

Default 1

Parameters *minimum-free* — Specifies the minimum number of free addresses.

0 to 255

percent — Specifies that the value indicates a percentage.

event-when-depleted — This parameter enables a system-generate event when all available addresses in the pool/subnet of local DHCP server are depleted.

nak-non-matching-subnet

Syntax [**no**] **nak-non-matching-subnet**

Context config>router>dhcp>server>pool

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

With this command, if the local DHCPv4 server receives a DHCP request with option 50 (means client try to request a previous allocated message as described in section 3.2 of RFC 2131, *Dynamic Host Configuration Protocol*) and the address allocation algorithm ends up using a pool and the address in option50 is not in pool, then system will return a DHCP NAK, otherwise system just drop the DHCP packet.

Default no nak-non-matching-subnet

offer-time

Syntax **offer-time** [**min** *minutes*] [**sec** *seconds*]
no offer-time

Context config>service>vprn>dhcp>server>pool

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the offer time.

The **no** form of the command returns the value to the default.

Default 1 minute

Parameters *time* — Specifies the offer time.

Values	minutes:	0 — 10
	seconds	0 — 59

options

Syntax **options**

Context config>service>vprn>dhcp>server>pool

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure pool options. The options defined here can be overruled by defining the same option in the local user database.

Default none

custom-option

Syntax **custom-option** *option-number* **address** [*ip-address*...(up to 4 max)]
custom-option *option-number* **hex** *hex-string*
custom-option *option-number* **string** *ascii-string*
no custom-option *option-number*

Context config>service>vprn>dhcp>server>pool>options
config>service>vprn>dhcp>server>pool>subnet>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures specific DHCP options. The options defined here can overrule options in the local user database.

The **no** form of the removes the option from the configuration.

Default none

Parameters *option-number* — specifies the option number that the DHCP server uses to send the identification strings to the DHCP client.

Values 1 — 254

address *ip-address* **—** Specifies the IP address of this host.

hex *hex-string* **—** Specifies the hex value of this option.

Values 0x0..0xFFFFFFFF...(maximum 254 hex nibbles)

string *ascii-string* **—** Specifies the value of this option.

Values Up to 127 characters maximum.

dns-server

Syntax **dns-server** *ip-address* [*ip-address*...(up to 4 max)]

Context config>service>vprn>dhcp>server>pool>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the IP address of the DNS server.

Default none

Parameters *ip-address* — The IP address of the DNS server. 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).

domain-name

Syntax **domain-name** *domain-name*
no domain-name

Context config>service>vprn>dhcp>server>pool>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the default domain for a DHCP client that the router uses to complete unqualified hostnames (without a dotted-decimal domain name).

The **no** form of the command removes the name from the configuration.

Default none

Parameters *domain-name* — Specifies the domain name for the client.

Values Up to 127 characters

lease-rebind-time

Syntax **lease-rebind-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no lease-rebind-time

Context config>service>vprn>dhcp>server>pool>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the time the client transitions to a rebinding state.

The **no** form of the command removes the time from the configuration.

Default none

Parameters *time* — Specifies the lease rebind time.

Values

days:	0 — 3650
hours:	0 — 23
minutes:	0 — 59
seconds	0 — 59

lease-renew-time

Syntax **lease-renew-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no lease-renew-time

Context config>service>vprn>dhcp>server>pool>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the time the client transitions to a renew state.

The **no** form of the command removes the time from the configuration.

Default none

Parameters *time* — Specifies the lease renew time.

Values

days:	0 — 3650
hours:	0 — 23
minutes:	0 — 59
seconds	0 — 59

lease-time

Syntax **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no lease-time

Context config>service>vprn>dhcp>server>pool>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the amount of time that the DHCP server grants to the DHCP client permission to use a particular IP address.

The **no** form of the command removes the lease time parameters from the configuration.

Default none

Parameters *time* — Specifies the lease time.

Values		
	days :	0 — 3650
	hours	0 — 23
	minutes:	0 — 59
	seconds	0 — 59

subnet

Syntax **subnet** {*ip-address*/*mask* | *ip-address netmask*} [**create**]
no subnet {*ip-address*/*mask* | *ip-address netmask*}

Context config>service>vprn>dhcp>server>pool

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command creates a subnet of IP addresses to be served from the pool. The subnet cannot include any addresses that were assigned to subscribers without those addresses specifically excluded. When the subnet is created no IP addresses are made available until a range is defined.

Default none

Parameters *ip-address* — Specifies the base IP address of 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).

mask — The subnet mask in dotted decimal notation. 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.

netmask — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

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

address-range

Syntax [**no**] **address-range** *start-ip-address end-ip-address*

Context config>service>vprn>dhcp>server>pool>subnet

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures a range of IP addresses to be served from the pool. All IP addresses between the start and end IP addresses will be included (other than specific excluded addresses).

Default none

Parameters *start-ip-address* — Specifies the start address of this range to include. 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).

end-ip-address — Specifies the end address of this range to include. 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).

exclude-addresses

Syntax [**no**] **exclude-addresses** *start-ip-address* [*end-ip-address*]

Context config>service>vprn>dhcp>server>pool>subnet

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies a range of IP addresses that excluded from the pool of IP addresses in this subnet.

Default none

Parameters *start-ip-address* — Specifies the start address of this range to exclude. 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).

end-ip-address — Specifies the end address of this range to exclude. 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).

maximum-declined

Syntax **maximum-declined** *maximum-declined*
no maximum-declined

Context config>service>vprn>dhcp>server>pool>subnet

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the maximum number of declined addresses allowed.

Default 64

Parameters *maximum-declined* — Specifies the maximum number of declined addresses allowed.

Values 0 — 4294967295

minimum-free

Syntax **minimum-free** *minimum-free*
no minimum-free

Context config>service>vprn>dhcp>server>pool>subnet

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the minimum number of free addresses in this subnet. If the actual number of free addresses in this subnet falls below this configured minimum, a notification is generated.

Default 1

Parameters *minimum-free* — Specifies the minimum number of free addresses in this subnet.

Values 0 — 255

default-router

Syntax **default-router** *ip-address* [*ip-address*...(up to 4 max)]
no default-router

Context config>service>vprn>dhcp>server>pool>subnet

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the IP address of the default router for a DHCP client. Up to four IP addresses can be specified.

The **no** form of the command removes the address(es) from the configuration.

Default none

Parameters *ip-address* — Specifies the IP address of the default router. 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).

subnet-mask

Syntax **subnet-mask** *ip-address*
no subnet-mask

Context config>service>vprn>dhcp>server>pool>subnet

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the subnet-mask option to the client. The mask can either be defined (for supernetting) or taken from the pool address.

The **no** form of the command removes the address from the configuration.

Default none

Parameters *ip-address* — Specifies the IP address of the subnet mask. 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).

use-gi-address

Syntax [**no**] **use-gi-address**

Context config>service>vprn>dhcp>server

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the use of gi-address matching. If the gi-address flag is enabled, a pool can be used even if a subnets is not found. If the local-user-db-name is not used, the gi-address flag is used and addresses are handed out by GI only. If a user must be blocked from getting an address the server maps to a local user database and configures the user with no address.

A pool can include multiple subnets. Since the GI is shared by multiple subnets in a subscriber-interface the pool may provide IP addresses from any of the subnets included when the GI is matched to any of its subnets. This allows a pool to be created that represents a sub-int.

Default no use-gi-address

user-db

Syntax **user-db** *local-user-db-name*
no user-db

Context config>service>vprn>dhcp>server

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures a local user database for authentication.

Default not enabled

Parameters *local-user-db-name* — Specifies the name of a local user database.

vrf-target

Syntax **vrf-target {***ext-community***|export** *ext-community***|import** *ext-community***}**
no vrf-target

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the route target to be added to the advertised routes or compared against the received routes from other VRFs on the same or remote PE routers. vrf-import or vrf-export policies override the vrf-target policy.

The **no** form of the command removes the vrf-target.

Default no vrf-target

Parameters *ext-comm* — 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-address*:*comm-val* |*2byte-asnumber:ext-comm-val*|*4byte-asnumber:comm-val*}

ip-address:	a.b.c.d
comm-val:	0 — 65535
2byte-asnumber:	1 — 65535
4byte-asnumber	0 — 4294967295

import *ext-community* **—** Specify communities allowed to be accepted from remote PE neighbors.

export *ext-community* **—** Specify communities allowed to be sent to remote PE neighbors.

IPOE Commands

subscriber-mgmt

Syntax **subscriber-mgmt**

Context config

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure subscriber management entities. A subscriber is uniquely identified by a subscriber identification string. Each subscriber can have one DHCP session active at any time. Each session is referred to as a subscriber host and is identified by its IP address and MAC address.

NOTE: On 7210 SAS-K 2F4T6C, the 7x50 subscriber management CLI is inherited to allow users to configure DHCP address for specific hosts when using DHCP server for local craft access. 7210 SAS platforms do not support subscriber management. In all the commands under this section, in the context of 7210 usage, DHCP subscriber refers to a IP host requesting a IP address.

local-user-db

Syntax **local-user-db** *local-user-db-name* **[create]**

Context config>subscr-mgmt

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure a local user database.

Default not enabled

Parameters *local-user-db-name* — Specifies the name of a local user database.

ipoe

Syntax **ipoe**

Context config>subscr-mgmt>loc-user-db

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures IPoE host parameters.

mask

Syntax **mask type** *dhcp-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}
no mask type *dhcp-match-type*

Context config>subscr-mgmt>loc-user-db>ipoe

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the mask.

Parameters *dhcp-match-type* — Specifies up to four matching types to identify a host.

Values DHCP: circuit-id, option60, remote-id, sap-id, string, system-id
PPP: circuit-id, remote-id, service-name, username

Values prefix-string *prefix-string*

Specifies a substring that is stripped of the start of the incoming circuit ID before it is matched against the value configured in the DHCP or PPPOE circuit ID.

This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum, *' is wildcard.

prefix-length *prefix-length* **—** Specifies the number of characters to remove from the start of the incoming circuit ID before it is matched against the value configured in the DHCP circuit ID.

Values 1to 127

suffix-string *suffix-string* **—** Specifies a substring that is stripped of the end of the incoming circuit ID before it is matched against the value configured in DHCP circuit ID.

This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum

suffix-length *suffix-length* **—** Specifies the number of characters to remove from the end of the incoming circuit ID before it is matched against the value configured in the DHCP circuit ID.

Values 1to 127

host

Syntax **host** *host-name* [**create**]
no host *host-name*

Context config>subscr-mgmt>loc-user-db>ipoe

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command defines a DHCP host /subscriber.

Parameters *host-name* — Specifies the host name.

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

address

Syntax **address gi-address** [**scope** *scope*]
address *ip-address*[*/prefix-length*]
address pool *pool-name* [**secondary-pool** *sec-pool-name*] [**delimiter** *delimiter*]
address use-pool-from-client [**delimiter** *delimiter*]
no address

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures how the IP address is defined for this host.

When the user-db is used from a local-dhcp-server, then this command defines how to define the IP address the server will "offer" to the DHCP-client.

When the user-db is used for PPPoE authentication, the **gi-address** parameter cannot be used. A fixed IP address will then cause PPPoE to use this IP address. If no IP address is specified, the PPPoE will look for IP address by other means (DHCP). If a pool name is given, this pool will be sent in the DHCP request so that it can be used in by the DHCP server to determine which address to give to the host.

The **no** form of the command causes no IP address to be assigned to this host. In a user-db referred to from a local-dhcp-server, creating a host without address information will cause the matching client never to get an IP address.

Default no address

Parameters **gi-address —** When specified, the gi-address of the DHCP message is taken to look for a subnet in the local DHCP server. The first available free address of the subnet is taken and "offered" to the host. When **local-user-db** is used for PPPoE authentication, this has the same result as **no address**.

ip-address — Specifies the fixed IP address to use for this host.

pool-name/sec-pool-name — Specifies the primary (and secondary) pool (in the local DHCP server) to use to look for an available address. The first available IP address from any subnet in the pool will be used. When the local user database is used for PPPoE authentication, this causes the specified pool name to be sent to the DHCP server in a vendor-specific sub-option under Option 82

use-pool-from-client — Use the pool-name in the Option 82 vendor-specific sub-option.

delimiter *delimiter* **—** A single ascii character specifies the delimiter of separating primary and secondary pool names in option82 VSO.

gi-address

Syntax **gi-address** *ip-address*
no gi-address

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description This command allows selection of gi-addresses based on the host entry in LUDB.

The gi-address must be a valid address (associated with an interface) within the routing context that received the DHCP message on the access side.

Default no gi-address

Parameters *ip-address* — Specifies the IPv4 gi-address.

auth-domain-name

Syntax **auth-domain-name** *domain-name*
no auth-domain-name

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command sets the domain name which can be appended to user-name in RADIUS-authentication-request message for the given host.

Parameters *domain-name* — Specifies the domain name to be appended to user-name in RADIUS-authentication-request message for the given host.

host-identification

Syntax **host-identification**

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure host identification parameters.

server

Syntax **server** *ip-address*
no server

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the IP address of the DHCP server to relay to.

The **no** form of the command removes the DHCP server IP address from the configuration.

The configured DHCP server IP address must reference one of the addresses configured under the DHCP CLI context of an IES/VPRN subscriber.

Default no server

Parameters *ip-address* — Specifies the IP address of the DHCP server.

circuit-id

Syntax **circuit-id string** *ascii-string*
circuit-id hex *hex-string*
no circuit-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the circuit-id to match.

Parameters *ascii-string* — specifies the circuit ID from the Option 82.

hex-string — Specifies the circuit ID in hexadecimal format from the Option 82.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

encap-tag-range

Syntax **encap-tag-range start-tag** *start-tag* **end-tag** *end-tag*
no encap-tag-range

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies a range of encapsulation tag as the host identifications. The encapsulation tag is dot1q or qinq on Ethernet port.

For dot1q, the start/end-tag is single number, range from 0-4094; for QinQ, the start/end-tag format is x.y, x or y could be "*", which means ignore inner or outer tag;

Note: This command will only be used when **encap-tag-range** is configured as one of the match-list parameters.

The **no** form of the command removes the values from the configuration.

Default none

Parameters **start-tag** *start-tag* — Specifies the value of the start label in the range of SAP's allowed on this host.

Values

start-tag	dot1q	qtag1	
	qinq	(qtag1.qtag2 | qtag1.* | *.qtag2)	
		qtag1	[0 to 4094]
		qtag2	[0 to 4094]
		vpi	[0 to 4095] (NNI)
			[0 to 255] (UNI)
		vci	[1 to 65535]

end-tag *end-tag* — Specifies the value of the end label in the range of SAP's allowed on this host.

Values

end-tag	dot1q	qtag1	
	qinq	(qtag1.qtag2 | qtag1.* | *.qtag2)	
		qtag1	[0 to 4094]
		qtag2	[0 to 4094]
		vpi	[0 to 4095] (NNI)
			[0 to 255] (UNI)
		vci	[1 to 65535]

mac

Syntax **mac** *ieee-address*
no mac

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the MAC address to match.

Parameters *ieee-address* — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

option60

Syntax **option60** *hex-string*
no option60

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the Vendor-Identifying Vendor Option to match. Option 60 is encoded as Type-Length-Value (TLV). The *hex-string* portion of Option 60 in the received DHCP request is used for matching. Only the first 32 bytes can be defined here. If Option 60 from the message is longer, those bytes are ignored.

Default no option60

Parameters *hex-string* — Specifies the hex value of this option.

Values 0x0 to 0xFFFFFFFF(maximum 254 hex nibbles)

remote-id

Syntax **remote-id** *hex-string*
remote-id *ascii-string*
no remote-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the remote id of this host.

The **no** form of this command returns the system to the default.

Default no remote-id

Parameters *remote-id* — Specifies the remote-id.

sap-id

Syntax **sap-id** *sap-id*
no sap-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the SAP ID from the Nokia vendor-specific sub-option in Option 82 to match.

Parameters *sap-id* — Specifies a SAP ID, up to 255 characters maximum.

service-id

Syntax **service-id** *service-id*
no service-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies an existing service ID from the Nokia vendor-specific sub-option in Option 82 to match.

Parameters *service-id* — Specifies an existing service ID.

Values 1 to 2147483647

string

Syntax **string** *string*
no string

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the string from the Nokia vendor-specific sub-option in Option 82 to match.

Parameters *string* — Specifies the string, up to 255 characters maximum.

system-id

Syntax **system-id** *system-id*
no system-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the system ID from the Nokia vendor-specific sub-option in Option 82 to match.

Parameters *system-id* — Specifies the system ID, up to 255 characters maximum.

options

Syntax **options**

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure pool options. The options defined here can be overruled by defining the same option in the local user database.

Default none

custom-option

Syntax **custom-option** *option-number* **address** [*ip-address*...(up to 4 max)]
custom-option *option-number* **hex** *hex-string*
custom-option *option-number* **string** *ascii-string*
no custom-option *option-number*

Context config>subscr-mgmt>loc-user-db>ipoe>host>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures specific DHCP options. The options defined here can overrule options in the local user database.

The **no** form of the removes the option from the configuration.

Default none

Parameters *option-number* — specifies the option number that the DHCP server uses to send the identification strings to the DHCP client.

Values 1 to 254

address *ip-address* **—** Specifies the IP address of this host.

hex *hex-string* **—** Specifies the hex value of this option.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

string *ascii-string* **—** Specifies the value of this option.

Values Up to 127 characters maximum.

default-router

Syntax **default-router** *ip-address* [*ip-address*...(up to 4 max)]
no default-router

Context config>subscr-mgmt>loc-user-db>ipoe>host>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the IP address of the default router for a DHCP client. Up to four IP addresses can be specified.

The **no** form of the command removes the address(es) from the configuration.

Default none

Parameters *ip-address* — Specifies the IP address of the default router. 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).

dns-server

Syntax **dns-server address** [*ip-address*...(up to 4 max)]
no dns-server

Context config>subscr-mgmt>loc-user-db>ipoe>host>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the IP address of the DNS server.

Default none

Parameters *ipv6-address* — Specifies the IPv4 address of the DNS server. 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).

domain-name

Syntax **domain-name** *domain-name*
no domain-name

Context config>subscr-mgmt>loc-user-db>ipoe>host>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the default domain for a DHCP client that the router uses to complete unqualified host names (without a dotted-decimal domain name).

The **no** form of the command removes the name from the configuration.

Default none

Parameters *domain-name* — Specifies the domain name for the client.

Values Up to 127 characters

lease-rebind-time

Syntax **lease-rebind-time [days** *days***] [hrs** *hours***] [min** *minutes***] [sec** *seconds***]**
no lease-rebind-time

Context config>subscr-mgmt>loc-user-db>ipoe>host>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the time the client transitions to a rebinding state.

The **no** form of the command removes the time from the configuration.

Default none

Parameters *time* — Specifies the lease rebind time.

Values

days *days*	0 to 3650
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 59

lease-renew-time

Syntax **lease-renew-time [days** *days***] [hrs** *hours***] [min** *minutes***] [sec** *seconds***]**
no lease-renew-time

Context config>subscr-mgmt>loc-user-db>ipoe>host>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the time the client transitions to a renew state.

The **no** form of the command removes the time from the configuration.

Default none

Parameters *time* — Specifies the lease renew time.

Values

days:	0 to 3650
hours:	0 to 23
minutes:	0 to 59
seconds	0 to 59

lease-time

Syntax **lease-time [days** *days***] [hrs** *hours***] [min** *minutes***] [sec** *seconds***]**
no lease-time

Context config>subscr-mgmt>loc-user-db>ipoe>host>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the amount of time that the DHCP server grants to the DHCP client permission to use a particular IP address.

The **no** form of the command removes the lease time parameters from the configuration.

Default none

Parameters *time* — Specifies the lease time.

Values

days *days*	0 to 3650
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 59

subnet-mask

Syntax **subnet-mask** *ip-address*
no subnet-mask

Context config>subscr-mgmt>loc-user-db>ipoe>host>options

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the subnet-mask option to the client. The mask can either be defined (for supernetting) or taken from the pool address.

The **no** form of the command removes the address from the configuration.

Default none

Parameters *ip-address* — Specifies the IP address of the subnet mask. 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).

mask

Syntax **mask type** *dhcp-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}
no mask type *dhcp-match-type*

Context config>subscr-mgmt>loc-user-db>ipoe

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the mask.

Parameters *dhcp-match-type* — The data type represents the type of matching done to identify a DHCP host.

Values circuit-id, option60, remote-id, sap-id, string, system-id

prefix-string *prefix-string* **—** Specifies a substring that is stripped of the start of the incoming circuit ID before it is matched against the value configured in the DHCP circuit ID.

This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum, "*"

prefix-length *prefix-length* **—** Specifies the number of characters to remove from the start of the incoming circuit ID before it is matched against the value configured in the circuit ID.

Values 1 to 127

suffix-string *suffix-string* **—** Specifies a substring that is stripped of the end of the incoming circuit ID before it is matched against the value configured in circuit ID.

This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum

suffix-length *suffix-length* **—** Specifies the number of characters to remove from the end of the incoming circuit ID before it is matched against the value configured in the circuit ID.

Values 1 to 127

match-list

Syntax **match-list** *match-type-1* [*match-type-2*...(up to 4 max)]
no match-list

Context config>subscr-mgmt>loc-user-db>ipoe

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the type of matching done to identify a host.

Default no match-list

Parameters *match-type-x* — Specifies up to four matching types to identify a host.

Values circuit-id, mac, option60, remote-id, sap-id, service-id, string, system-id

circuit-id — Specifies the DHCP4 option (82,1) or DHCP6 option 18.

mac — Specifies the MAC address of the client. Chaddr in DHCP4 and DUID in IPv6.

option60 — Specifies the DHCP4 option 60.

remote-id — Specifies the DHCP4 option (82,2) or DHCP6 option 37

sap-id — Specifies the SAP ID on which DHCPv4 packet are received. The sap-id is inserted as ALU VSO (82,9,4) by the DHCPv4 relay in router. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. Since the dhcp-relay configuration is enabled under the group-interface CLI hierarchy, the group-interface and the service-id must be known before the sap-id can be used for LUDB match.

service-id — Specifies the service-id of the ingress SAP for DHCPv4 packets. The service-id is inserted as ALU VSO (82,9,3) by the DHCPv4 relay in router. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay.

string — Specifies the custom string configured under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. The string is inserted as ALU VSO (82,9,5) by the DHCPv4 relay in router. Since the dhcp-relay configuration is enabled under the group-interface CLI hierarchy, the group-interface and the service-id must be known before the string can be used for LUDB match.

system-id — Specifies the system-id of the node name configured under the **system>name** CLI hierarchy. The system-id is inserted as ALU VSO (82,9,1) by the DHCPv4 relay in router. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. Since the dhcp-relay configuration is enabled under the group-interface CLI hierarchy, the group-interface and the service-id must be known before the system-id can be used for LUDB match.

shutdown

Syntax [**no**] **shutdown**

Context config>subscr-mgmt>loc-user-db
config>subscr-mgmt>loc-user-db>ipoe>host

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command places the entity into an administratively enabled state.

SDP Commands

spoke-sdp

Syntax [**no**] **spoke-sdp** *sdp-id*

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command binds a service to an existing Service Distribution Point (SDP).

The SDP has an operational state that determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service will be down.

The SDP must already be defined in the config>service>sdp context in order to associate an SDP with a VPRN service. If the sdp sdp-id is not already configured, an error message is generated. If the sdp-id exists, 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 routers can participate in the service. Alternatively, auto-bind can be used. With auto-bind, no vprn>spoke-sdp configuration is required. When both auto-bind and spoke-sdp are configured, spoke-sdp takes precedence. Spoke-sdp must be deconfigured for auto-bind to take effect.

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 is affected. Once the SDP is removed, no packets are forwarded to the far-end router.

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

Special Cases **VPRN —** Several SDPs can be bound to a VPRN service. Each SDP must be destined to a different 7210 SAS router. If two *sdp-id* bindings terminate on the same 7210 SAS, an error occurs and the second SDP binding is rejected.

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

Values 1 to 17407

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

Interface Commands

interface

Syntax **interface** *ip-int-name*
no interface *ip-int-name*

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command creates a logical IP routing interface for a Virtual Private Routed Network (VPRN). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within VPRN service IDs. The **interface** command can be executed in the context of an VPRN service ID. The IP interface created is associated with the service core network routing instance and default routing table. The typical use for IP interfaces created in this manner is for subscriber internet access.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config router interface** and **config service vprn interface**. Interface names must not be in the dotted decimal notation of an IP address. 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. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

The available IP address space for local subnets and routes is controlled with the **config router service-prefix** command. The **service-prefix** command administers the allowed subnets that can be defined on service IP interfaces. It also controls the prefixes that may be learned or statically defined with the service IP interface as the egress interface. This allows segmenting the IP address space into **config router** and **config service** domains.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All VPRN IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The **no** form of this command removes IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For VPRN services, the IP interface must be shutdown before the SAP on that interface may be removed. VPRN services do not have the **shutdown** command in the SAP CLI context. VPRN service SAPs rely on the interface status to enable and disable them.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and config service vprn interface commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

If *ip-int-name* already exists within the service ID, the context will be changed to maintain that IP interface. If *ip-int-name* already exists within another service ID or is an IP interface defined within the **config router** commands, an error will occur and context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

address

Syntax **address** *ip-address/mask* | *ip-address netmask*} [**broadcast** [**all-ones** | **host-ones**]
no address

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Assigns an IP address, IP subnet, and broadcast address format to a VPRN IP router interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each VPRN 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 7210 SAS.

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 adminstratively 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 addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

/ — The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "/" and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.

mask-length — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values 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.

broadcast — The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indictates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) will be received by the IP interface.

Default host-ones

all-ones — The **all-ones** keyword following the **broadcast** parameter specifies the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

host-ones — The **host-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

allow-directed-broadcasts

Syntax [**no**] **allow-directed-broadcasts**

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command controls the forwarding of directed broadcasts out of the IP interface.

A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The **allow-directed-broadcasts** command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface.

When enabled, a frame destined to the local subnet on this IP interface will be sent as a subnet broadcast out this interface. Care should be exercised when allowing directed broadcasts as it is a well-known mechanism used for denial-of-service attacks.

When disabled, directed broadcast packets discarded at this egress IP interface will be counted in the normal discard counters for the egress SAP.

By default, directed broadcasts are not allowed and will be discarded at this egress IP interface.

The **no** form of this command disables the forwarding of directed broadcasts out of the IP interface.

Default **no allow-directed-broadcasts** — Directed broadcasts are dropped.

ip-mtu

Syntax **ip-mtu octets**
no ip-mtu

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the IP maximum transmit unit (packet) for this interface.

The no form of the command returns the default value.

Default no ip-mtu

loopback

Syntax [**no**] **loopback**

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated interface cannot be bound to a SAP.

When using mtrace/mstat in a Layer 3 VPN context then the configuration for the VPRN should have a loopback address configured which has the same address as the core instance's system address (BGP next-hop).

Default None

local-proxy-arp

Syntax **[no] local-proxy-arp**

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus will become the forwarding point for all traffic between hosts in that subnet. When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

Default no local-proxy-arp

mac

Syntax **[no] mac ieee-mac-address**

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command assigns a specific MAC address to a VPRN IP interface.

The no form of this command returns the MAC address of the IP interface to the default value.

Default The physical MAC address associated with the Ethernet interface that the SAP is configured on.

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.

proxy-arp

Syntax **[no] proxy-arp**

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables proxy ARP on the interface.

Default no proxy-arp

proxy-arp-policy

Syntax **[no] proxy-arp-policy** *policy-name [policy-name...(up to 5 max)]*

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables a proxy ARP policy for the interface.

The no form of this command disables the proxy ARP capability.

Default no proxy-arp

Parameters *policy-name* — The export route 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.

remote-proxy-arp

Syntax **[no] remote-proxy-arp**

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables remote proxy ARP on the interface.

Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.

static-arp

Syntax [**no**] **static-arp** *ip-address ieee-mac-address*

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP will appear in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface. If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of this command removes a static ARP entry.

Default none

Parameters *ip-address* — Specifies the IP address for the static ARP in IP address dotted decimal notation.

ieee-mac-address — Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa*-*bb*-*cc*-*dd*-*ee*-*ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

DHCP Server Commands

dhcp

Syntax **dhcp**

Context config>service>vprn>if

Platforms Supported: 7210 SAS-K 2F4T6C.

This command enables the context to configure DHCP parameters.

action

Syntax **action** {**replace** | **drop** | **keep**}
no action

Context config>service>vprn>if>dhcp>option

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures the processing required when the SR-Series router receives a DHCP request that already has a Relay Agent Information Option (Option 82) field in the packet.

The **no** form of this command returns the system to the default value.

Default Per RFC 3046, *DHCP Relay Agent Information Option* , section 2.1.1, *Reforwarded DHCP requests*, the default is to keep the existing information intact. The exception to this is if the giaddr of the received packet is the same as the ingress address on the router. In that case the packet is dropped and an error is logged.

Parameters **replace —** In the upstream direction (from the user), the existing Option 82 field is replaced with the Option 82 field from the router. In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop — The packet is dropped, and an error is logged.

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.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

circuit-id

Syntax **circuit-id** [**ascii-tuple** | **ifindex** | **sap-id** | **vlan-ascii-tuple**]
no circuit-id

Context config>service>vprn>if>dhcp>option

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

When enabled, the router sends the interface index (If Index) in the **circuit-id** sub-option of the DHCP packet. The If Index of a router interface can be displayed using the command **show>router>interface>detail**. This option specifies data that must be unique to the router that is relaying the circuit.

If disabled, the **circuit-id** sub-option of the DHCP packet will be left empty.

The **no** form of this command returns the system to the default.

Default circuit-id

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 "|".

ifindex — Specifies that the interface index will be used. The If Index of a router interface can be displayed using the command **show>router>interface>detail**.

sap-id — Specifies that the SAP ID will be 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.

gi-address

Syntax **gi-address** *ip-address* [*src-ip-addr*]
no gi-address

Context config>service>vprn>if>dhcp

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures the gateway interface address for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. The GI address is needed, when the router functions as a DHCP relay, to distinguish between the different subscriber interfaces and potentially between the group interfaces defined.

Default no gi-address

Parameters *ip-address* — Specifies the host IP address to be used for DHCP relay packets.

src-ip-address — Specifies the source IP address to be used for DHCP relay packets.

option

Syntax [**no**] **option**

Context config>service>vprn>if>dhcp

Platforms Supported: 7210 SAS-K 2F4T6C.

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

remote-id

Syntax **remote-id** [**mac** | **string** *string*]
no remote-id

Context config>service>vprn>sub-if>grp-if>dhcp>option
config>service>vprn>nw-if>dhcp>option

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

When enabled, the router sends the MAC address of the remote end (typically the DHCP client) in the **remote-id** sub-option of the DHCP packet. This command identifies the host at the other end of the circuit. If disabled, the **remote-id** sub-option of the DHCP packet will be left empty.

The **no** form of this command returns the system to the default.

Default remote-id

Parameters **mac —** This keyword specifies the MAC address of the remote end is encoded in the sub-option.

string *string* **—** Specifies the remote-id.

vendor-specific-option

Syntax [**no**] **vendor-specific-option**

Context config>service>vprn>if>dhcp>option

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

client-mac-address

Syntax [**no**] **client-mac-address**
config>service>vprn>if>dhcp>option>vendor

Platforms Supported: 7210 SAS-K 2F4T6C.

This command enables the sending of the MAC address in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

sap-id

Syntax [**no**] **sap-id**

Context config>service>vprn>if>dhcp>option>vendor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the sending of the SAP ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

service-id

Syntax [**no**] **service-id**

Context config>service>vprn>if>dhcp>option>vendor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the sending of the service ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the Alcatel-Lucent vendor specific sub-option of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

Context config>service>vprn>if>dhcp>option>vendor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the vendor specific sub-option string of the DHCP relay packet.

The **no** form of the command returns the default value.

Parameters *text* — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

system-id

Syntax [**no**] **system-id**

Context config>service>vprn>if>dhcp>option>vendor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies whether the system-id is encoded in the Alcatel-Lucent vendor specific sub-option of Option 82.

Default None

server

Syntax **server** *server1* [*server2*...(up to 8 max)]

Context config>service>vprn>if>dhcp

Platforms Supported: 7210 SAS-K 2F4T6C.

This command specifies a list of servers where requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list. There can be a maximum of 8 DHCP servers configured.

The flood command is applicable only in the VPLS case. There is a scenario with VPLS where the VPLS node only wants to add Option 82 information to the DHCP request to provider per-subscriber information, but it does not do full DHCP relay. In this case, the server is set to "flood". This means the DHCP request is still a broadcast and is sent through the VPLS domain. A node running at L3 further upstream then can perform the full L3 DHCP relay function.

Default no server

Parameters *server* — Specifies the DHCP server IP address.

trusted

Syntax [**no**] **trusted**

Context config>service>vprn>if>dhcp

Platforms Supported: 7210 SAS-K 2F4T6C.

According to RFC 3046, *DHCP Relay Agent Information Option*, a DHCP request where the giaddr is 0.0.0.0 and which contains a Option 82 field in the packet, should be discarded, unless it arrives on a "trusted" circuit.

If trusted mode is enabled on an IP interface, the relay agent (the SR-Series) will modify the request's giaddr to be equal to the ingress interface and forward the request.

Note that this behavior only applies when the action in the Relay Agent Information Option is "keep". In the case where the Option 82 field is being replaced by the relay agent (action = "replace"), the

original Option 82 information is lost anyway, and there is thus no reason for enabling the trusted option.

The **no** form of this command returns the system to the default.

Default not enabled

Interface ICMP Commands

icmp

Syntax **icmp**

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures Internet Control Message Protocol (ICMP) parameters on a VPRN service.

mask-reply

Syntax [**no**] **mask-reply**

Context config>service>vprn>if>icmp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance will reply to mask requests.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

Default **mask-reply** — Reply to ICMP mask requests.

redirects

Syntax **redirects** [*number seconds*]
no redirects

Context config>service>vprn>if>icmp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This commad configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.

When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

The **redirects** command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default **redirects 100 10** — Maximum of 100 redirect messages in 10 seconds.

Parameters *number* — The maximum number of ICMP redirect messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *seconds* of ICMP redirect messages that can be issued.

Values 1 — 60

ttl-expired

Syntax **ttl-expired** *number seconds*
no ttl-expired

Context config>service>vprn>if>icmp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default ttl-expired 100 10

Parameters *number* — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 — 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>vprn>if>icmp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables and configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The **unreachables** command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.

By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 10 per 10 second time interval.

The **no** form of this command disables the generation of icmp destination unreachable messages on the router interface.

Default **unreachables 100 10**

Parameters *number* — The maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 — 1000

seconds — The time frame in seconds used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 — 60

Interface SAP Commands

sap

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

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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 7210 SAS. 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.

Default No SAPs are defined.

Special Cases **VPRN —** A VPRN SAP must be defined on an Ethernet interface.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 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 2/3 specifies port 3 on MDA 2 in slot .

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

create — Keyword used to create a SAP instance.

tod-suite

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

Context config>service>vprn>if>sap

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command applies a time-based policy (filter or QoS policy) to the SAP. The suite name must already exist in the **config>cron** context.

Default no tod-suite

Parameters *tod-suite-name* — Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP or a subscriber. The suite can be applied to more than one SAP.

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>service>vprn>if>sap

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command creates the accounting policy context that can be applied to an interface SAP or interface SAP spoke SDP.

An accounting policy must be defined before it can be associated with a SAP.
If the *policy-id* does not exist, an error message is generated.

A maximum of one accounting policy can be associated with a SAP at one time. Accounting policies are configured in the **config>log** context.

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

Default Default accounting policy.

Parameters *acct-policy-id* — Enter the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 — 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>vprn>if>sap

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables accounting and statistical data collection for either an interface SAP or interface SAP spoke SDP, or network port. 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

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

The **no** form of this command restores **arp-timeout** to the default value.

Default 14400 seconds

Parameters *seconds* — The minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 — 65535

delayed-enable

Syntax **delayed-enable** *seconds* [**init-only**]
no delayed-enable

Context config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command delays making interface operational by the specified number of seconds.

In environments with many subscribers, it can take time to synchronize the subscriber state between peers when the subscriber-interface is enabled (perhaps, after a reboot). To ensure that the state has time to be synchronized, the **delayed-enable** timer can be specified. The optional parameter **init-only** can be added to use this timer only after a reboot.

Default no delayed-enable

Parameters *seconds* — Specifies the number of seconds to delay before the interface is operational.

Values 1 — 1200

init-only — Delays the initialization of the subscriber-interface to give the rest of the system time to complete necessary tasks such as allowing routing protocols to converge and/or to allow MCS to sync the subscriber information. The delay only occurs immediately after a reboot.

Default

vpls

Syntax vpls *service-name*

Context config>service
config>service>vprn>if

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

The vpls command, within the IP interface context, is used to bind the IP interface to the specified service name.

The system does not attempt to resolve the service name provided until the IP interface is placed into the administratively up state (no shutdown). Once the IP interface is administratively up, the system scans the available VPLS services that have the allow-ip-int-binding flag set for a VPLS service associated with the name. If the service name is bound to the service name when the IP interface is already in the administratively up state, the system will immediately attempt to resolve the given name.

If a VPLS service is found associated with the name and with the allow-ip-int-binding flag set, the IP interface will be attached to the VPLS service allowing routing to and from the service virtual ports once the IP interface is operational.

A VPLS service associated with the specified name that does not have the allow-ip-int-binding flag set or a non-VPLS service associated with the name will be ignored and will not be attached to the IP interface.

If the service name is applied to a VPLS service after the service name is bound to an IP interface and the VPLS service allow-ip-int-binding flag is set at the time the name is applied, the VPLS service is automatically resolved to the IP interface if the interface is administratively up or when the interface is placed in the administratively up state.

If the service name is applied to a VPLS service without the allow-ip-int-binding flag set, the system does not attempt to resolve the applied service name to an existing IP interface bound to the name. To rectify this condition, the flag must first be set and then the IP interface must enter or reenter the administratively up state.

While the specified service name may be assigned to only one service context in the system, it is possible to bind the same service name to more than one IP interface. If two or more IP interfaces are bound to the same service name, the first IP interface to enter the administratively up state (if currently administratively down) or to reenter the administratively up state (if currently administratively up) when a VPLS service is configured with the name and has the allow-ip-

intbinding flag set will be attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

Once an IP interface is attached to a VPLS service, the name associated with the service cannot be removed or changed until the IP interface name binding is removed. Also, the allow-ip-int-binding flag cannot be removed until the attached IP interface is unbound from the service name. Unbinding the service name from the IP interface causes the IP interface to detach from the VPLS service context. The IP interface may then be bound to another service name or a SAP or SDP binding may be created for the interface using the sap or spoke-sdp commands on the interface.

Default none

Parameters *service-name* — The service-name parameter is required when using the IP interface VPLS command and specifies the service name that the system will attempt to resolve to an allow-ip-int-binding enabled VPLS service associated with the name. The specified name is expressed as an ASCII string comprised of up to 32 characters. It does not need to already be associated with a service and the system does not check to ensure that multiple IP interfaces are not bound to the same name.

ingress

Syntax ingress

Context config>service>vprn>if>vpls

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

The ingress node in this context under the VPLS binding is used to define the routed ip-filter-id optional filter overrides.

v4-routed-override-filter

Syntax v4-routed-override-filter *ip-filter-id*
no v4-routed-override-filter

Context config>service>ies>if>vprn>ingress

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

The v4-routed-override-filter command is used to specify an IP filter ID that is applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IP filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and when it is not defined or it is removed, the IP routed packets uses the any existing ingress IP filter on the VPLS virtual port.

The no form of the command is used to remove the IP routed override filter from the ingress IP interface. When removed, the IP ingress routed packets within a VPLS service attached to the IP interface uses the IP ingress filter applied to the packets virtual port when defined.

Default none

Parameters *ip-filter-id* — Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the **configure>filter>ip-filter** context.

Values 1 — 65535

egress

Syntax **egress**

Context config>service>vprn>if>sap

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure egress SAP Quality of Service (QoS) policies and filter policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

ingress

Syntax **ingress**

Context config>service>vprn>if>sap

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

agg-shaper-rate

Syntax **agg-shaper-rate** *agg-rate*
no agg-shaper-rate

Context config>service>vprn>sap>ingress

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command allows user to specify the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues. User can specify the CIR rate and the PIR rate. User must not oversubscribe the total bandwidth available for use by ingress queues.
The no form of the command disables the use of SAP aggregate rate shaper. In other words, the SAP can use up the maximum bandwidth available.

Default no agg-shaper-rate

Parameters ***agg-rate*** **—** Specifies the rate in kilobits per second.

Values 50 — 3000000 | max

Default max

agg-shaper-rate

Syntax **agg-shaper-rate** *agg-rate*
no agg-shaper-rate

Context config>service>vprn>sap>egress

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.
This command allows user to specify the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues. User can specify the CIR rate and the PIR rate. User must not oversubscribe the total bandwidth available for use by ingress queues.
The no form of the command disables the use of SAP aggregate rate shaper. In other words, the SAP can use up the maximum bandwidth available.

Default no agg-shaper-rate

Parameters ***agg-rate*** **—** Specifies the rate in kilobits per second.

Values 50 — 1000000 | max

Default max

filter

Syntax **filter ip** *ip-filter-id*
filter [mac *mac-filter-id***]**
no filter [*ip ip-filter-id***]**
no filter [mac *mac-filter-id***]**
no filter

Context config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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.

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

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local.

Parameters **ip** *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 — 65535

mac *mac-filter-id* — Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1— 65535

Context

Description

qos

Syntax **qos** *policy-id*
no qos

Context config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP) or IP interface. 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 or IP interface ingress and egress policies on SAP or IP interface 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 or IP interface 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 an ingress IP interface that is bound to a VPRN, the policy becomes associated with every SAP on the VPRN and augments the QoS policy that is defined on each SAP. Packets that are bridged will be processed using the policy defined on the VPRN SAP; packets that are routed will be processed using the policy defined in the IES IP interface-binding context.

When an egress QoS policy is associated with an IP interface that has been bound to a VPRN, the policy becomes associated with every SAP on the VPRN 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 VPRN SAP; packets that are routed will be processed using the policy defined in the IP interface-binding context.

By default, no specific QoS policy is associated with the SAP or IP interface 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 or IP interface, and the QoS policy reverts to the default.

Parameters *policy-id* — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.

Values 1 — 65535

BGP Commands

bgp

Syntax [**no**] **bgp**

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the BGP protocol with the VPRN service.

The **no** form of the command disables the BGP protocol from the given VPRN service.

Default no bgp

advertise-inactive

Syntax [**no**] **advertise-inactive**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables or disables the advertising of inactive BGP routers to other BGP peers.

By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination.

Default no advertise-inactive

aggregator-id-zero

Syntax [**no**] **aggregator-id-zero**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command is used to set the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths.

When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute.

When this command is enabled, BGP adds the router ID to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level, while this command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of the command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute.

The **no** form of the command used at the group level reverts to the value defined at the group level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default no aggregator-id-zero — BGP adds the AS number and router ID to the aggregator path attribute.

best-path-selection

Syntax **best-path-selection**

Context config>service>vprn>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables path selection configuration.

always-compare-med

Syntax **always-compare-med** {**zero** | **infinity**}
no always-compare-med

Context config>service>vprn>bgp>best-path-selection

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies how the Multi-Exit Discriminator (MED) path attribute is used in the BGP route selection process. The MED attribute is always used in the route selection process regardless of the peer AS that advertised the route. This parameter determines what MED value is inserted in the RIB-IN. If this parameter is not configured, only the MEDs of routes that have the same peer ASs are compared.

The **no** form of the command removes the parameter from the configuration.

Default no always-compare-med — Only compare MEDs of routes that have the same peer AS.

Parameters **zero —** Specifies that for routes learned without a MED attribute that a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.

infinity — Specifies for routes learned without a MED attribute that a value of infinity (4294967295) is used in the MED comparison. This in effect makes these routes the least desirable.

as-path-ignore

Syntax [**no**] **as-path-ignore**

Context config>service>vprn>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command determines whether the AS path is used to determine the best BGP route.

If this option is present, the AS paths of incoming routes are not used in the route selection process.

The **no** form of the command removes the parameter from the configuration.

Default no as-path-ignore

ignore-nh-metric

Syntax **ignore-nh-metric**
no ignore-nh-metric

Context config>service>vprn
config>service>vprn>bgp>best-path-selection

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command instructs BGP to disregard the resolved distance to the BGP next-hop in its decision process for selecting the best route to a destination. When configured in the config>router>bgp>bestpath- selection context, this command applies to the comparison of two BGP routes with the same NLRI learned from base router BGP peers. When configured in the config>service>vprn context, this command applies to the comparison of two BGP-VPN routes for the same IP prefix imported into the VPRN from the base router BGP instance. When configured in the config>service>vprn>bgp>best-path-selection context, this command applies to the comparison of two BGP routes for the same IP prefix learned from VPRN BGP peers.

The no form of the command (no ignore-nh-metric) restores the default behavior whereby BGP factors distance to the next-hop into its decision process.

Default no ignore-nh-metric

ignore-router-id

Syntax **ignore-router-id**
no ignore-router-id

Context config>service>vprn>bgp>best-path-selection

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

When the ignore-router-id command is present and the current best path to a destination was learned from EBGP peer X with BGP identifier x and a new path is received from EBGP peer Y with BGP identifier y the best path remains unchanged if the new path is equivalent to the current best path up to

the BGP identifier comparison – even if y is less than x. The no form of the command restores the default behavior of selecting the route with the lowest BGP identifier (y) as best.

Default no ignore-router-id

as-override

Syntax **[no] as-override**

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command replaces all instances of the peer's AS number with the local AS number in a BGP route's AS_PATH.

This command breaks BGP's loop detection mechanism. It should be used carefully.

Default as-override is not enabled by default.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the BGP authentication key.

Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD-5 message-based digest. The authentication key can be any combination of letters or numbers from 1 to 16.

The no form of the command removes the authentication password from the configuration and effectively disables authentication.

Default Authentication is disabled and the authentication password is empty.

Parameters *authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash-key — The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

auth-keychain

Syntax **auth-keychain** *name*

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the BGP authentication key for all peers.

The keychain allows the rollover of authentication keys during the lifetime of a session.

Default **no auth-keychain**

Parameters *name* — Specifies the name of an existing keychain, up to 32 characters, to use for the specified TCP session or sessions.

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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 BGP protocol peering.

Default **no bfd-enable**

connect-retry

Syntax **connect-retry** *seconds*
no connect-retry

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the BGP connect retry timer value in seconds.

When this timer expires, BGP tries to reconnect to the configured peer. This configuration parameter can be set at three levels: global level (applies to all peers), peer-group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to the default value.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default 120 seconds

Parameters *seconds* — The BGP Connect Retry timer value in seconds, expressed as a decimal integer.

Values 1 — 65535

damping

Syntax [**no**] **damping**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables BGP route damping for learned routes which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route policy definition.

The **no** form of the command used at the global level disables route damping.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:

Half-life: 15 minutes
Max-suppress: 60 minutes
Suppress-threshold:3000
Reuse-threshold 750

Default no damping — Learned route damping is disabled.

disable-4byte-asn

Syntax [**no**] **disable-4byte-asn**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command disables the use of 4-byte ASNs. It can be configured at all 3 level of the hierarchy so it can be specified down to the per peer basis.

If this command is enabled 4-byte ASN support should not be negotiated with the associated remote peer(s).

The **no** form of the command resets the behavior to the default which is to enable the use of 4-byte ASN.

disable-capability-negotiation

Syntax [**no**] **disable-capability-negotiation**

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command disables the exchange of capbilities. When command is enabled and after the peering is flapped, any new capabilities are not negotiated and will strictly support IPv4 routing exchanges with that peer.

The **no** form of the command removes this command from the configuration and restores the normal behavior.

Default no disable-capability-negotiation

disable-client-reflect

Syntax [**no**] **disable-client-reflect**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command disables the reflection of routes by the route reflector to the group or neighbor. This only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients.

The **no** form re-enables client reflection of routes.

Default no disable-client-reflect — Client routes are reflected to all client peers.

disable-communities

Syntax **disable-communities** [**standard**] [**extended**]
no disable-communities

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures BGP to disable sending communities.

Parameters **standard —** Specifies standard communities that existed before VPRNs or 2547.

extended — Specifies BGP communities used were expanded after the concept of 2547 was introduced, to include handling the VRF target.

disable-fast-external-failover

Syntax [**no**] **disable-fast-external-failover**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures BGP fast external failover.

enable-peer-tracking

Syntax [**no**] **enable-peer-tracking**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables BGP peer tracking.

Default no enable-peer-tracking

export

Syntax **export** *policy* [*policy...*]
no export

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the export policies to be used to control routes advertised to BGP neighbors.

When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be configured. The first policy that matches is applied.

Note that if a non-existent route policy is applied to a VPRN instance, the CLI generates a warning message. This message is only generated at an interactive CLI session and the route policy association is made. No warning message is generated when a non-existent route policy is applied to a VPRN instance in a configuration file or when SNMP is used.

The **no** form of this command removes all route policy names from the export list.

Default no export — BGP advertises routes from other BGP routes but does not advertise any routes from other protocols unless directed by an export policy.

Parameters *policy* — A route policy statement name.

family

Syntax **family** [**ipv4**]
no family

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the IP family capability.

The **no** form of the command reverts to the default.

Default no family

Parameters *ipv4* — Provisions IPv4 support.

group

Syntax **group** *name* [**dynamic-peer**]
no group

Context config>service>vprn>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command creates a context to configure a BGP peer group.

The **no** form of the command deletes the specified peer group and all configurations associated with the peer group. The group must be shutdown before it can be deleted.

Default None — No peer groups are defined.

Parameters *name* — The peer group name. Allowed values is a 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.

dynamic-peer — This flag designates that the given BGP group will be used by BGP peers created dynamically based on subscriber-hosts pointing to corresponding BGP peering policy. There can be only one BGP group with this flag set in any given VPRN. No bBGP neighbous can be manually configured in a BGP group with this flag set.

Default disabled

neighbor

Syntax [**no**] **neighbor** *ip-address*

Context config>service>vprn>bgp>group

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command creates a BGP peer/neighbor instance within the context of the BGP group.

This command can be issued repeatedly to create multiple peers and their associated configuration.

The **no** form of the command is used to remove the specified neighbor and the entire configuration associated with the neighbor. The neighbor must be administratively **shutdown** before attempting to delete it. If the neighbor is not shutdown, the command will not result in any action except a warning message on the console indicating that neighbor is still administratively up.

Default none — No neighbors are defined.

Parameters *ip-address* — The IP address of the BGP peer router in dotted decimal notation.

Values ipv4-address : a.b.c.d

family

Syntax **family** [**ipv4**]
no family

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the address family or families to be supported over BGP peerings in the base router. This command is additive so issuing the **family** command adds the specified address family to the list.

The **no** form of the command removes the specified address family from the associated BGP peerings. If an address family is not specified, then reset the supported address family back to the default.

Default ipv4

Parameters **ipv4 —** Provisions support for IPv4 routing information.

hold-time

Syntax **hold-time** *seconds* [**strict**]
no hold-time

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the BGP hold time, expressed in seconds.

The BGP hold time specifies the maximum time BGP waits between successive messages (either keepalive or update) from its peer, before closing the connection. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

Even though the router OS implementation allows setting the **keepalive** time separately, the configured **keepalive** timer is overridden by the **hold-time** value under the following circumstances:

1. If the specified hold-time is less than the configured **keepalive** time, then the operational **keepalive** time is set to a third of the **hold-time**; the configured **keepalive** time is not changed.
2. If the **hold-time** is set to zero, then the operational value of the **keepalive** time is set to zero; the configured **keepalive** time is not changed. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default 90 seconds

Parameters *seconds* — The hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.

Values 0, 3 — 65535

strict — When this parameter is specified, the advertised BGP hold-time from the far-end BGP peer must be greater than or equal to the specified value.

import

Syntax **import** *policy* [*policy...*]
no import

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the import policies to be used to control routes advertised to BGP neighbors. Route policies are configured in the **config>router>policy-options** context. When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be specified. The first policy that matches is applied.

The **no** form of this command removes all route policy names from the import list.

Default no import — BGP accepts all routes from configured BGP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

Parameters *policy* — A route policy statement name.

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires. The *seconds* parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **keepalive** value is generally one-third of the **hold-time** interval. Even though the OS implementation allows the **keepalive** value and the **hold-time** interval to be independently set, under the following circumstances, the configured **keepalive** value is overridden by the **hold-time** value:

If the specified **keepalive** value is greater than the configured **hold-time**, then the specified value is ignored, and the **keepalive** is set to one third of the current **hold-time** value.

If the specified **hold-time** interval is less than the configured **keepalive** value, then the **keepalive** value is reset to one third of the specified **hold-time** interval.

If the **hold-time** interval is set to zero, then the configured value of the **keepalive** value is ignored. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default 30 seconds

Parameters *seconds* — The keepalive timer in seconds, expressed as a decimal integer.

Values 0 — 21845

local-address

Syntax **local-address** *ip-address*
no local-address

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Configures the local IP address used by the group or neighbor when communicating with BGP peers.

Outgoing connections use the **local-address** as the source of the TCP connection when initiating connections with a peer.

When a local address is not specified, the 7210 SAS uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of the command removes the configured local-address for BGP.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no local-address** — The router ID is used when communicating with IBGP peers and the interface address is used for directly connected EBGP peers.

ip-address — The local address expressed in dotted decimal notation. Allowed values are a valid routable IP address on the router, either an interface or system IP address.

local-as

Syntax **local-as** *as-number* [**private**]
no local-as

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures a BGP virtual autonomous system (AS) number.

In addition to the AS number configured for BGP in the `config>router>autonomous-system` context, a virtual (local) AS number is configured.The virtual AS number is added to the as-path message before the router's AS number makes the virtual AS the second AS in the as-path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). Thus, by specifying this at each neighbor level, it is possible to have a separate as-number per EBGP session.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The **private** attribute can be added or removed dynamically by reissuing the command.

Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number. Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local AS number.

This is an optional command and can be used in the following circumstance:

Provider router P is moved from AS1 to AS2. The customer router that is connected to P, however, is configured to belong to AS1. To avoid reconfiguring the customer router, the **local-as** value on router P can be set to AS1. Thus, router P adds AS1 to the as-path message for routes it advertises to the customer router.

The **no** form of the command used at the global level will remove any virtual AS number configured. The **no** form of the command used at the group level reverts to the value defined at the global level. The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default no local-as

Parameters *as-number* — The virtual autonomous system number, expressed as a decimal integer.

Values 1 — 65535

private — Specifies the local-as is hidden in paths learned from the peering.

local-preference

Syntax **local-preference** *local-preference*
no local-preference

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute. This value is used if the BGP route arrives from a BGP peer without the **local-preference** integer set.

The specified value can be overridden by any value set via a route policy. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no local-preference** — Does not override the local-preference value set in arriving routes and analyze routes without local preference with value of 100.

Parameters *local-preference* — The local preference value to be used as the override value, expressed as a decimal integer.

Values 0 — 4294967295

loop-detect

Syntax **loop-detect** {**drop-peer** | **discard-route** | **ignore-loop**| **off**}
no loop-detect

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures how the BGP peer session handles loop detection in the AS path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

Note that dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of the command used at the global level reverts to default, which is **loop-detect ignore-loop**.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default loop-detect ignore-loop

Parameters **drop-peer —** Sends a notification to the remote peer and drops the session.

discard-route — Discards routes received with loops in the AS path.

ignore-loop — Ignores routes with loops in the AS path but maintains peering.

off — Disables loop detection.

med-out

Syntax **med-out** {**number | igp-cost**}
no med-out

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set via a route policy.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to default where the MED is not advertised.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default no med-out

Parameters *number* — The MED path attribute value, expressed as a decimal integer.

Values 0 — 4294967295

igp-cost — The MED is set to the IGP cost of the given IP prefix.

min-as-origination

Syntax **min-as-origination** *seconds*
no min-as-origination

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the minimum interval, in seconds, at which a path attribute, originated by the local router, can be advertised to a peer.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to default.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default 15 seconds

Parameters *seconds* — The minimum path attribute advertising interval in seconds, expressed as a decimal integer.

Values 2 — 255

min-route-advertisement

Syntax **min-route-advertisement** *seconds*
no min-route-advertisement

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the minimum interval, in seconds, at which a prefix can be advertised to a peer.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command reverts to default values.

Default 30 seconds

Parameters *seconds* — The minimum route advertising interval, in seconds, expressed as a decimal integer.

Values 1— 255

multihop

Syntax **multihop** *ttl-value*
no multihop

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGP peer multiple hops away.

This parameter is meaningful only when configuring EBGP peers. It is ignored if set for an IBGP peer.

The **no** form of the command is used to convey to the BGP instance that the EBGP peers are directly connected.
The **no** form of the command reverts to default values.

Default **1** — EBGP peers are directly connected.

64 — IBGP

Parameters *ttl-value* — The TTL value, expressed as a decimal integer.

Values 1 — 255

passive

Syntax [**no**] **passive**

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables passive mode for the BGP group or neighbor.

When in passive mode, BGP will not attempt to actively connect to the configured BGP peers but responds only when it receives a connect open request from the peer.

The **no** form of the command used at the group level disables passive mode where BGP actively attempts to connect to its peers.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no passive —** BGP will actively try to connect to all the configured peers.

peer-as

Syntax **peer-as** *as-number*

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.

For EBGP peers, the peer AS number configured must be different from the autonomous system number configured for this router under the global level since the peer will be in a different autonomous system than this router

For IBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level.

This is a required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.

Default No AS numbers are defined.

Parameters *as-number* — The autonomous system number, expressed as a decimal integer.

Values 1 — 65535

preference

Syntax [**no**] **preference** *preference*

Context config>service>vprn>bgp
config>service>vprn>bgp>group

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the route preference for routes learned from the configured peer(s).

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The lower the preference the higher the chance of the route being the active route. The OS assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF.

The **no** form of the command used at the global level reverts to default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default 170

Parameters *preference* — The route preference, expressed as a decimal integer.

Values 1 — 255

prefix-limit

Syntax **prefix-limit** *limit* [**log-only**] [**threshold** *percent*]
no prefix-limit

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the maximum number of routes BGP can learn from a peer.

When the number of routes reaches a certain percentage (default is 90% of this limit), an SNMP trap is sent. When the limit is exceeded, the BGP peering is dropped and disabled.

The **no** form of the command removes the **prefix-limit**.

Default no prefix-limit

Parameters *limit* — The number of routes that can be learned from a peer, expressed as a decimal integer.

Values 1 — 4294967295

log-only — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, the BGP peering is not dropped.

percent — The threshold value (as a percentage) that triggers a warning message to be sent. The default value is 90%.

rapid-withdrawal

Syntax [**no**] **rapid-withdrawal**

Context config>service>vprn>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command disables the delay (Minimum Route Advertisement) on sending BGP withdrawals. Normal route withdrawals may be delayed up to the mininum route advertisement to allow for efficient packing of BGP updates.

The **no** form of the command removes this command from the configuration and returns withdrawal processing to the normal behavior.

Default no rapid-withdrawal

remove-private

Syntax [**no**] **remove-private**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command allows private AS numbers to be removed from the AS path before advertising them to BGP peers.

When the **remove-private** parameter is set at the global level, it applies to all peers regardless of group or neighbor configuration. When the parameter is set at the group level, it applies to all peers in the group regardless of the neighbor configuration.

The OS software recognizes the set of AS numbers that are defined by IANA as private. These are AS numbers in the range 64512 through 65535, inclusive.

The **no** form of the command used at the global level reverts to default value. The **no** form of the command used at the group level reverts to the value defined at the global level. The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no remove-private** — Private AS numbers will be included in the AS path attribute.

type

Syntax [**no**] **type** {**internal** | **external**}

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command designates the BGP peer as type internal or external.

The type of **internal** indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.

By default, the OS derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, then the peer is considered **external**.

The **no** form of the command used at the group level reverts to the default value.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no type** — Type of neighbor is derived on the local AS specified.

Parameters **internal** — Configures the peer as internal.

external — Configures the peer as external.

ospf

Syntax [**no**] **ospf**

Context config>service>vprn

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables access to the context to enable an OSPF protocol instance.

When an OSPF instance is created, the protocol is enabled. To start or suspend execution of the OSPF protocol without affecting the configuration, use the **no shutdown** command.

The **no** form of the command deletes the OSPF protocol instance removing all associated configuration parameters.

Default **no ospf** — The OSPF protocol is not enabled.

area

Syntax [**no**] **area** *area-id*

Context config>service>vprn>ospf

Platforms Supported: 7210 SAS-K 2F4T6C.

This command creates the context to configure an OSPF area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted decimal notation or as a 32-bit decimal integer.

The **no** form of the command deletes the specified area from the configuration. Deleting the area also removes the OSPF configuration of all the interfaces, virtual-links, sham-links, and address-ranges etc., that are currently assigned to this area.

Default **no area** — No OSPF areas are defined.

Parameters *area-id* — The OSPF area ID expressed in dotted decimal notation or as a 32-bit decimal integer.

Values 0.0.0.0 — 255.255.255.255 (dotted decimal)
0 — 4294967295 (decimal integer)

area-range

Syntax **area-range** *ip-prefix/prefix-length* [**advertise** | **not-advertise**]
no area-range *ip-prefixImask*
no area-range *ip-prefixImask*

Context config>service>vprn>ospf>area
ospf>service>vprn>nssa

Platforms Supported: 7210 SAS-K 2F4T6C.

This command creates ranges of addresses on an Area Border Router (ABR) for the purpose of route summarization or suppression. When a range is created, the range is configured to be advertised or

not advertised into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.

ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.

The **no** form of the command deletes the range (non) advertisement.

Default **no area-range** — No range of addresses are defined.

Special Cases **NSSA Context —** In the NSSA context, the option specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs.

Area Context — If this command is not entered under the NSSA context, the range applies to summary LSAs even if the area is an NSSA.

Parameters *ipv6-prefix/prefix-length —* The IP prefix in dotted decimal notation for the range used by the ABR to advertise that summarizes the area into another area.

Values ipv6-prefix - 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
prefix-length - [0..128]

mask — The subnet mask for the range expressed as a decimal integer mask length or in dotted decimal notation.

Values 0 — 32 (mask length), 0.0.0.0 — 255.255.255.255 (dotted decimal)

advertise | **not-advertise —** Specifies whether or not to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range will be advertised, and the keyword **not-advertise** indicates the range will not be advertised.
The default is **advertise**.

blackhole-aggregate

Syntax [**no**] **blackhole-aggregate**

Context config>service>vprn>ospf>area

Platforms Supported: 7210 SAS-K 2F4T6C.

This command installs a low priority blackhole route for the entire aggregate. Exisiting routes that make up the aggregate will have a higher priority and only the components of the range for which no route exists are blackholed.

It is possible that when performing area aggregation, addresses may be included in the range for which no actual route exists. This can cause routing loops. To avoid this problem configure the blackhole aggregate option.

The **no** form of this command removes this option.

Default blackhole-aggregate

interface

Syntax [**no**] **interface** *ip-int-name* [**secondary**]

Context config>service>vprn>ospf>area

Platforms Supported: 7210 SAS-K 2F4T6C.

This command creates a context to configure an OSPF interface.

By default interfaces are not activated in any interior gateway protocol such as OSPF unless explicitly configured.

The **no** form of the command deletes the OSPF interface configuration for this interface. The **shutdown** command in the **config>router>ospf>interface** context can be used to disable an interface without removing the configuration for the interface.

Default **no interface** — No OSPF interfaces are defined.

Parameters *ip-int-name* — The IP interface name. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service vprn interface** commands. An interface name cannot be in the form of an IP address. Interface names can be 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.

If the IP interface name does not exist or does not have an IP address configured an error message will be returned.

If the IP interface exists in a different area it will be moved to this area.

secondary — Allows multiple secondary adjacencies to be established over a single IP interface.

sham-link

Syntax **sham-link** *ip-int-name ip-address*

Context config>service>vprn>ospf>area

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command is similar to a virtual link with the exception that metric must be included in order to distinguish the cost between the MPLS-VPRN link and the backdoor.

Parameters *ip-int-name* — The local interface name used for the sham-link. This is a mandatory parameter and interface names must be unique within the group of defined IP interfaces for **config>router>interface**, **config>service>ies>interface** and **config>service>vprn>interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters, the entire string must be enclosed within double quotes. If the IP interface name does not exist or does not have an IP address configured, an error message will be returned.

ip-address — The IP address of the SHAM-link neighbor in IP address dotted decimal notation. This parameter is the remote peer of the sham link's IP address used to set up the SHAM link. This is a mandatory parameter and must be a valid IP address.

advertise-subnet

Syntax [**no**] **advertise-subnet**

Context config>service>vprn>ospf>area>*if*

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables advertising point-to-point interfaces as subnet routes (network number and mask). When disabled, point-to-point interfaces are advertised as host routes.

Note that this command is not valid in the OSPF3 context.

The **no** form of the command disables advertising point-to-point interfaces as subnet routes meaning they are advertised as host routes.

Default **advertise-subnet** — Advertises point-to-point interfaces as subnet routes.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>service>vprn>ospf>area>*if*
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.

Note that this command is not valid in the OSPF3 context.

All neighboring routers must use the same type of authentication and password for proper protocol communication. If the **authentication-type** is configured as password, then this key must be configured.

By default, no authentication key is configured.

The **no** form of the command removes the authentication key.

Default **no authentication-key** — No authentication key is defined.

Parameters *authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 8 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash-key — The hash key. The key can be any combination of ASCII characters up to 22 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

authentication-type

Syntax **authentication-type** {**password** | **message-digest**}
no authentication-type

Context config>service>vprn>ospf>area>*if*
config>service>vprn>ospf>area>virtual-link

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables authentication and specifies the type of authentication to be used on the OSPF interface, virtual-link, and sham-link.

Note that this command is not valid in the OSPF3 context.

Both simple **password** and **message-digest** authentication are supported.

By default, authentication is not enabled on an interface.

The **no** form of the command disables authentication on the interface.

Default **no authentication** — No authentication is enabled on an interface.

Parameters **password —** This keyword enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

message-digest — This keyword enables message digest MD5 authentication in accordance with RFC1321. If this option is configured, then at least one message-digest-key must be configured

bfd-enable

Syntax [**no**] **bfd-enable** [**remain-down-on-failure**]

Context config>router>ospf>area>interface

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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.

NOTES:

- BFD is not supported for IPv6 interfaces.

- For more information about the protocols and platforms that support BFD, see the "7210 SAS Router Configuration User Guide".

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

Default **no bfd-enable**

Parameters **remain-down-on-failure** — Forces adjacency down on BFD failure.

dead-interval

Syntax **dead-interval** *seconds*
no dead-interval

Context config>service>vprn>ospf>area>*if*
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the time, in seconds, that OSPF waits before declaring a neighbor router down. If no hello packets are received from a neighbor for the duration of the dead interval, the router is assumed to be down. The minimum interval must be two times the hello interval.

The **no** form of the command reverts to the default value.

Default **40**

Special Cases **OSPF Interface —** If the **dead-interval** configured applies to an interface, then all nodes on the subnet must have the same dead interval.

Virtual Link — If the **dead-interval** configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same dead interval.

Sham-link — If the **dead-interval** configured applies to a sham-link, then the interval on both endpoints of the sham-link must have the same dead interval.

Parameters *seconds* — The dead interval expressed as a decimal integer.

Values 2 — 2147483647 seconds

hello-interval

Syntax **hello-interval** *seconds*
no hello-interval

Context config>service>vprn>ospf>area>*if*
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the interval between OSPF hellos issued on the interface, virtual link, or sham-link.

The hello interval, in combination with the dead-interval, is used to establish and maintain the adjacency. Use this parameter to edit the frequency that hello packets are sent.

Reducing the interval, in combination with an appropriate reduction in the associated **dead-interval**, allows for faster detection of link and/or router failures at the cost of higher processing costs.

The **no** form of this command reverts to the default value.

Default **hello-interval 10** — A 10-second hello interval.

Special Cases **OSPF Interface —** If the **hello-interval** configured applies to an interface, then all nodes on the subnet must have the same hello interval.

Virtual Link — If the **hello-interval** configured applies to a virtual link, then the interval on both termination points of the virtual link must have the same hello interval.

Sham Link — If the hello-interval configured applies to a sham-link, then the interval on both endpoints of the sham-link must have the same hello interval

Parameters *seconds* — The hello interval in seconds expressed as a decimal integer.

Values 1 — 65535

interface-type

Syntax **interface-type** {**broadcast** | **point-to-point**}
no interface-type

Context config>service>vprn>ospf>area>*if*

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures the interface type to be either broadcast or point-to-point.

Use this command to set the interface type of an Ethernet link to point-to-point to avoid having to carry the broadcast adjacency maintenance overhead if the Ethernet link provided the link is used as a point-to-point.

If the interface type is not known at the time the interface is added to OSPF and subsequently the IP interface is bound (or moved) to a different interface type, this command must be entered manually.

The **no** form of the command reverts to the default value.

Default **point-to-point** — If the physical interface is SONET.

broadcast — If the physical interface is Ethernet or unknown.

Special Cases **Virtual-Link —** A virtual link is always regarded as a point-to-point interface and not configurable.

Parameters **broadcast —** Configures the interface to maintain this link as a broadcast network. To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.

point-to-point — Configures the interface to maintain this link as a point-to-point link.

message-digest-key

Syntax **message-digest-key** *keyid* **md5** [*key* | *hash-key*] [**hash**]
no message-digest-key *keyid*

Context config>service>vprn>ospf>area>*if*
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures a message digest key when MD5 authentication is enabled on the interface, virtual-link or sham-link. Multiple message digest keys can be configured.

Note that this command is not valid in the OSPF3 context.

The **no** form of the command removes the message digest key identified by the *key-id*.

Default No message digest keys are defined.

Parameters **keyid —** The *keyid* is expressed as a decimal integer.

Values 1 — 255

md5 *key* **—** The MD5 key. The *key* can be any alphanumeric string up to 16 characters in length.

md5 *hash-key* — The MD5 hash key. The key can be any combination of ASCII characters up to 32 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

metric

Syntax **metric** *metric*
no metric

Context config>service>vprn>ospf>area>if
config>service>vprn>ospf>area>sham-link

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures an explicit route cost metric for the OSPF interface that overrides the metrics calculated based on the speed of the underlying link.

The **no** form of the command deletes the manually configured interface metric, so the interface uses the computed metric based on the **reference-bandwidth** command setting and the speed of the underlying link.

Default **no metric** — The metric is based on **reference-bandwidth** setting and the link speed.

Parameters *metric* — The metric to be applied to the interface expressed as a decimal integer.

Values 1 — 65535

mtu

Syntax **mtu** *bytes*
no mtu

Context config>service>vprn>ospf>area>if

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures the OSPF packet size used on this interface. If this parameter is not configured OSPF derives the MTU value from the MTU configured (default or explicitly) in the following contexts:

config>port>ethernet
config>port>sonet-sdh>path
config>port>tdm>t3-e3
config>port>tdm>t1-e1>channel-group

If this parameter is configured, the smaller value between the value configured here and the MTU configured (default or explicitly) in an above-mentioned context is used.

To determine the actual packet size add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured in this command.

Use the **no** form of this command to revert to default.

Default **no mtu** — Uses the value derived from the MTU configured in the **config>port** context.

Parameters *bytes* — The MTU to be used by OSPF for this logical interface in bytes.

Values 512 — 9198 (9212-14) (Depends on the physical media)

passive

Syntax [**no**] **passive**

Context config>service>vprn>ospf>area>if

Platforms Supported: 7210 SAS-K 2F4T6C.

This command adds the passive property to the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol.

By default, only interface addresses that are configured for OSPF will be advertised as OSPF interfaces. The **passive** parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol.

While in passive mode, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets.

The **no** form of the command removes the passive property from the OSPF interface.

Default Service interfaces defined in **config>router>service-prefix** are passive.

All other interfaces are not passive.

priority

Syntax **priority** *number*
no priority

Context config>service>vprn>ospf>area>if

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures the priority of the OSPF interface that is used an election of the designated router on on the subnet.

This parameter is only used if the interface is of type broadcast. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be Designated Router or Backup Designated Router.

The **no** form of the command reverts the interface priority to the default value.

Default **priority 1**

Parameters *number* — The interface priority expressed as a decimal integer. A value of 0 indicates the router is not eligible to be the Designated Router of Backup Designated Router on the interface subnet.

Values 0 — 255

retransmit-interval

Syntax **retransmit-interval** *seconds*
no retransmit-interval

Context config>service>vprn>ospf>area>if
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the length of time, in seconds, that OSPF will wait before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor.

The value should be longer than the expected round trip delay between any two routers on the attached network. Once the retransmit-interval expires and no acknowledgement has been received, the LSA will be retransmitted.

The **no** form of this command reverts to the default interval.

Default **retransmit-interval 5**

Parameters *seconds* — The retransmit interval in seconds expressed as a decimal integer.

Values 1 — 3600

transit-delay

Syntax **transit-delay** *seconds*
no transit-delay

Context config>service>vprn>ospf>area>if
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf>area>sham-link

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command configures the estimated time, in seconds, that it takes to transmit a link state advertisement (LSA) on the interface or virtual link or sham-link.

The **no** form of this command reverts to the default delay time.

Default **transit-delay 1**

Parameters *seconds* — The transit delay in seconds expressed as a decimal integer.

Values 0 — 3600

nssa

Syntax [**no**] **nssa**

Context config>service>vprn>ospf>area

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command creates the context to configure an OSPF Not So Stubby Area (NSSA) and adds/removes the NSSA designation from the area.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF domain.

Existing virtual links of a non-stub or NSSA area will be removed when the designation is changed to NSSA or stub.

An area can be designated as stub or NSSA but never both at the same time.

By default, an area is not configured as an NSSA area.

The **no** form of the command removes the NSSA designation and configuration context from the area.

Default **no nssa** — The OSPF area is not an NSSA.

originate-default-route

Syntax **originate-default-route** [**type-7**]
no originate-default-route

Context config>service>vprn>ospf>area>nssa

Platforms Supported: 7210 SAS-K 2F4T6C.

This command enables the generation of a default route and its LSA type (3 or 7) into a Not So Stubby Area (NSSA) by an NSSA Area Border Router (ABR).

When configuring an NSSA with no summaries, the ABR will inject a type 3 LSA default route into the NSSA area. Some older implementations expect a type 7 LSA default route.

The **no** form of the command disables origination of a default route.

Default **no originate-default-route** — A default route is not originated.

Parameters **type-7 —** Specifies a type 7 LSA should be used for the default route.

Configure this parameter to inject a type-7 LSA default route instead the type 3 LSA into the NSSA configured with no summaries.
To revert to a type 3 LSA, enter **originate-default-route** without the **type-7** parameter.

Default Type 3 LSA for the default route.

redistribute-external

Syntax [**no**] **redistribute-external**

Context config>service>vprn>ospf>area>nssa

Platforms Supported: 7210 SAS-K 2F4T6C.

This command enables the redistribution of external routes into the Not So Stubby Area (NSSA) or an NSSA area border router (ABR) that is exporting the routes into non-NSSA areas.

NSSA or Not So Stubby Areas are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an Area Border Router to the entire OSPF domain.

The **no** form of the command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.

Default **redistribute-external** — External routes are redistributed into the NSSA.

summaries

Syntax [**no**] **summaries**

Context config>service>vprn>ospf>area>nssa
config>service>vprn>ospf>area>stub

Platforms Supported: 7210 SAS-K 2F4T6C.

This command enables sending summary (type 3) advertisements into a stub area or Not So Stubby Area (NSSA) on an Area Border Router (ABR). This parameter is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub or nssa area. By default, summary route advertisements are sent into the stub area or NSSA.

The **no** form of the command disables sending summary route advertisements and, for stub areas, only the default route is advertised by the ABR.

Default **summaries** — Summary routes are advertised by the ABR into the stub area or NSSA.

stub

Syntax [**no**] **stub**

Context config>service>vprn>ospf>area

Platforms Supported: 7210 SAS-K 2F4T6C.

This command enables access to the context to configure an OSPF stub area and adds/removes the stub designation from the area. External routing information is not flooded into stub areas. All routers in the stub area must be configured with the **stub** command. An OSPF area cannot be both an NSSA and a stub area. Existing virtual links of a non STUB or NSSA area will be removed when its designation is changed to NSSA or STUB.

By default, an area is not a stub area.

The **no** form of the command removes the stub designation and configuration context from the area.

Default **no stub** — The area is not configured as a stub area.

default-metric

Syntax **default-metric** *metric*
no default-metric

Context config>service>vprn>ospf>area>stub

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures the metric used by the area border router (ABR) for the default route into a stub area. The default metric should only be configured on an ABR of a stub area. An ABR generates a default route if the area is a **stub** area.

The **no** form of the command reverts to the default value.

Default default-metric 1

Parameters *metric* — The metric expressed as a decimal integer for the default route cost to be advertised into the stub area.

Values 1 — 16777215

virtual-link

Syntax [**no**] **virtual-link** *router-id* **transit-area** *area-id*

Context config>service>vprn>ospf>area

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures a virtual link to connect area border routers to the backbone via a virtual link. The backbone area (area 0.0.0.0) must be contiguous and all other areas must be connected to the backbone area. If it is not practical to connect an area to the backbone (see area 0.0.0.2 in the picture below) then the area border routers (routers 1 and 2 in the picture below) must be connected via a virtual link. The two area border routers will form a point-to-point like adjacency across the transit area (area 0.0.0.1 in the picture below). A virtual link can only be configured while in the area 0.0.0.0 context.

The *router-id* specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or a Not So Stubby Area (NSSA).

The **no** form of the command deletes the virtual link.

Default No virtual link is defined.

Parameters *router-id* — The router ID of the virtual neighbor in IP address dotted decimal notation.

transit-area *area-id* — The area-id specified identifies the transit area that links the backbone area with the area that has no physical connection with the backbone.

The OSPF backbone area, area 0.0.0.0, must be contiguous and all other areas must be connected to the backbone area. The backbone distributes routing information between areas. If it is not practical to connect an area to the backbone (see Area 0.0.0.5 in Figure 75) then the area border routers (such as routers Y and Z) must be connected via a virtual link. The two area border routers form a point-to-point-like adjacency across the transit area (see Area 0.0.0.4).

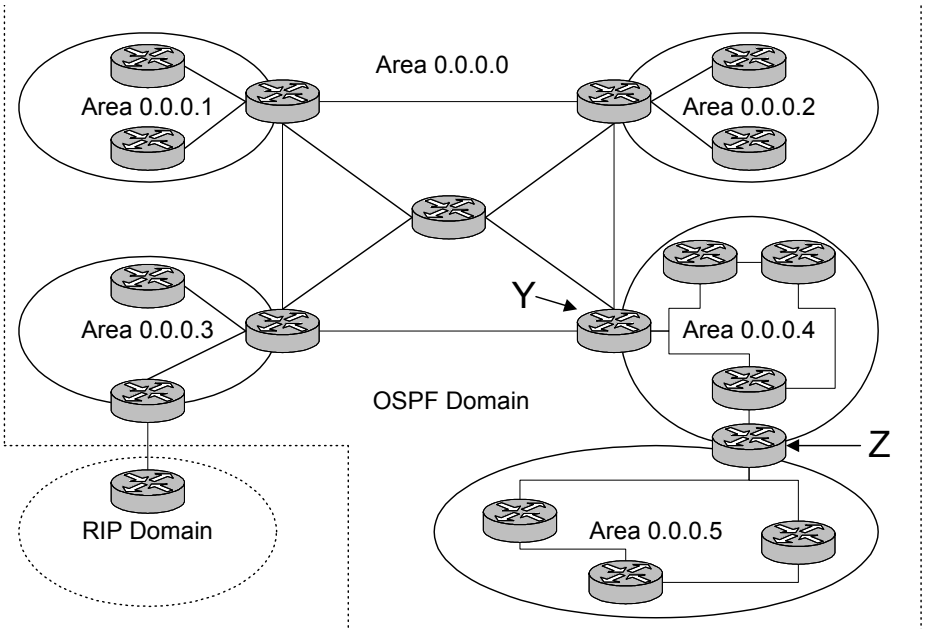

Figure 75: OSPF Areas

compatible-rfc1583

Syntax [**no**] **compatible-rfc1583**

Context config>service>vprn>ospf

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables OSPF summary and external route calculations in compliance with RFC1583 and earlier RFCs.

RFC1583 and earlier RFCs use a different method to calculate summary and external route costs. To avoid routing loops, all routers in an OSPF domain should perform the same calculation method.

Although it would be favorable to require all routers to run a more current compliancy level, this command allows the router to use obsolete methods of calculation.

The **no** form of the command enables the post-RFC1583 method of summary and external route calculation.

Default **compatible-rfc1583** — RFC1583 compliance is enabled.

export

Syntax **export** *policy-name* [*policy-name…*]
no export

Context config>service>vprn>ospf

Platforms Supported: 7210 SAS-K 2F4T6C.

This command associates export route policies to determine which routes are exported from the route table to OSPF. Export polices are only in effect if OSPF is configured as an ASBR.

If no export policy is specified, non-OSPF routes are not exported from the routing table manager to OSPF.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of the command removes all policies from the configuration.

Default **no export** — No export route policies specified.

Parameters *policy-name* — The export route 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.

The specified name(s) must already be defined.

external-db-overflow

Syntax **external-db-overflow** *limit interval*
no external-db-overflow

Context config>service>vprn>ospf

Platforms Supported: 7210 SAS-K 2F4T6C.

This command enables limits on the number of non-default AS-external-LSA entries that can be stored in the LSDB and specifies a wait timer before processing these after the limit is exceeded.

The *limit* value specifies the maximum number of non-default AS-external-LSA entries that can be stored in the link-state database (LSDB). Placing a limit on the non-default AS-external-LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reach or exceed the *limit*, the table is in an overflow state. When in an overflow state, the router will not originate any new AS-external-LSAs. In fact, it withdraws all the self-originated non-default external LSAs.

The *interval* specifies the amount of time to wait after an overflow state before regenerating and processing non-default AS-external-LSAs. The waiting period acts like a dampening period preventing the router from continuously running Shortest Path First (SPF) calculations caused by the excessive number of non-default AS-external LSAs.

The **external-db-overflow** must be set identically on all routers attached to any regular OSPF area. OSPF stub areas and not-so-stubby areas (NSSAs) are excluded.

The **no** form of the command disables limiting the number of non-default AS-external-LSA entries.

Default **no external-db-overflow** — No limit on non-default AS-external-LSA entries.

Parameters *limit* — The maximum number of non-default AS-external-LSA entries that can be stored in the LSDB before going into an overflow state expressed as a decimal integer.

Values -1 — 2147483647

interval — The number of seconds after entering an overflow state before attempting to process non-default AS-external-LSAs expressed as a decimal integer.

Values 0 — 2147483647

external-preference

Syntax **external-preference** *preference*
no external-preference

Context config>service>vprn>ospf

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures the preference for OSPF external routes.

A route can be learned by the router from different protocols in which case the costs are not comparable; when this occurs the preference is used to decide which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the config>router context.

The **no** form of the command reverts to the default value.

Default **external-preference 150** — OSPF external routes have a default preference of 150.

Parameters *preference* — The preference for external routes expressed as a decimal integer.

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes[a]
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes

Route Type	Preference	Configurable (Continued)
IS-IS level 2 external	165	Yes
BGP	170	Yes

a. Preference for OSPF internal routes is configured with the **preference** command.

Values 1 — 255

ignore-dn-bit

Syntax [**no**] **ignore-dn-bit**

Context config>service>vprn>ospf

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies whether to ignore the DN bit for OSPF LSA packets for this instance of OSPF on the router. When enabled, the DN bit for OSPF LSA packets will be ignored. When disabled, the DN bit will not be ignored for OSPF LSA packets.

overload

Syntax **overload** [**timeout** *seconds*]
no overload

Context config>service>vprn>ospf

Platforms Supported: 7210 SAS-K 2F4T6C.

This command changes the overload state of the local router so that it appears to be overloaded. When overload is enabled, the router can participate in OSPF routing, but is not used for transit traffic. Traffic destined to directly attached interfaces continue to reach the router.

To put the IGP in an overload state enter a timeout value. The IGP will enter the overload state until the timeout timer expires or a **no overload** command is executed.

If the **overload** command is encountered during the execution of an **overload-on-boot** command then this command takes precedence. This could occur as a result of a saved configuration file where both parameters are saved. When the file is saved by the system the **overload-on-boot** command is saved after the **overload** command.

Use the **no** form of this command to return to the default. When the **no overload** command is executed, the overload state is terminated regardless the reason the protocol entered overload state.

Default no overload

Parameters **timeout** *seconds* — Specifies the number of seconds to reset overloading.

Values 60 —1800

Default 60

overload-include-stub

Syntax [**no**] **overload-include-stub**

Context config>service>vprn>ospf

Platforms Supported: 7210 SAS-K 2F4T6C.

This command is used to to determine if the OSPF stub networks should be advertised with a maximum metric value when the system goes into overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, will be advertised at the maximum metric.

Default no overload-include-stub

overload-on-boot

Syntax **overload-on-boot** [**timeout** *seconds*]
no overload

Context config>service>vprn>ospf

Platforms Supported: 7210 SAS-K 2F4T6C.

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:

- The timeout timer expires.
- A manual override of the current overload state is entered with the **no overload** command.

  The **no overload** command does not affect the **overload-on-boot** function.

The **no** form of the command removes the overload-on-boot functionality from the configuration.

Default no overload-on-boot

Parameters **timeout** *seconds* — Specifies the number of seconds to reset overloading.

Values 60 —1800

Default 60

preference

Syntax **preference** *preference*
no preference

Context config>service>vprn>ospf

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures the preference for OSPF internal routes.

A route can be learned by the router from different protocols in which case the costs are not comparable, when this occurs the preference is used to decide to which route will be used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, then the decision of what route to use is determined by the configuration of the **ecmp** in the config>router context.

The **no** form of the command reverts to the default value.

Default **preference 10** — OSPF internal routes have a preference of 10.

Parameters *preference* — The preference for internal routes expressed as a decimal integer. Defaults for different route types are listed in the following table.

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes[a]
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

a. Preference for OSPF internal routes is configured with the **preference** command.

Values 1 — 255

reference-bandwidth

Syntax **reference-bandwidth** *reference-bandwidth*
no reference-bandwidth

Context config>service>vprn>ospf

Platforms Supported: 7210 SAS-K 2F4T6C.

This command configures the reference bandwidth in kilobits per second (Kbps) that provides the reference for the default costing of interfaces based on their underlying link speed.

The default interface cost is calculated as follows:

cost = reference-bandwidth ÷ bandwidth

The default *reference-bandwidth* is 100,000,000 Kbps or 100 Gbps, so the default auto-cost metrics for various link speeds are as as follows:

- 10 Mbs link default cost of 10000
- 100 Mbs link default cost of 1000
- 1 Gbps link default cost of 100
- 10 Gbps link default cost of 10

The **reference-bandwidth** command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the **metric** *metric* command in the **config>router>ospf>area>interface** *ip-int-name* context.

The **no** form of the command reverts the reference-bandwidth to the default value.

Default **reference-bandwidth 100000000** — Reference bandwidth of 100 Gbps.

Parameters *reference-bandwidth* — The reference bandwidth in kilobits per second expressed as a decimal integer.

Values 1 — 1000000000

super-backbone

Syntax [**no**] **super-backbone**

Context config>service>vprn>ospf

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies whether CE-PE functionality is required or not. The OSPF super backbone indicates the type of the LSA generated as a result of routes redistributed into OSPF. When enabled, the redistributed routes are injected as summary, external or NSSA LSAs. When disabled, the redistributed routes are injected as either external or NSSA LSAs only.

Default no super-backbone

suppress-dn-bit

Syntax [**no**] **suppress-dn-bit**

Context config>service>vprn>ospf

Platforms Supported: 7210 SAS-K 2F4T6C.

This command specifies whether to suppress the setting of the DN bit for OSPF LSA packets generated by this instance of OSPF on the router. When enabled, the DN bit for OSPF LSA packets generated by this instance of the OSPF router will not be set. When disabled, this instance of the OSPF router will follow the normal proceedure to determine whether to set the DN bit.

Default no suppress-dn-bit

timers

Syntax **timers**

Context config>service>vprn>ospf

Platforms Supported: 7210 SAS-K 2F4T6C.

This command enables the context that allows for the configuration of OSPF timers. Timers control the delay between receipt of a link state advertisement (LSA) requiring a Dijkstra (Shortest Path First (SPF)) calculation and the minimum time between successive SPF calculations.

Changing the timers affect CPU utilization and network reconvergence times. Lower values reduce convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase reconvergence time.

Default none

spf-wait

Syntax **spf-wait** *max-spf-wait* [*spf-initial-wait* [*spf-second-wait*]]
no spf-wait

Context config>service>vprn>ospf>timers

Platforms Supported: 7210 SAS-K 2F4T6C.

This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, then the next SPF will run after 2000 milliseconds, and then next SPF will run after 4000 milliseconds, etc., until it reaches the **spf-wait** value. The SPF interval will stay at the **spf-wait** value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

The timer must be entered in increments of 100 milliseconds. Values entered that do not match this requirement will be rejected.

Use the **no** form of this command to return to the default.

Default no spf-wait

Parameters *max-spf-wait* — Specifies the maximum interval in milliseconds between two consecutive SPF calculations.

Values 1 — 120000

Default 1000

spf-initial-wait — Specifies the initial SPF calculation delay in milliseconds after a topology change.

Values 10 — 100000

Default 1000

spf-second-wait — Specifies the hold time in milliseconds between the first and second SPF calculation.

Values 10 — 100000

Default 1000

vpn-domain

Syntax **vpn-domain** [*type* {**0005** | **0105** | **0205** | **8005**}] **id** *id*
no vpn-domain

Context config>service>vprn>ospf

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID. This applies to VPRN instances of OSPF only. An attempt to modify the value of this object will result in an inconsistent value error when is not a VPRN instance. The parameters are mandatory and can be entered in either order. This command is not applicable in the **config>service>vprn>ospf3** context.

Default no vpn-domain

Parameters *id* — Specifies the OSPF VPN domain in the "xxxx.xxxx.xxxx" format. This is exchanged using BGP in the extended community attribute associated with a prefix. This object applies to VPRN instances of OSPF only.

type — Specifies the type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID.

Values 0005, 0105, 0205, 8005

vpn-tag

Syntax **vpn-tag** *vpn-tag*
no vpn-tag

Context config>service>vprn>ospf

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command specifies the route tag for an OSPF VPN on a PE router. This field is set in the tag field of the OSPF external LSAs generated by the PE. This is mainly used to prevent routing loops. This applies to VPRN instances of OSPF only. An attempt to modify the value of this object will result in an inconsistent value error when is not a VPRN instance.

Default vpn-tag 0

lsa-arrival

Syntax **lsa-arrival** *lsa-arrival-time*
no lsa-arrival

Context config>service>vprn>ospf>timers

Platforms Supported: 7210 SAS-K 2F4T6C.

This parameter defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors. It is recommended that the neighbors configured (**lsa-generate**) *lsa-second-wait* interval is equal or greater then the **lsa-arrival** timer configured here.

Use the **no** form of this command to return to the default.

Default no lsa-arrival

Parameters *lsa-arrival-time* — Specifies the timer in milliseconds. Values entered that do not match this requirement will be rejected.

Values 0 — 600000

lsa-generate

Syntax **lsa-generate** *max-lsa-wait* [*lsa-initial-wait* [*lsa-second-wait*]]
no lsa-generate-interval

Context config>service>vprn>ospf>timers

Platforms Supported: 7210 SAS-K 2F4T6C.

This parameter customizes the throttling of OSPF LSA-generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the *lsa-second-wait* timer until a maximum value is reached. Configuring the **lsa-arrival** interval to equal or less than the *lsa-second-wait* interval configured in the **lsa-generate** command is recommended.

Use the **no** form of this command to return to the default.

Default no lsa-generate

Parameters *max-lsa-wait* — Specifies the maximum interval, in milliseconds, between two consecutive ocurrences of an LSA being generated.

The timer must be entered as either 1 or in millisecond increments. Values entered that do not match this requirement will be rejected.

Values 1 — 600000

egress-label

Syntax **egress-label** *start-label* **[***end-label***]**

Context show>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Display services using the range of egress labels.

If only the mandatory *start-label* parameter is specified, only services using the specified label are displayed.

If both *start-label* and *end-label* parameters are specified, the services using the range of labels X where *start-label* <= X <= *end-label* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *start-label* — The starting egress label value for which to display services using the label range. If only *egress-label1* is specified, services only using *egress-label1* are displayed.

Values 0 | 2048 — 131071

end-label — The ending egress label value for which to display services using the label range.

Default The *egress-label1* value.

Values 2049 — 131071

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#
```

ingress-label

Syntax **ingress-label** *start-label* [*end-label*]

Context show>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Display services using the range of ingress labels.

If only the mandatory *start-label* parameter is specified, only services using the specified label are displayed.

If both *start-label* and *end-label* parameters are specified, the services using the range of labels X where *start-label* <= X <= *end-label* are displayed.

Use the **show router** *vprn-service-id* **ldp bindings** command to display dynamic labels.

Parameters *start-label* — The starting ingress label value for which to display services using the label range. If only *start-label* is specified, services only using *start-label* are displayed.

Values 0, 2048 — 131071

end-label — The ending ingress label value for which to display services using the label range.

Default The *start-label* value.

Values 2048 — 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 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** [**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*

Context show>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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 *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 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.

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

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.
E.QoS	The SAP egress QoS policy number specified on the egress SAP.
E.Mac/IP	The MAC or IP filter policy ID applied to the egress SAP
A.Pol	The accounting policy ID assigned to the SAP.

Label	Description (Continued)
Adm	The desired state of the SAP.
Opr	The actual state of the SAP.

Sample Output

```
*A:ALA-12# show service sap-using sap 1/1
===============================================================================
Service Access Points
===============================================================================
PortId          SvcId      SapMTU I.QoS I.Mac/IP E.QoS E.Mac/IP A.Pol  Adm  Opr
-------------------------------------------------------------------------------
1/1/7:0         1          1518   10    8        10    none     none   Up   Up
1/1/11:0        100        1514   1     none     1     none     none   Down Down
1/1/7:300       300        1518   10    none     10    none     1000   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 3
-------------------------------------------------------------------------------
*A:ALA-12#
```

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]

Context show>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Displays SDP information.

If no optional parameters are specified, a summary SDP output for all SDPs is displayed.

Parameters *sdp-id* — The SDP ID for which to display information.

Default All SDPs.

Values 1 — 17407

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.

Output **Show Service SDP —** The following table describes show service SDP output fields:

Label	Description
SDP Id	The SDP identifier.
Adm MTU	Specifies the 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 MPLS tunnel defined by this SDP.
Adm Admin State	Specifies the state of the SDP.
Opr Oper State	Specifies the operating state of the SDP.
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.
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.
Deliver Delivered	Specifies the type of delivery used by the SDP: MPLS.
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.

Label	Description (Continued)
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
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.
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.

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
-------------------------------------------------------------------------------
Number of SDPs : 5
===============================================================================
*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
===============================================================================
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#


*A:dut-a>show>service# sdp

===============================================================================
Services: Service Destination Points
===============================================================================
SdpId  AdmMTU  OprMTU  IP address       Adm  Opr         Del     LSP   Sig
-------------------------------------------------------------------------------
17406  0       9166    2.2.2.2          Up   Up          MPLS    L/B   TLDP
17407  0       9166    3.3.3.3          Up   Up          MPLS    L/B   TLDP
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
        I = SR-ISIS, O = SR-OSPF
===============================================================================
*A:dut-a>show>service#
```

sdp-using

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

Context show>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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 X —** 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#


A:ALA-48# show service sdp-using
===============================================================================
SDP Using
===============================================================================
SvcId       SdpId               Type Far End        Opr State I.Label  E.Label
-------------------------------------------------------------------------------
3           2:3                 Spok 10.20.1.2      Up        n/a      n/a
103         3:103               Spok 10.20.1.3      Up        131067   131068
103         4:103               Spok 10.20.1.2      Up        131065   131069
105         3:105               Spok 10.20.1.3      Up        131066   131067
-------------------------------------------------------------------------------
Number of SDPs : 4
-------------------------------------------------------------------------------
A:ALA-48

*A:dut-a>show>service# sdp-using
```

```
===============================================================================
SDP Using
===============================================================================
SvcId      SdpId              Type   Far End             Opr   I.Label E.Label
                                                        State
-------------------------------------------------------------------------------
5          17406:4294967294   BgpAd  2.2.2.2             Up    131067  131066
5          17407:4294967295   BgpAd  3.3.3.3             Up    131066  131066
-------------------------------------------------------------------------------
Number of SDPs : 2
-------------------------------------------------------------------------------
===============================================================================
*A:dut-a>show>service#
```

service-using

Syntax **service-using** [**epipe**] [**ies**] [**vpls**] [**vprn**][**sdp** *sdp-id*] [**customer** *customer-id*]

Context show>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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.

vprn — Displays matching VPRN 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 an customer.

Values 1 — 2147483647

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.

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

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
900          VPRN      Up     Up         2                 11/04/2006 04:55:12
-------------------------------------------------------------------------------
Matching Services : 4
===============================================================================
*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#


A:de14# 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:de14#
```

id

Syntax **id** *service-id* {**all** | **arp** | **base** | **bgp-ad | fdb** | **labels** | **mfib** | **sap** | **sdp** | **split-horizon-group** | **stp**}

Context show>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays information for a particular service-id.

Parameters *service-id* — The unique service identification number that identifies the service in the service domain.

all — Display detailed information about the service.

arp — Display ARP entries for the service.

base — Display basic service information.

bgp-ad — Display BGP AD related information for the service.

fdb — Display FDB entries.

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.

all

Syntax **all**

Context show>service>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Displays detailed information for all aspects of the service.

Output **Sample Output**

```
*A:duth# show service  id 502 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 502                Vpn Id            : 0
```

```
Service Type      : VPRN
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 12/23/2014 04:57:18
Last Mgmt Change  : 12/23/2014 04:57:18
Admin State       : Up                  Oper State        : Up

Route Dist.       : 502:502             VPRN Type         : regular
AS Number         : 100                 Router Id         : 8.8.8.8
ECMP              : Enabled             ECMP Max Routes   : 1

Auto Bind Tunnel
Resolution        : filter
Filter Protocol   : rsvp

Max IPv6 Routes   : No Limit
Ignore NH Metric  : Disabled
Vrf Target        : target:502:502
Vrf Import        : None
Vrf Export        : None
MVPN Vrf Target   : None
MVPN Vrf Import   : None
MVPN Vrf Export   : None
Label mode        : vrf
BGP VPN Backup    : Disabled

SAP Count         : 0                   SDP Bind Count    : 5
Allow IP Intf Bind: Disabled
VSD Domain        : <none>

-------------------------------------------------------------------------------
ETH-CFM service specifics
-------------------------------------------------------------------------------
Tunnel Faults     : ignore


-------------------------------------------------------------------------------
VPRN service Network Specifics
-------------------------------------------------------------------------------
Ing Net QoS Policy : none
Ingress FP QGrp    : (none)                          Ing FP QGrp Inst : (none)
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Sdp Id 502:502  -(9.9.9.9)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 502:502                 Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type           : n/a                     VC Tag            : n/a
Admin Path MTU    : 0                       Oper Path MTU     : 9186
Delivery          : MPLS
Far End           : 9.9.9.9
Tunnel Far End    : n/a                     LSP Types         : RSVP

Admin State       : Up                      Oper State        : Up
Acct. Pol         : None                    Collect Stats     : Disabled
```

```
Ingress Label      : n/a                        Egress Label      : n/a
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
BGP IPv4 FlowSpec  : Disabled
BGP IPv6 FlowSpec  : Disabled
Admin ControlWord  : Not Preferred              Oper ControlWord  : False
BFD Template       : None
BFD-Enabled        : no                         BFD-Encap         : ipv4
Last Status Change : 12/23/2014 18:12:34        Signaling         : n/a
Last Mgmt Change   : 12/21/2014 19:39:56
Class Fwding State : Down
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

Application Profile: None

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.      : n/a                        I. Fwd. Octs.     : n/a
E. Fwd. Pkts.      : n/a                        E. Fwd. Octets    : n/a
Extra-Tag-Drop-Pkts: n/a                        Extra-Tag-Drop-Oc*: n/a

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

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
Lsp Name           : i2
Admin State        : Up                         Oper State        : Up
Time Since Last Tr*: 04d06h44m

-------------------------------------------------------------------------------
Sdp Id 602:502  -(2.2.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 602:502                    Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type            : n/a                        VC Tag            : n/a
Admin Path MTU     : 0                          Oper Path MTU     : 9186
Delivery           : MPLS
Far End            : 2.2.2.2
Tunnel Far End     : n/a                        LSP Types         : RSVP

Admin State        : Up                         Oper State        : Up
Acct. Pol          : None                       Collect Stats     : Disabled
Ingress Label      : n/a                        Egress Label      : n/a
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
BGP IPv4 FlowSpec  : Disabled
BGP IPv6 FlowSpec  : Disabled
Admin ControlWord  : Not Preferred           Oper ControlWord  : False
BFD Template       : None
BFD-Enabled        : no                      BFD-Encap         : ipv4
Last Status Change : 12/23/2014 18:12:24     Signaling         : n/a
Last Mgmt Change   : 12/21/2014 19:39:56
Class Fwding State : Down
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

Application Profile: None

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.      : n/a                     I. Fwd. Octs.     : n/a
E. Fwd. Pkts.      : n/a                     E. Fwd. Octets    : n/a
Extra-Tag-Drop-Pkts: n/a                     Extra-Tag-Drop-Oc*: n/a

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

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
Lsp Name           : b2
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 04d06h44m

-------------------------------------------------------------------------------
Sdp Id 702:502  -(7.7.7.7)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 702:502                 Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type            : n/a                     VC Tag            : n/a
Admin Path MTU     : 0                       Oper Path MTU     : 9186
Delivery           : MPLS
Far End            : 7.7.7.7
Tunnel Far End     : n/a                     LSP Types         : RSVP

Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : n/a                     Egress Label      : n/a
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
BGP IPv4 FlowSpec  : Disabled
BGP IPv6 FlowSpec  : Disabled
```

```
Admin ControlWord  : Not Preferred           Oper ControlWord  : False
BFD Template       : None
BFD-Enabled        : no                      BFD-Encap         : ipv4
Last Status Change : 12/23/2014 18:12:27     Signaling         : n/a
Last Mgmt Change   : 12/21/2014 19:39:56
Class Fwding State : Down
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

Application Profile: None

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.      : n/a                     I. Fwd. Octs.     : n/a
E. Fwd. Pkts.      : n/a                     E. Fwd. Octets    : n/a
Extra-Tag-Drop-Pkts: n/a                     Extra-Tag-Drop-Oc*: n/a

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

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
Lsp Name           : g2
Admin State        : Up                      Oper State        : Up
Time Since Last Tr*: 04d06h44m

-------------------------------------------------------------------------------
Sdp Id 802:502  -(99.99.99.99)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 802:502                 Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type            : n/a                     VC Tag            : n/a
Admin Path MTU     : 0                       Oper Path MTU     : 9182
Delivery           : MPLS
Far End            : 99.99.99.99
Tunnel Far End     : n/a                     LSP Types         : BGP

Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : n/a                     Egress Label      : n/a
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
BGP IPv4 FlowSpec  : Disabled
BGP IPv6 FlowSpec  : Disabled
Admin ControlWord  : Not Preferred           Oper ControlWord  : False
BFD Template       : None
BFD-Enabled        : no                      BFD-Encap         : ipv4
```

```
Last Status Change : 12/23/2014 18:12:41      Signaling         : n/a
Last Mgmt Change   : 12/21/2014 19:39:56
Class Fwding State : Down
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

Application Profile: None

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.      : n/a                      I. Fwd. Octs.     : n/a
E. Fwd. Pkts.      : n/a                      E. Fwd. Octets    : n/a
Extra-Tag-Drop-Pkts: n/a                      Extra-Tag-Drop-Oc*: n/a

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

-------------------------------------------------------------------------------
BGP Tunnel Information :
-------------------------------------------------------------------------------
BGP LSP Id         : 262147

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

-------------------------------------------------------------------------------
Sdp Id 902:502  -(22.22.22.22)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 902:502                  Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type            : n/a                      VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 9182
Delivery           : MPLS
Far End            : 22.22.22.22
Tunnel Far End     : n/a                      LSP Types         : BGP

Admin State        : Up                       Oper State        : Up
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : n/a                      Egress Label      : n/a
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
BGP IPv4 FlowSpec  : Disabled
BGP IPv6 FlowSpec  : Disabled
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
BFD Template       : None
BFD-Enabled        : no                       BFD-Encap         : ipv4
```

```
Last Status Change : 12/23/2014 19:25:37      Signaling         : n/a
Last Mgmt Change   : 12/21/2014 19:39:56
Class Fwding State : Down
Flags              : None
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

Application Profile: None

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.      : n/a                     I. Fwd. Octs.     : n/a
E. Fwd. Pkts.      : n/a                     E. Fwd. Octets    : n/a
Extra-Tag-Drop-Pkts: n/a                     Extra-Tag-Drop-Oc*: n/a

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

-------------------------------------------------------------------------------
BGP Tunnel Information :
-------------------------------------------------------------------------------
BGP LSP Id         : 262387

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

-------------------------------------------------------------------------------
Number of SDPs : 5
-------------------------------------------------------------------------------
* indicates that the corresponding row element may have been truncated.
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
No Sap Associations

-------------------------------------------------------------------------------
Service Interfaces
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Interface
-------------------------------------------------------------------------------
If Name            : toixia1
Admin State        : Up                  Oper (v4/v6)      : Up/Down
Protocols          : None
IP Addr/mask       : 108.2.1.1/24        Address Type      : Primary
IGP Inhibit        : Disabled            Broadcast Address : Host-ones
HoldUp-Time        : 0                   Track Srrp Inst   : 0
Description        : N/A
```

```
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Description       : (Not Specified)
If Index          : 2                    Virt. If Index    : 2
Last Oper Chg     : 12/23/2014 20:36:08 Global If Index   : 48
Mon Oper Grp      : None
Srrp En Rtng      : Disabled             Hold time         : N/A
Port Id           : rvpls
TOS Marking       : Trusted              If Type           : VPRN
SNTP B.Cast       : False
MAC Address       : 00:00:0b:00:01:00    Mac Accounting    : Disabled
Ingress stats     : Disabled             IPv6 DAD          : Enabled
ARP Timeout       : 14400s               IPv6 Nbr ReachTime: 30s
ARP Retry Timer   : 5000ms               IPv6 stale time   : 14400s
ARP Limit         : Disabled             IPv6 Nbr Limit    : Disabled
ARP Threshold     : Disabled             IPv6 Nbr Threshold: Disabled
ARP Limit Log Only: Disabled             IPv6 Nbr Log Only : Disabled
IP MTU            : (default)
IP Oper MTU       : 1500
LdpSyncTimer      : None
LSR Load Balance  : system
EGR Load Balance  : both
Vas If Type       : none
TEID Load Balance : Disabled
SPI Load Balance  : Disabled
uRPF Chk          : disabled
uRPF Ipv6 Chk     : disabled
PTP HW Assist     : Disabled
Mpls Rx Pkts      : 0                    Mpls Rx Bytes     : 0
Mpls Tx Pkts      : 0                    Mpls Tx Bytes     : 0

DHCP6 Relay Details
Description       : (Not Specified)
Admin State       : Down                 Lease Populate    : 0
Oper State        : Down                 Nbr Resolution    : Disabled
If-Id Option      : None                 Remote Id         : Disabled
Src Addr          : Not configured
Python plcy       : (Not Specified)

DHCP6 Server Details
Admin State       : Down                 Max. Lease States : 8000

ICMP Details
Redirects    : Number - 100                    Time (seconds)   - 10
Unreachables : Number - 100                    Time (seconds)   - 10
TTL Expired  : Number - 100                    Time (seconds)   - 10

Routed VPLS Details
VPLS Name        : rvpls-502
Binding Status   : Up
Reason           : (Not Specified)
Egr Reclass Plcy : 0
Ing Filter       : none
Ingr IPv6 Flt    : none
EVPN Tunnel      : false

-------------------------------------------------------------------------------
Admin Groups
-------------------------------------------------------------------------------
No Matching Entries
```

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

-------------------------------------------------------------------------------
Srlg Groups
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Interface
-------------------------------------------------------------------------------
If Name           : toB-Ixia
Admin State       : Up                  Oper (v4/v6)      : Up/Down
Protocols         : OSPFv2
IP Addr/mask      : 200.41.2.8/24       Address Type      : Primary
IGP Inhibit       : Disabled            Broadcast Address : Host-ones
HoldUp-Time       : 0                   Track Srrp Inst   : 0
Description       : N/A
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Description       : (Not Specified)
If Index          : 3                   Virt. If Index    : 3
Last Oper Chg     : 12/23/2014 20:00:55 Global If Index   : 49
Mon Oper Grp      : None
Srrp En Rtng      : Disabled            Hold time         : N/A
Port Id           : rvpls
TOS Marking       : Trusted             If Type           : VPRN
SNTP B.Cast       : False
MAC Address       : c4:08:4a:6f:99:65   Mac Accounting    : Disabled
Ingress stats     : Disabled            IPv6 DAD          : Enabled
ARP Timeout       : 14400s              IPv6 Nbr ReachTime: 30s
ARP Retry Timer   : 5000ms              IPv6 stale time   : 14400s
ARP Limit         : Disabled            IPv6 Nbr Limit    : Disabled
ARP Threshold     : Disabled            IPv6 Nbr Threshold: Disabled
ARP Limit Log Only: Disabled            IPv6 Nbr Log Only : Disabled
IP MTU            : (default)
IP Oper MTU       : 1500
LdpSyncTimer      : None
LSR Load Balance  : system
EGR Load Balance  : both
Vas If Type       : none
TEID Load Balance : Disabled
SPI Load Balance  : Disabled
uRPF Chk          : disabled
uRPF Ipv6 Chk     : disabled
PTP HW Assist     : Disabled
Mpls Rx Pkts      : 0                   Mpls Rx Bytes     : 0
Mpls Tx Pkts      : 0                   Mpls Tx Bytes     : 0

DHCP6 Relay Details
Description       : (Not Specified)
Admin State       : Down                Lease Populate    : 0
Oper State        : Down                Nbr Resolution    : Disabled
If-Id Option      : None                Remote Id         : Disabled
Src Addr          : Not configured
Python plcy       : (Not Specified)

DHCP6 Server Details
Admin State       : Down                Max. Lease States : 8000
```

```
ICMP Details
Redirects    : Number - 100                    Time (seconds)   - 10
Unreachables : Number - 100                    Time (seconds)   - 10
TTL Expired  : Number - 100                    Time (seconds)   - 10

Routed VPLS Details
VPLS Name        : rvpls-622
Binding Status   : Up
Reason           : (Not Specified)
Egr Reclass Plcy : 0
Ing Filter       : none
Ingr IPv6 Flt    : none
EVPN Tunnel      : false

-------------------------------------------------------------------------------
Admin Groups
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Srlg Groups
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
===============================================================================
*A:duth#
```

Show All Service-ID Output — The following table describes the show all service-id command output fields:

Label	Description
Service Detailed Information	
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.
Customer Id	The customer identifier.
Last Status Change	The date and time of the most recent change in the administrative or operating status of the service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Admin State	The current administrative state.
Oper State	The current operational state.
Route Dist.	Displays the route distribution number.
AS Number	Displays the autonomous system number.
Router Id	Displays the router ID for this service.

Label	Description
Auto Bind	Specifies the automatic binding type for the SDP assigned to this service.
Vrf Target	Specifies the VRF target applied to this service.
Vrf Import	Specifies the VRF import policy applied to this service.
Vrf Export	Specifies the VRF export policy applied to this service.
Description	Generic information about the service.
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.
Service Destination Points (SDPs)	
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.
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 MPLS tunnel defined by this SDP.

Label	Description
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the keepalive protocol.
Oper State	The current status of the keepalive protocol.
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.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Number of SDPs	The total number SDPs applied to this service ID.
Service Access Points	
Service Id	The service identifier.
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Admin State	The desired 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
Ingress qos-policy	The SAP ingress QoS policy ID.
Egress qos-policy	The SAP egress QoS policy ID. This command is applicable only to 7210 SAS X.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
Spoke SDPs	
Managed by Service	Specifies the service-id of the management VPLS managing this spoke SDP.
Managed by Spoke	Specifies the sap-id inside the management VPLS managing this spoke SDP.
Prune state	Specifies the STP state inherited from the management VPLS.
Peer Pw Bits	Indicates the bits set by the LDP peer when there is a fault on its side of the pseudowire. LAC failures occur on the SAP that has been configured on the pipe service, PSN bits are set by SDP-binding failures on the pipe service. The pwNotForwarding bit is set when none of the above failures apply, such as an MTU mismatch failure. This value is only applicable if the peer is using the pseudowire status signalling method to indicate faults. 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
Max IPv4 Routes	Maximum IPv4 routes configured for use with the service.
Last Changed	The date and time of the most recent management-initiated change.
Dot1Q Ethertype	The Dot1q ethertype in use by the SAP.
Ingr IP Fltr-Id	The policy ID of the IP filter applied at ingress.
Ingr Mac Fltr-Id	The policy ID of the MAC filter applied at ingress.
Egr IP Fltr-Id	The policy ID of the IP filter applied at egress.
Egr Mac Fltr-Id	The policy ID of the MAC filter applied at egress.
tod-suite	The TOD suite applied for use by this SAP.
rate	Specifies the SAP aggregate rate configured for the aggregate policer/meter used by this SAP.
burst	Specifies the burst to be used with SAP aggregate policer/meter used by this SAP.

Label	Description
Classifiers Allocated	Number of SAP ingress QoS resources allocated for use by this SAP.
Classifiers Used	Number of SAP ingress QoS resources in use by this SAP.
Meters Allocated	Number of SAP ingress meter resources allocated for use by this SAP. This is set to half the number of classifiers allocated to this SAP.
Meters Used	Number of SAP ingress meters in use.
Ingress Stats	The number of received packets/octets for this SAP.
Egress Stats	The number of packets/octets forwarded out of this SAP.
Ingress Drop Stats	Number of packets/octets dropped by the system.
Extra-Tag Drop Stats	Number of packets received with the count of VLAN tags exceeding the count of VLAN tags implied by the SAP encapsulation.
Ingress Meter 1	The index of the ingress QoS meter of this SAP.
For. InProf	Number of in-profile packets/octets received on this SAP.
For. OutProf	Number of out-of-profile packets/octets received on this SAP.
If Name	IP interface name assigned by user.
Protocols	Protocols enabled for use on this interface.
Oper (v4/v6)	Operational status of this interface for IPv4 and IPv6.
IP Addr/mask	IPv4 address and Mask assigned to this interface.
Address Type	Whether the address is a primary or secondary address.
Broadcast Address	Type of broadcast address used. It can be host-ones or all-ones.
If Index	The interface Index assigned by the system. It is used with SNMP IfTable.
Virt. If Index	The interface index assigned by the system. It is used with SNMP.
Last Oper Chg	Timestamp associated with the last operational change.
Global If Index	This is the system wide Interface index allotted by the system.
If Type	Network — The IP interface is a network/core IP interface. Service — The IP interface is a service IP interface.
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled.
Arp Timeout	Specifies the timeout for an ARP entry learned on the interface.
IP Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the port to the far-end router, without requiring the packet to be fragmented.

Label	Description
LdpSyncTimer	Specifies the value used for IGP-LDP synchronization.
Redirects	Specifies the rate for ICMP redirect messages.
Unreachables	Specifies the rate for ICMP unreachable messages.
TTL Expired	Specifies the rate for ICMP TTL messages.
MAC Address	Specifies the 48-bit IEEE 802.3 MAC address.

authentication

Syntax **authentication**

Context show>service>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context to display subscriber authentication information.

statistics

Syntax **statistics** [**policy** *name*] [**sap** *sap-id*]

Context show>service>id>authentication

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays session authentication statistics for this service.

Parameters **policy** *name* — Specifies the subscriber authentication policy statistics to display.

sap *sap-id* — Specifies the SAP ID statistics to display. See Common CLI Command Descriptions on page 945 for command syntax.

Sample Output

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

arp

Syntax **arp** [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*] [**sdp** *sdp-id:vc-id*] [**summary**]

Context show>service>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Displays the ARP table for the IES instance.

Parameters *ip-address* — Displays only ARP entries in the ARP table with the specified IP address.

Default All IP addresses.

mac *ieee-address* — Displays only ARP entries in the ARP table with the specified 48-bit MAC address. The MAC address can be expressed in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

Default All MAC addresses.

sap *sap-id* — Displays SAP information for the specified SAP ID. See Common CLI Command Descriptions on page 945 for command syntax.

port id — Specifies matching service ARP entries associated with the specified 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
Service ID	The service ID number.
MAC	The specified MAC address
Source-Identifier	The location the MAC is defined.
Type	Static – FDB entries created by management. Learned – Dynamic entries created by the learning process. OAM – Entries created by the OAM process.
Age	The time elapsed since the service was enabled.
Interface	The interface applied to the service.
Port	The port where the SAP is applied.

Sample Output

```
*A:ALA-12# show service id 2 arp
===============================================================================
ARP Table
```

```
===============================================================================
IP Address      MAC Address       Type    Age       Interface         Port
-------------------------------------------------------------------------------
190.11.1.1      00:03:fa:00:08:22 Other   00:00:00  ies-100-190.11.1  1/1/11:0
===============================================================================
*A:ALA-12#
```

base

Syntax **base**

Context show>service>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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	Specifies the type of service.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
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.

Label	Description
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.
OprMTU	Specifies the actual 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.
Opr	The operating state of the SDP.

Sample Output

```
*A:ALA-12# show service id 1 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPRN
Customer Id       : 1
Last Status Change: 02/01/2007 09:11:39
Last Mgmt Change  : 02/01/2007 09:11:46
Admin State       : Up                  Oper State        : Down
Route Dist.       : 10001:1
AS Number         : 10000               Router Id         : 10.10.10.103
ECMP              : Enabled             ECMP Max Routes   : 8
Max Routes        : No Limit            Auto Bind         : LDP
Vrf Target        : target:10001:1
Vrf Import        : vrfImpPolCust1
Vrf Export        : vrfExpPolCust1
SAP Count         : 1                   SDP Bind Count    : 18
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                      Type         AdmMTU  OprMTU  Adm     Opr
-------------------------------------------------------------------------------
sap:1/1/7:0                     q-tag        1518    1518    Up      Up
sdp:10:1 M(10.20.1.3)           TLDP         4462    4462    Up      TLDP Down
sdp:20:1 M(10.20.1.4)           TLDP         4462    4462    Up      TLDP Down
sdp:30:1 M(10.20.1.5)           TLDP         4462    4462    Up      TLDP Down
sdp:40:1 M(10.20.1.20)          TLDP         1534    4462    Up      Up
sdp:200:1 M(10.20.1.30)         TLDP         1514    4462    Up      Up
sdp:300:1 M(10.20.1.31)         TLDP         4462    4462    Up      TLDP Down
sdp:500:1 M(10.20.1.50)         TLDP         4462    4462    Up      TLDP Down
===============================================================================
*A:ALA-12#
```

statistics

Syntax **statistics** [**sap** *sap-id*]
statistics [**sdp** *sdp-id*:*vc-id*]

statistics [**interface** *interface-name*]

Context show>service>id>dhcp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Displays DHCP statistics information.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 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 *interface-name* — Displays information for the specified IP interface.

Show DHCP Statistics Output — The following table describes the output fields for DHCP statistics.

Label	Description
Received Packets	The number of packets received from the DHCP clients.
Transmitted Packets	The number of packets transmitted to the DHCP clients.
Received Malformed Packets	The number of corrupted/invalid packets received from the DHCP clients.
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients. In this case, a frame is dropped due to the client sending a DHCP packet with Option 82 filled in before “trust” is set under the DHCP interface command.
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded.
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded.
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Server Packets Discarded	The number of packets received from the DHCP server that were discarded.
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded.
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.

Sample Output

```
A:sim1# show service id 11 dhcp statistics
===============================================================
DHCP Global Statistics, service 11
===============================================================
Rx Packets                        : 32
Tx Packets                        : 12
Rx Malformed Packets              : 0
Rx Untrusted Packets              : 0
Client Packets Discarded          : 0
Client Packets Relayed            : 11
Client Packets Snooped            : 21
Server Packets Discarded          : 0
Server Packets Relayed            : 0
Server Packets Snooped            : 0
===============================================================
A:sim1#
```

interface

Syntax **interface** [*ip-address* | *ip-int-name*] [**detail**]

Context show>service>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Displays information for the IP interfaces associated with the service.

If no optional parameters are specified, a summary of all IP interfaces associated to the service are displayed.

Parameters *ip-address* — The IP address of the interface for which to display information.

Values 1.0.0.0 — 223.255.255.255

ip-int-name — The IP interface name for which to display information.

detail — Displays detailed IP interface information.

Default IP interface summary output.

Output **Show Service-ID Interface —** The following table describes show service-id interface output fields:

Label	Description
Interface-Name	The name used to refer to the interface.
Type	Specifies the interface type.
IP-Address	Specifies the IP address/IP subnet/broadcast address of the interface.
Adm	The desired state of the interface.
Opr	The operating state of the interface.

Label	Description (Continued)
Interface	
If Name	The name used to refer to the interface.
Admin State	The desired state of the interface.
Oper State	The operating state of the interface.
IP Addr/mask	Specifies the IP address/IP subnet/broadcast address of the interface.
Details	
If Index	The index corresponding to this interface. The primary index is 1. For example, all interfaces are defined in the Base virtual router context.
If Type	Specifies the interface type.
Port Id	Specifies the SAP's port ID.
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled.
Arp Timeout	Specifies the timeout for an ARP entry learned on the interface.
MAC Address	Specifies the 48-bit IEEE 802.3 MAC address.
ICMP Mask Reply	Specifies whether ICMP mask reply is enabled or disabled.
ICMP Details	
Redirects	Specifies the rate for ICMP redirect messages.
Unreachables	Specifies the rate for ICMP unreachable messages.
TTL Expired	Specifies the rate for ICMP TTL messages.

Sample Output

```
*A:ALA-12# show service id 321 interface
===============================================================================
Interface Table
===============================================================================
Interface-Name                   Type IP-Address         Adm    Opr  Type
-------------------------------------------------------------------------------
test                             Pri  190.11.1.1/24      Up     Up   IES
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
*A:ALA-12#


A:ALA-49# show service id 88 interface detail
===============================================================================
Interface Table
===============================================================================
Interface
-------------------------------------------------------------------------------
If Name      : Sector A
```

```
Admin State  : Up                                Oper State       : Down
Protocols    : None

IP Addr/mask : Not Assigned
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Description  :
If Index     : 26                                Virt. If Index   : 26
SAP Id       : 71/1/1.2.2
TOS Marking  : Untrusted                         If Type          : IES
SNTP B.Cast  : False                             IES ID           : 88
MAC Address  : Not configured.                   Arp Timeout      : 14400
IP MTU       : 1500                              ICMP Mask Reply  : True
Arp Populate : Disabled
Cflowd       : None

Proxy ARP Details
Proxy ARP    : Enabled                           Local Proxy ARP  : Disabled
Policies     : ProxyARP

DHCP Details
Admin State  : Up                                Lease Populate   : 0
Action       : Keep                              Trusted          : Disabled
ICMP Details
Redirects    : Number - 100                      Time (seconds)   - 10
Unreachables : Number - 100                      Time (seconds)   - 10
TTL Expired  : Number - 100                      Time (seconds)   - 10
-------------------------------------------------------------------------------
Interface
-------------------------------------------------------------------------------
If Name      : test
Admin State  : Up                                Oper State       : Down
Protocols    : None
IP Addr/mask : Not Assigned
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Description  :
If Index     : 27                                Virt. If Index   : 27
SAP Id       : 101/1/2:0
TOS Marking  : Untrusted                         If Type          : IES
SNTP B.Cast  : False                             IES ID           : 88
MAC Address  : Not configured.                   Arp Timeout      : 14400
Arp Populate : Disabled

Proxy ARP Details
Proxy ARP    : Disabled                          Local Proxy ARP  : Disabled

ICMP Details
Redirects    : Number - 100                      Time (seconds)   - 10
Unreachables : Number - 100                      Time (seconds)   - 10
TTL Expired  : Number - 100                      Time (seconds)   - 10
-------------------------------------------------------------------------------
Interfaces : 2
===============================================================================
A:ALA-49#
```

sap

Syntax **sap** *sap-id* [**detail**]]

Context show>service>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Displays information for the SAPs associated with the service.

If no optional parameters are specified, a summary of all associated SAPs is displayed.

Parameters *sap-id* — The ID that displays SAPs for the service. See Common CLI Command Descriptions on page 945 for command syntax.

detail — Displays detailed information for the SAP.

Output **Show Service-ID SAP —** The following table describes show service SAP fields:

Sample Output

Label	Description
Service Id	The service identifier.
SAP	The SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch,RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
Last Status Change	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 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.

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.
Dropped	The number of packets and octets dropped due to SAP state, ingress MAC or IP filter, same segment discard, bad checksum, etc.
Off. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Dro. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip 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 by the Qchip due to: MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the ingress Qchip.
For. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
Dro. InProf	The number of in-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
Dro. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, etc.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the egress Qchip.
For. OutProf	The number of out-of-profile packets and octets (rate above CIR) forwarded by the egress Qchip.

Label	Description (Continued)
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.

```
*A:ALA-12# show service id 321 sap 1/1/4:0
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 321
SAP                : 1/1/4:0                 Encap             : q-tag
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100
Admin State        : Up                       Oper State        : Down
Flags              : PortOperDown
                     SapIngressQoSMismatch
Last Status Change : 02/03/2007 12:58:37
Last Mgmt Change   : 02/03/2007 12:59:10
Admin MTU          : 1518                     Oper MTU          : 1518
Ingress qos-policy : 100                      Egress qos-policy : 1
Ingress Filter-Id  : n/a                      Egress Filter-Id  : n/a
Multi Svc Site     : None
Acct. Pol          : None                     Collect Stats     : Disabled
===============================================================================
*A:ALA-12#


*A:ALA-12# show service id 321 sap 1/1/4:0 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 321
SAP                : 1/1/4:0                 Encap             : q-tag
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Admin State        : Up                       Oper State        : Down
Flags              : PortOperDown
                     SapIngressQoSMismatch
Last Status Change : 02/03/2007 12:58:37
Last Mgmt Change   : 02/03/2007 12:59:10
Admin MTU          : 1518                     Oper MTU          : 1518
Ingress qos-policy : 100                      Egress qos-policy : 1
Ingress Filter-Id  : n/a                      Egress Filter-Id  : n/a
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 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
'''
===============================================================================
*A:ALA-12#
```

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail**]

Context show>service>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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 a spoke or a mesh.
Split Horizon Group	Name of the split horizon group that the SDP belongs to.
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 MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.

Label	Description (Continued)
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	he 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.
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.

Sample Output

```
A:Dut-A# show service id  1 sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
 Sdp Id 1:1  -(10.20.1.2)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id            : 1:1                       Type              : Spoke
VC Type           : Ether                     VC Tag            : n/a
Admin Path MTU    : 0                         Oper Path MTU     : 9186
Far End           : 10.20.1.2                 Delivery          : MPLS
```

```
Admin State        : Up                       Oper State        : Up
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 2048                     Egress Label      : 2048
Ing mac Fltr       : n/a                      Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                      Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                      Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 05/31/2007 00:45:43      Signaling         : None
Last Mgmt Change   : 05/31/2007 00:45:43
Class Fwding State : Up
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0

MAC Learning       : Enabled                  Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled
MAC Pinning        : Disabled

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

Statistics           :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0
MCAC Policy Name   :
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited

Associated LSP LIST :
Lsp Name           : A_B_1
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m35s

Lsp Name           : A_B_2
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m35s

Lsp Name           : A_B_3
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_4
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_5
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_6
Admin State        : Up                       Oper State        : Up
Time Since Last Tr*: 00h26m34s
```

```
Lsp Name           : A_B_7
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_8
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 00h26m35s

Lsp Name           : A_B_9
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 00h26m34s

Lsp Name           : A_B_10
Admin State        : Up                        Oper State        : Up
Time Since Last Tr*: 00h26m34s
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : enabled
Default LSP        : A_B_10                    Multicast LSP     : A_B_9
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
af                  A_B_3
be                  A_B_1
ef                  A_B_6
h1                  A_B_7
h2                  A_B_5
l1                  A_B_4
l2                  A_B_2
nc                  A_B_8
===============================================================================
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Up                        Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                       Port State        : Forwarding
Port Number        : 2049                      Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : N/A
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                       Designated Port Id: 0

Fwd Transitions    : 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
* indicates that the corresponding row element may have been truncated.
-------------------------------------------------------------------------------
A:Dut-A#
```

aggregate

Syntax **aggregate** [**active**]

Context show>router

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays aggregated routes.

Parameters **active** — This keyword filters out inactive aggregates.

Output **Show Aggregate Output Fields —** The following table describes router aggregate output fields.

Label	Description
Prefix	Displays the destination address of the aggregate route in dotted decimal notation.
Summary	Specifies whether the aggregate or more specific components are advertised.
AS Set	Displays an aggregate where the path advertised for the route consists of all elements contained in all paths that are being summarized.
Aggr AS	Displays the aggregator path attribute to the aggregate route.
Aggr IP-Address	The IP address of the aggregated route.
State	The operational state of the aggregated route.
No. of Aggregates	The total number of aggregated routes.

Sample Output

```
*A:ALA-12# show router 3 aggregate
===============================================================================
Aggregates (Service: 3)
===============================================================================
Prefix                Summary  AS Set   Aggr AS     Aggr IP-Address   State
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of Aggregates: 0
-------------------------------------------------------------
*A:ALA-12#
```

arp

Syntax **arp** [*ip-address* | *ip-int-name* | **mac** *ieee-mac-addr*]

Context show>router

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays the router ARP table sorted by IP address.

If no command line options are specified, all ARP entries are displayed.

Parameters *ip-addr* — Only displays ARP entries associated with the specified IP address.

ip-int-name — Only displays ARP entries associated with the specified IP interface name.

mac *ieee-mac-addr* **—** Only displays ARP entries associated with the specified MAC address.

Output **ARP Table Output —** The following table describes ARP table output fields:

Label	Description
IP Address	The IP address of the ARP entry.
MAC Address	The MAC address of the ARP entry.
Expiry	The age of the ARP entry.
Type	Dyn — The ARP entry is a dynamic ARP entry.
	Inv — The ARP entry is an inactive static ARP entry (invalid).
	Oth — The ARP entry is a local or system ARP entry.
	Sta — The ARP entry is an active static ARP entry.
Interface	The IP interface name associated with the ARP entry.
No. of ARP Entries	The number of ARP entries displayed in the list.

Sample Output

```
*A:ALA-12# show router 3 arp
===============================================================================
ARP Table (Service: 3)
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.10.10.103    04:67:ff:00:00:01 00h00m00s Oth    system
10.10.4.3       00:00:00:00:00:00 00h00m00s Oth    ALA-1-2
10.10.5.3       00:00:00:00:00:00 00h00m00s Oth    ALA-1-3
10.10.7.3       00:00:00:00:00:00 00h00m00s Oth    ALA-1-5
10.10.0.16      00:00:00:00:00:00 00h00m00s Oth    bozo
10.10.3.3       00:00:00:00:00:00 00h00m00s Oth    gizmo
10.10.2.3       00:00:00:00:00:00 00h00m00s Oth    hobo
10.10.1.17      00:00:00:00:00:00 00h00m00s Oth    int-cflowd
10.0.0.92       00:00:00:00:00:00 04h00m00s Dyn    to-104
10.0.0.103      04:67:01:01:00:01 00h00m00s Oth[I] to-104
10.0.0.104      04:68:01:01:00:01 03h59m49s Dyn[I] to-104
10.10.36.2      00:00:00:00:00:00 00h00m00s Oth    tuesday
192.168.2.98    00:03:47:c8:b4:86 00h14m37s Dyn[I] management
192.168.2.103   00:03:47:dc:98:1d 00h00m00s Oth[I] management
-------------------------------------------------------------------------------
No. of ARP Entries: 14
===============================================================================
*A:ALA-12#


*A:ALA-12# show router 3 arp 10.10.0.3
```

```
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.10.0.3       04:5d:ff:00:00:00 00:00:00  Oth    system
===============================================================================
*A:ALA-12#


*A:ALA-12# show router 3 arp to-ser1
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Expiry    Type   Interface
-------------------------------------------------------------------------------
10.10.13.1      04:5b:01:01:00:02 03:53:09  Dyn    to-ser1
===============================================================================
*A:ALA-12#
```

damping

Syntax **damping** [*ip-prefix/mask* | *ip-address*] [**detail**]
damping [*damp-type*] [**detail**]

Context show>router>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays BGP routes with have been dampened due to route flapping. This command can be entered with or without a route parameter.

When the keyword **detail** is included, more detailed information displays.

When only the command is entered (without any parameters included except **detail**), then all dampened routes are listed.

When a parameter is specified, then the matching route or routes are listed.

When a **decayed**, **history,** or **suppressed** keyword is specified, only those types of dampened routes are listed.

Parameters *ip-prefix/mask* — Displays damping information for the specified IP prefix and mask length.

ip-address — Displays damping entry for the best match route for the specified IP address.

damp-type — Displays damping type for the specified IP address.

decayed — Displays damping entries that are decayed but are not suppressed.

history — Displays damping entries that are withdrawn but have history.

suppressed — Displays damping entries suppressed because of route damping.

detail — Displays detailed information.

Output **Show Damping Output Fields —** The following table describes BGP damping output fields:

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured or inherited local AS for the specified peer group. If not configured, then it is the same value as the AS.
Network	Route IP prefix and mask length for the route.
Flag(s)	Legend: Status codes: u- used, s-suppressed, h-history, d-decayed, *-valid. If a * is not present, then the status is invalid. Origin codes: i-IGP, e-EGP, ?-incomplete, >-best
Network	The IP prefix and mask length for the route.
From	The originator ID path attribute value.

Label	Description (Continued)
Reuse time	The time when a suppressed route can be used again.
AS Path	The BGP AS path for the route.
Peer	The router ID of the advertising router.
NextHop	BGP nexthop for the route.
Peer AS	The autonomous system number of the advertising router.
Peer Router-Id	The router ID of the advertising router.
Local Pref	BGP local preference path attribute for the route.
Age	The time elapsed since the service was enabled.
Last update	The time when BGP was updated last in second/minute/hour (SS:MM:HH) format.
FOM Present	The current Figure of Merit (FOM) value.
Number of Flaps	The number of flaps in the neighbor connection.
Reuse time	The time when the route can be reused.
Path	The BGP AS path for the route.
Applied Policy	The applied route policy name.

Sample Output

```
*A:ALA-12# show router 3 bgp damping
===============================================================================
  BGP Router ID : 10.0.0.14         AS : 65206   Local AS : 65206
===============================================================================
  Legend -
  Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
  Origin codes  : i - IGP, e - EGP, ? - incomplete,  - best
===============================================================================
BGP Damped Routes
===============================================================================
Flag  Network            From            Reuse       AS-Path
-------------------------------------------------------------------------------
ud*i  12.149.7.0/24      10.0.28.1       00h00m00s   60203 65001 19855 3356
                                                      1239  22406
si    24.155.6.0/23      10.0.28.1       00h43m41s   60203 65001 19855 3356
                                                      2914  7459
si    24.155.8.0/22      10.0.28.1       00h38m31s   60203 65001 19855 3356
                                                      2914  7459
si    24.155.12.0/22     10.0.28.1       00h35m41s   60203 65001 19855 3356
                                                      2914  7459
si    24.155.22.0/23     10.0.28.1       00h35m41s   60203 65001 19855 3356
                                                      2914  7459
si    24.155.24.0/22     10.0.28.1       00h35m41s   60203 65001 19855 3356
                                                      2914  7459
si    24.155.28.0/22     10.0.28.1       00h34m31s   60203 65001 19855 3356
```

```
                                                        2914  7459
si    24.155.40.0/21     10.0.28.1       00h28m24s   60203 65001 19855 3356
                                                        7911  7459
si    24.155.48.0/20     10.0.28.1       00h28m24s   60203 65001 19855 3356
                                                        7911  7459
ud*i  61.8.140.0/24      10.0.28.1       00h00m00s   60203 65001 19855 3356
                                                        4637  17447
ud*i  61.8.141.0/24      10.0.28.1       00h00m00s   60203 65001 19855 3356
                                                        4637  17447
ud*i  61.9.0.0/18        10.0.28.1       00h00m00s   60203 65001 19855 3356
                                                        3561  9658  6163
. . .
ud*i  62.213.184.0/23    10.0.28.1       00h00m00s   60203 65001 19855 3356
                                                        6774  6774  9154
-------------------------------------------------------------------------------
*A:ALA-12#
```

```
*A:ALA-12# show router 3 bgp damping detail
===============================================================================
  BGP Router ID : 10.0.0.14         AS : 65206   Local AS : 65206
===============================================================================
  Legend -
  Status codes  : u - used, s - suppressed, h - history, d - decayed, * -
 valid
  Origin codes  : i - IGP, e - EGP, ? - incomplete,  - best
===============================================================================
BGP Damped Routes
===============================================================================
-------------------------------------------------------------------------------
Network : 12.149.7.0/24
-------------------------------------------------------------------------------
Network          : 12.149.7.0/24        Peer             : 10.0.28.1
NextHop          : 10.0.28.1            Reuse time       : 00h00m00s
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h22m09s            Last update      : 02d00h58m
FOM Present      : 738                  FOM Last upd.    : 2039
Number of Flaps  : 2                    Flags            : ud*i
Path             : 60203 65001 19855 3356  1239  22406
Applied Policy   : default-damping-profile
-------------------------------------------------------------------------------
Network : 15.142.48.0/20
-------------------------------------------------------------------------------
Network          : 15.142.48.0/20       Peer             : 10.0.28.1
NextHop          : 10.0.28.1            Reuse time       : 00h00m00s
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h00m38s            Last update      : 02d01h20m
FOM Present      : 2011                 FOM Last upd.    : 2023
Number of Flaps  : 2                    Flags            : ud*i
Path             : 60203 65001 19855 3356  3561  5551  1889
Applied Policy   : default-damping-profile
-------------------------------------------------------------------------------
Network : 15.200.128.0/19
-------------------------------------------------------------------------------
Network          : 15.200.128.0/19      Peer             : 10.0.28.1
NextHop          : 10.0.28.1            Reuse time       : 00h00m00s
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h00m38s            Last update      : 02d01h20m
```

```
FOM Present      : 2011                 FOM Last upd.    : 2023
Number of Flaps  : 2                    Flags            : ud*i
Path             : 60203 65001 19855 1299  702   1889
Applied Policy   : default-damping-profile
-------------------------------------------------------------------------------
Network : 15.203.192.0/18
-------------------------------------------------------------------------------
Network          : 15.203.192.0/18      Peer             : 10.0.28.1
NextHop          : 10.0.28.1            Reuse time       : 00h00m00s
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h00m07s            Last update      : 02d01h20m
FOM Present      : 1018                 FOM Last upd.    : 1024
Number of Flaps  : 1                    Flags            : ud*i
Path             : 60203 65001 19855 1299  702   1889
Applied Policy   : default-damping-profile
-------------------------------------------------------------------------------
*A:ALA-12#
```

```
*A:ALA-12# show router 3 bgp damping 15.203.192.0/18 detail
===============================================================================
  BGP Router ID : 10.0.0.14         AS : 65206   Local AS : 65206
===============================================================================
  Legend -
  Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
  Origin codes  : i - IGP, e - EGP, ? - incomplete,  - best
===============================================================================
BGP Damped Routes 15.203.192.0/18
===============================================================================
Network : 15.203.192.0/18
-------------------------------------------------------------------------------
Network          : 15.203.192.0/18      Peer             : 10.0.28.1
NextHop          : 10.0.28.1            Reuse time       : 00h00m00s
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h00m42s            Last update      : 02d01h20m
FOM Present      : 2003                 FOM Last upd.    : 2025
Number of Flaps  : 2                    Flags            : ud*i
Path             : 60203 65001 19855 3356  702   1889
Applied Policy   : default-damping-profile
-------------------------------------------------------------------------------
Paths : 1
===============================================================================
*A:ALA-12#
*A:ALA-12# show router 3 bgp damping suppressed detail
===============================================================================
  BGP Router ID : 10.0.0.14         AS : 65206   Local AS : 65206
===============================================================================
  Legend -
  Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
  Origin codes  : i - IGP, e - EGP, ? - incomplete,  - best
===============================================================================
BGP Damped Routes (Suppressed)
===============================================================================
Network : 15.142.48.0/20
-------------------------------------------------------------------------------
Network          : 15.142.48.0/20       Peer             : 10.0.28.1
NextHop          : 10.0.28.1            Reuse time       : 00h29m22s
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
```

```
Age              : 00h01m28s            Last update      : 02d01h20m
FOM Present      : 2936                 FOM Last upd.    : 3001
Number of Flaps  : 3                    Flags            : si
Path             : 60203 65001 19855 3356  702   1889
Applied Policy   : default-damping-profile
-------------------------------------------------------------------------------
Network : 15.200.128.0/19
-------------------------------------------------------------------------------
Network          : 15.200.128.0/19      Peer             : 10.0.28.1
NextHop          : 10.0.28.1            Reuse time       : 00h29m22s
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h01m28s            Last update      : 02d01h20m
FOM Present      : 2936                 FOM Last upd.    : 3001
Number of Flaps  : 3                    Flags            : si
Path             : 60203 65001 19855 3356  702   1889
Applied Policy   : default-damping-profile
-------------------------------------------------------------------------------
Network : 15.203.240.0/20
-------------------------------------------------------------------------------
Network          : 15.203.240.0/20      Peer             : 10.0.28.1
NextHop          : 10.0.28.1            Reuse time       : 00h29m22s
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h01m28s            Last update      : 02d01h20m
FOM Present      : 2936                 FOM Last upd.    : 3001
Number of Flaps  : 3                    Flags            : si
Path             : 60203 65001 19855 3356  702   1889
Applied Policy   : default-damping-profile
-------------------------------------------------------------------------------
Network : 15.206.0.0/17
-------------------------------------------------------------------------------
Network          : 15.206.0.0/17        Peer             : 10.0.28.1
NextHop          : 10.0.28.1            Reuse time       : 00h29m22s
Peer AS          : 60203                Peer Router-Id   : 32.32.27.203
Local Pref       : none
Age              : 00h01m28s            Last update      : 02d01h20m
FOM Present      : 2936                 FOM Last upd.    : 3001
Number of Flaps  : 3                    Flags            : si
Path             : 60203 65001 19855 3356  702   1889
Applied Policy   : default-damping-profile
-------------------------------------------------------------------------------
*A:ALA-12#
```

group

Syntax **group** [*name*] [**detail**]

Context show>router>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays group information for a BGP peer group. This command can be entered with or without parameters.

When this command is entered without a group name, information about all peer groups displays.

When the command is issued with a specific group name, information only pertaining to that specific peer group displays.

The 'State' field displays the BGP group's operational state. Other valid states are:

Up - BGP global process is configured and running.

Down - BGP global process is administratively shutdown and not running.

Disabled - BGP global process is operationally disabled. The process must be restarted by the operator.

Parameters *name* — Displays information for the BGP group specified.

detail — Displays detailed information.

Output **Standard and Detailed Group Output —** The following table describes the standard and detailed command output fields for a BGP group:

Sample Output

Label	Description
Group	BGP group name
Group Type	No Type – Peer type not configured.
	External – Peer type configured as external BGP peers.
	Internal – Peer type configured as internal BGP peers.
State	Disabled – The BGP peer group has been operationally disabled.
	Down – The BGP peer group is operationally inactive.
	Up – The BGP peer group is operationally active.
Peer AS	The configured or inherited peer AS for the specified peer group.
Local AS	The configured or inherited local AS for the specified peer group.
Local Address	The configured or inherited local address for originating peering for the specified peer group.
Loop Detect	The configured or inherited loop detect setting for the specified peer group.

Label	Description (Continued)
Connect Retry	The configured or inherited connect retry timer value.
	Authentication
	None — No authentication is configured.
	MD5 — MD5 authentication is configured.
Local Pref	The configured or inherited local preference value.
MED Out	The configured or inherited MED value assigned to advertised routes without a MED attribute.
Min Route Advt.	The minimum amount of time that must pass between route updates for the same IP prefix.
Min AS Originate	The minimum amount of time that must pass between updates for a route originated by the local router.
Multihop	The maximum number of router hops a BGP connection can traverse.
Multipath	The configured or inherited multipath value, determining the maximum number of ECMP routes BGP can advertise to the RTM.
Prefix Limit	No Limit — No route limit assigned to the BGP peer group.
	1 - 4294967295 — The maximum number of routes BGP can learn from a peer.
Passive	Disabled — BGP attempts to establish BGP connections with neighbors in the specified peer group.
	Enabled — BGP will not actively attempt to establish BGP connections with neighbors in the specified peer group.
Next Hop Self	Disabled — BGP is not configured to send only its own IP address as the BGP nexthop in route updates to neighbors in the peer group.
	Enabled — BGP sends only its own IP address as the BGP nexthop in route updates to neighbors in the specified peer group.
Aggregator ID 0	Disabled — BGP is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group.
	Enabled — BGP is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group.
Remove Private	Disabled — BGP will not remove all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group.
	Enabled — BGP removes all private AS numbers from the AS path attribute in updates sent to the neighbor in the peer group.

Label	Description (Continued)
Damping	Disabled — The peer group is configured not to dampen route flaps. Enabled — The peer group is configured to dampen route flaps.
Export Policy	The configured export policies for the peer group.
Import Policy	The configured import policies for the peer group.
Hold Time	The configured hold time setting.
Keep Alive	The configured keepalive setting.
Cluster Id	None — No cluster ID has been configured.
Client Reflect	Disabled — The BGP route reflector will not reflect routes to this neighbor. Enabled — The BGP route reflector is configured to reflect routes to this neighbor.
NLRI	The type of NLRI information that the specified peer group can accept. Unicast — IPv4 unicast routing information can be carried.
Preference	The configured route preference value for the peer group.
List of Peers	A list of BGP peers configured under the peer group.
Total Peers	The total number of peers configured under the peer group.
Established	The total number of peers that are in an established state.

```
*A:ALA-12# show router 3 bgp group
===============================================================================
BGP Groups
===============================================================================
Group            : To_AS_40000
-------------------------------------------------------------------------------
Description      : Not Available
Group Type       : No Type              State            : Up
Peer AS          : 40000                Local AS         : 65206
Local Address    : n/a                  Loop Detect      : Ignore
Export Policy    : direct2bgp
Hold Time        : 90                   Keep Alive       : 30
Cluster Id       : None                 Client Reflect   : Enabled
NLRI             : Unicast              Preference       : 170

List of Peers
- 10.0.0.1       : To_Jukebox
- 10.0.0.12      : Not Available
- 10.0.0.13      : Not Available
- 10.0.0.14      : To_ALA-1
- 10.0.0.15      : To_H-215
Total Peers      : 5                    Established      : 2
===============================================================================
*A:ALA-12#
```

neighbor

Syntax **neighbor** [*ip-address* [[**family** *family*] *filter1*]]
neighbor [*as-number* [[**family** *family*] filter2]]

Context show>router>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays BGP neighbor information. This command can be entered with or without any parameters.

When this command is issued without any parameters, information about all BGP peers displays.

When the command is issued with a specific IP address or ASN, information regarding only that specific peer or peers with the same AS display.

When either **received-routes** or **advertised-routes** is specified, then the routes received from or sent to the specified peer is listed (see second output example).
Note: This information is not available by SNMP.

When either **history** or **suppressed** is specified, then the routes learned from those peers that either have a history or are suppressed (respectively) are listed.

The 'State' field displays the BGP peer's protocol state. In additional to the standard protocol states, this field can also display the 'Disabled' operational state which indicates the peer is operationally disabled and must be restarted by the operator.

Parameters *ip-addr* — Displays the BGP neighbor with the specified IP address.

family *family* **—** Specifies the type of routing information to be distributed by the BGP instance.

Values ipv4, vpn-ipv4

filter1 — Specifies route criteria.

Values received-routes, advertised-routes, history, suppressed, detail

filter2 — Specifies route criteria.

Values history, suppressed, detail

Output **Standard and Detailed Neighbor —** The following table describes the standard and detailed command output fields for a BGP neighbor:

Label	Description
Peer	The IP address of the configured BGP peer.
Group	The BGP peer group to which this peer is assigned.
Peer AS	The configured or inherited peer AS for the peer group.

Label	Description (Continued)
Peer Address	The configured address for the BGP peer.
Peer Port	The TCP port number used on the far-end system.
Local AS	The configured or inherited local AS for the peer group.
Local Address	The configured or inherited local address for originating peering for the peer group.
Local Port	The TCP port number used on the local system.
Peer Type	External — Peer type configured as external BGP peers.
	Internal — Peer type configured as internal BGP peers.
State	Idle — The BGP peer is not accepting connections.
	Active — BGP is listening for and accepting TCP connections from this peer.
	Connect — BGP is attempting to establish a TCP connection from this peer.
	Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer.
	Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.
	Established — BGP has successfully established a peering and is exchanging routing information.
Last State	Idle — The BGP peer is not accepting connections.
	Active — BGP is listening for and accepting TCP connections from this peer.
	Connect — BGP is attempting to establish a TCP connection with this peer.
	Connect — BGP is attempting to establish a TCP connections from this peer.
	Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer.
	Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.
	Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.
Last Event	start — BGP has initialized the BGP neighbor.
	stop — BGP has disabled the BGP neighbor.

Label	Description (Continued)
	open — BGP transport connection opened. close — BGP transport connection closed. openFail — BGP transport connection failed to open. error — BGP transport connection error. connectRetry — Connect retry timer expired. holdTime — Hold time timer expired. keepAlive — Keepalive timer expired. recvOpen — Receive an OPEN message. revKeepalive — Receive an KEEPALIVE message. recvUpdate — Receive an UPDATE message. recvNotify — Receive an NOTIFICATION message. None — No events have occurred.
Last Error	Displays the last BGP error and sub-code to occur on the BGP neighbor.
Connect Retry	The configured or inherited connect retry timer value.
Local Pref.	The configured or inherited local preference value.
Min Route Advt.	The minimum amount of time that must pass between route updates for the same IP prefix.
Min AS Originate	The minimum amount of time that must pass between updates for a route originated by the local router.
Multihop	The maximum number of router hops a BGP connection can traverse.
Multipath	The configured or inherited multipath value, determining the maximum number of ECMP routes BGP can advertise to the RTM.
Damping	Disabled — BGP neighbor is configured not to dampen route flaps. Enabled — BGP neighbor is configured to dampen route flaps.
Loop Detect	Ignore — The BGP neighbor is configured to ignore routes with an AS loop. Drop — The BGP neighbor is configured to drop the BGP peering if an AS loop is detected. Off — AS loop detection is disabled for the neighbor.

Label	Description (Continued)
MED Out	The configured or inherited MED value assigned to advertised routes without a MED attribute.
Authentication	None — No authentication is configured. MD5 — MD5 authentication is configured.
Next Hop Self	Disabled — BGP is not configured to send only its own IP address as the BGP nexthop in route updates to the specified neighbor. Enabled — BGP will send only its own IP address as the BGP nexthop in route updates to the neighbor.
AggregatorID Zero	Disabled — The BGP Neighbor is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates. Enabled — The BGP Neighbor is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates.
Remove Private	Disabled — BGP will not remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor. Enabled — BGP will remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor.
Passive	Disabled — BGP will actively attempt to establish a BGP connection with the specified neighbor. Enabled — BGP will not actively attempt to establish a BGP connection with the specified neighbor.
Prefix Limit	No Limit — No route limit assigned to the BGP peer group. 1 - 4294967295 — The maximum number of routes BGP can learn from a peer.
Hold Time	The configured hold time setting.
Keep Alive	The configured keepalive setting.
Active Hold Time	The negotiated hold time, if the BGP neighbor is in an established state.
Active Keep Alive	The negotiated keepalive time, if the BGP neighbor is in an established state.
Cluster Id	The configured route reflector cluster ID. None — No cluster ID has been configured
Client Reflect	Disabled — The BGP route reflector is configured not to reflect routes to this neighbor. Enabled — The BGP route reflector is configured to reflect routes to this neighbor.

Label	Description (Continued)
Preference	The configured route preference value for the peer group.
Num of Flaps	The number of flaps in the neighbor connection.
Recd. Prefixes	The number of routes received from the BGP neighbor.
Active Prefixes	The number of routes received from the BGP neighbor and active in the forwarding table.
Recd. Paths	The number of unique sets of path attributes received from the BGP neighbor.
Suppressed Paths	The number of unique sets of path attributes received from the BGP neighbor and suppressed due to route damping.
Input Queue	The number of BGP messages to be processed.
Output Queue	The number of BGP messages to be transmitted.
i/p Messages	Total number of packets received from the BGP neighbor.

Sample Output

```
*A:ALA-12# show router 3 bgp neighbor
===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer : 10.0.0.15         Group : To_AS_40000
-------------------------------------------------------------------------------
Peer AS          : 65205
Peer Address     : 10.0.0.15            Peer Port        : 0
Local AS         : 65206
Local Address    : 10.0.0.16            Local Port       : 0
Peer Type        : External
State            : Active               Last State       : Connect
Last Event       : openFail
Last Error       : Hold Timer Expire
Hold Time        : 90                   Keep Alive       : 30
Active Hold Time : 0                    Active Keep Alive: 0
Cluster Id       : None
Preference       : 170                  Num of Flaps     : 0
Recd. Prefixes   : 0                    Active Prefixes  : 0
Recd. Paths      : 0                    Suppressed Paths : 0
Input Queue      : 0                    Output Queue     : 0
i/p Messages     : 0                    o/p Messages     : 0
i/p Octets       : 0                    o/p Octets       : 0
i/p Updates      : 0                    o/p Updates      : 0
Export Policy    : direct2bgp
===============================================================================
*A:ALA-12#


*A:ALA-12# show router 3 bgp neighbor detail
```

```
===============================================================================
BGP Neighbor (detail)
===============================================================================
-------------------------------------------------------------------------------
Peer : 10.0.0.15         Group : To_AS_40000
-------------------------------------------------------------------------------
Peer AS          : 65205
Peer Address     : 10.0.0.15            Peer Port        : 0
Local AS         : 65206
Local Address    : 10.0.0.16            Local Port       : 0
Peer Type        : External
State            : Active               Last State       : Connect
Last Event       : openFail
Last Error       : Hold Timer Expire
Connect Retry    : 20                   Local Pref.      : 100
Min Route Advt.  : 30                   Min AS Orig.     : 15
Multipath        : 1                    Multihop         : 5
Damping          : Disabled             Loop Detect      : Ignore
MED Out          : No MED Out           Authentication   : None
Next Hop Self    : Disabled             AggregatorID Zero: Disabled
Remove Private   : Disabled             Passive          : Disabled
Prefix Limit     : No Limit
Hold Time        : 90                   Keep Alive       : 30
Active Hold Time : 0                    Active Keep Alive: 0
Cluster Id       : None                 Client Reflect   : Enabled
Preference       : 170                  Num of Flaps     : 0
Recd. Prefixes   : 0                    Active Prefixes  : 0
Recd. Paths      : 0                    Suppressed Paths : 0
Input Queue      : 0                    Output Queue     : 0
i/p Messages     : 0                    o/p Messages     : 0
i/p Octets       : 0                    o/p Octets       : 0
i/p Updates      : 0                    o/p Updates      : 0
Export Policy    : direct2bgp

===============================================================================
*A:ALA-12#
```

Output **Show Advertised and Received Routes Output —** The following table describes the command output fields for both the standard and detailed information for a neighbor:

Sample Output

```
*A:ALA-12# show router 3 bgp neighbor 10.0.0.16 received-routes
===============================================================================
 BGP Router ID : 10.0.0.16         AS : 65206   Local AS : 65206
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP Neighbor
===============================================================================
Flag  Network            Nexthop         LocalPref  MED        As-Path
-------------------------------------------------------------------------------
?     10.0.0.16/32       10.0.0.16       100        none       No As-Path
?     10.0.6.0/24        10.0.0.16       100        none       No As-Path
```

```
?     10.0.8.0/24         10.0.0.16       100         none       No As-Path
?     10.0.12.0/24        10.0.0.16       100         none       No As-Path
?     10.0.13.0/24        10.0.0.16       100         none       No As-Path
?     10.0.204.0/24       10.0.0.16       100         none       No As-Path
===============================================================================
*A:ALA-12#
```

paths

Syntax **paths**

Context show>router>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays a summary of BGP path attributes.

Output **Show Path Output —** The following table describes the command output fields for a BGP path.

Sample Output

```
*A:ALA-12# show router 3 bgp paths
===============================================================================
 BGP Router ID : 10.0.0.14         AS : 65206   Local AS : 65206
===============================================================================
BGP Paths
-------------------------------------------------------------------------------
Path: 60203 65001 19855 3356  15412
-------------------------------------------------------------------------------
Origin           : IGP                  Next Hop         : 10.0.28.1
MED              : 60203                Local Preference : none
Refs             : 4                    ASes             : 5
Segments         : 1
Flags            : EBGP-learned
Aggregator       : 15412  62.216.140.1
-------------------------------------------------------------------------------
Path: 60203 65001 19855 3356  1     1236  1236  1236  1236
-------------------------------------------------------------------------------
Origin           : IGP                  Next Hop         : 10.0.28.1
MED              : 60203                Local Preference : none
Refs             : 2                    ASes             : 9
Segments         : 1
Flags            : EBGP-learned
-------------------------------------------------------------------------------
*A:ALA-12#
```

routes

Syntax **routes** [**family** *family*] [*prefix* [**detail** | **longer**]]
routes [**family** *family*] [*prefix* [**hunt | brief**]]
routes [**family** *family*] [**community** *comm-id*]
routes [**family** *family*] [**aspath-regex** *reg-ex1*]
routes [**family** *family*] [*ipv6-prefix*[/*prefix-length*] [**detail** | **longer**] | [**hunt** [**brief**]]]

Context show>router>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays BGP route information.

When this command is issued without any parameters, then the entire BGP routing table displays.

When this command is issued with an IP prefix/mask or IP address, then the best match for the parameter displays.

Parameters **family** *family* **—** Specifies the type of routing information to be distributed by the BGP instance.

Values **ipv4** — Displays only those BGP peers that have the IPv4 family enable and not those capable of exchanging IP-VPN routes.
vpn-ipv4 — Displays the BGP peers that are IP-VPN capable.
ipv6 — Displays the BGP peers that are IPv6 capable.
mcast-ipv4 — Displays the BGP peers that are mcast-ipv4 capable.

prefix — Specifies the type of routing information to display.

Values *rd*|[*rd*:]*ip-address*[/*mask*]

rd	{ip-address:number1 as-number1:number2 as-number2:number3}
number1	1 — 65535
as-number1	1 — 65535
number2	0 — 4294967295
as-number2	1 — 4294967295
number3	0 — 65535
ip-address	a.b.c.d
mask	0 — 32

filter — Specifies route criteria.

Values **hunt** Displays entries for the specified route in the RIB-In, RIB-Out, and RTM.

longer Displays the specified route and subsets of the route.
detail Display the longer, more detailed version of the output.

aspath-regex "*reg-exp***" —** Displays all routes with an AS path matching the specified regular expression *reg-exp*.

community *comm.-id* **—** Displays all routes with the specified BGP community.

Values [*as-number1:comm-val1* | *ext-comm* | *well-known-comm*]

ext-comm	type:{ip-address:comm-val1 | as-number1:comm-val2 | as-number2:comm-val1}
as-number1	0..65535
comm-val1	0..65535

type	keywords: target, origin
ip-address	a.b.c.d
comm-val2	0 — 4294967295
as-number2	0 — 4294967295

well-known-comm no-export, no-export-subconfed, no-advertise

Output **Show BGP Routes —** The following table describes the command output fields for BGP routes.

Sample Output

```
*A:ALA-12>config>router>bgp# show router 3 bgp routes family ipv4
===============================================================================
 BGP Router ID : 10.10.10.103      AS : 200     Local AS : 200
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP Routes
===============================================================================
Flag  Network                              Nexthop         LocalPref  MED
      VPN Label                            As-Path
-------------------------------------------------------------------------------
No Matching Entries Found
===============================================================================
*A:ALA-12>config>router>bgp#

A:SR-12# show router bgp routes 100.0.0.0/31 hunt
===============================================================================
 BGP Router ID : 10.20.1.1   AS : 100Local AS : 100
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP Routes
===============================================================================
RIB In Entries
-------------------------------------------------------------------------------
Network        : 100.0.0.0/31
Nexthop        : 10.20.1.2
Route Dist.    : 10.20.1.2:1         VPN Label     : 131070
From        : 10.20.1.2
Res. Nexthop   : 10.10.1.2
Local Pref.    : 100                 Interface Name: to-sr7
Aggregator AS  : none                Aggregator    : none
Atomic Aggr.   : Not Atomic          MED           : none
Community      : target:10.20.1.2:1
Cluster        : No Cluster Members
Originator Id  : None                Peer Router Id: 10.20.1.2
Flags       : Used  Valid  Best  IGP
AS-Path        : No As-Path
VPRN Imported  : 1 2 10 12
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Routes : 1
```

```
===============================================================================
A:SR-12#
```

summary

Syntax **summary [all]**

Context show>router>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays a summary of BGP neighbor information.

If confederations are not configured, that portion of the output will not display.

The “State” field displays the global BGP operational state. The valid values are:

Up — BGP global process is configured and running.
Down — BGP global process is administratively shutdown and not running.
Disabled — BGP global process is operationally disabled. The process must be restarted by the operator.

For example, if a BGP peer is operationally disabled, then the state in the summary table shows the state ‘Disabled’

Parameters **all** — Displays BGP peers in all instances.

Output **Show BGP Summary Output —** The following table describes the command output fields for a BGP summary:

Sample Output

```
*A:ALA-12# show router 3 bgp summary
===============================================================================
 BGP Router ID : 10.0.0.14         AS : 65206   Local AS : 65206
===============================================================================
 BGP Admin State         : Up            BGP Oper State       : Up
 Confederation AS        : 40000
 Member Confederations   : 65205 65206 65207 65208

 Number of Peer Groups   : 2             Number of Peers      : 7
 Total BGP Active Routes : 86689         Total BGP Routes     : 116999
 Total BGP Paths         : 35860         Total Path Memory    : 2749476
 Total Supressed Routes  : 0             Total History Routes : 0
 Total Decayed Routes    : 0
===============================================================================
BGP Summary
===============================================================================
Neighbor           AS PktRcvd PktSent InQ OutQ   Up/Down State|Recv/Actv/Sent
-------------------------------------------------------------------------------
10.0.0.1        65206       5   21849   0    0 00h01m29s 32/0/86683
10.0.0.12       65206       0       0   0    0 00h01m29s Active
10.0.0.13       65206       5   10545   0   50 00h01m29s 6/0/86683
10.0.0.15       65205       0       0   0    0 00h01m29s Active
10.0.0.16       65206       5    9636   0   50 00h01m29s 6/0/86683
10.0.27.1           2       0       0   0    0 00h01m29s Active
```

```
10.0.28.1       60203   22512      15   0    0 00h01m29s 116955/86689/9
===============================================================================
*A:ALA-12#
```

interface

Syntax **interface** [{[<ip-address|ip-int-name>][**detail**]}|**summary**]

Context show>router

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays the router IP interface table sorted by interface index.

Parameters *ip-address* — Only displays the interface information associated with the specified IP address.

ip-int-name — Only displays the interface information associated with the specified IP interface name.

detail — Displays detailed IP interface information.

summary — Displays summary IP interface information for the router.

— Standard IP Interface Output — The following table describes the standard output fields for an IP interface:

Sample Output

Detailed IP Interface Output — The following table describes the detailed output fields for an IP interface.

```
*A:ALA-12# show router 3 interface detail
===============================================================================
Interface Table
===============================================================================
Interface
-------------------------------------------------------------------------------
If Name      : to-ser1
Admin State  : Up                                Oper State       : Up

IP Addr/mask : 10.10.13.3/24                     Address Type     : Primary
IGP Inhibit  : Disabled                          Broadcast Address: Host-ones

IP Addr/mask : 10.200.0.1/16                     Address Type     : Secondary
IGP Inhibit  : Enabled                           Broadcast Address: Host-ones
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
If Index     : 2
Port Id     : 1/1/2                              If Type          : Network
Egress Filter: none                              Ingress Filter   : 100
QoS Policy   : 1                                 SNTP Broadcast   : False
MAC Address  : 04:5d:01:01:00:02                 Arp Timeout      : 14400

ICMP Details
Redirects    : Disabled
Unreachables : Number - 100                     Time (seconds)   - 10
TTL Expired  : Number - 100                     Time (seconds)   - 10
===============================================================================
*A:ALA-12#
```

Summary IP Interface Output — The following table describes the summary output fields for the router IP interfaces.

Label	Description
Instance	The router instance number.
Router Name	The name of the router instance.
Interfaces	The number of IP interfaces in the router instance.

Sample Output

```
*A:ALA-12# show router 3 interface summary
===============================================================================
Router Summary (Interfaces)
===============================================================================
Instance  Router Name                       Interfaces    Admin-Up   Oper-Up
-------------------------------------------------------------------------------
1         Base                              7             7          5
===============================================================================
*A:ALA-12#
```

route-table

Syntax **route-table** [*ip-prefix* [*/mask*] [**longer**] | [**protocol** *protocol*] | [**summary**]]

Context show>router

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays the active routes in the routing table.

If no command line arguments are specified, all routes are displayed, sorted by prefix.

Parameters *ip-prefix*[/*mask*] — Displays routes only matching the specified *ip-prefix* and optional *mask*.

longer — Displays routes matching the *ip-prefix*/*mask* and routes with longer masks.

protocol *protocol* — Displays routes learned from the specified protocol.

Values bgp, isis, local, ospf, rip, static, aggregate

summary — Displays a route table summary information.

Output **Standard Show Route Table Output —** The following table describes the standard output fields for the route table.

Sample Output

```
*A:ALA-12# show router 3 route-table
===============================================================================
```

```
Route Table
===============================================================================
Dest Address       Next Hop        Type    Protocol    Age       Metric  Pref
-------------------------------------------------------------------------------
10.10.0.1/32       10.10.13.1      Remote  OSPF        65844     1001    10
10.10.0.2/32       10.10.13.1      Remote  OSPF        65844     2001    10
10.10.0.3/32       0.0.0.0         Local   Local       1329261   0       0
10.10.0.4/32       10.10.34.4      Remote  OSPF        3523      1001    10
10.10.0.5/32       10.10.35.5      Remote  OSPF        1084022   1001    10
10.10.12.0/24      10.10.13.1      Remote  OSPF        65844     2000    10
10.10.13.0/24      0.0.0.0         Local   Local       65859     0       0
10.10.15.0/24      10.10.13.1      Remote  OSPF        58836     2000    10
10.10.24.0/24      10.10.34.4      Remote  OSPF        3523      2000    10
10.10.25.0/24      10.10.35.5      Remote  OSPF        399059    2000    10
10.10.34.0/24      0.0.0.0         Local   Local       3543      0       0
10.10.35.0/24      0.0.0.0         Local   Local       1329259   0       0
10.10.45.0/24      10.10.34.4      Remote  OSPF        3523      2000    10
10.200.0.0/16      0.0.0.0         Local   Local       4513      0       0
192.168.0.0/20     0.0.0.0         Local   Local       1329264   0       0
192.168.254.0/24   0.0.0.0         Remote  Static      11        1       5
-------------------------------------------------------------------------------
*A:ALA-12#
```

```
*A:ALA-12# show router 3 route-table 10.10.0.4
===============================================================================
Route Table
===============================================================================
Dest Address       Next Hop        Type    Protocol    Age       Metric  Pref
-------------------------------------------------------------------------------
10.10.0.4/32       10.10.34.4      Remote  OSPF        3523      1001    10
-------------------------------------------------------------------------------
*A:ALA-12#
```

```
*A:ALA-12# show router 3 route-table 10.10.0.4/32 longer
===============================================================================
Route Table
===============================================================================
Dest Address       Next Hop        Type    Protocol    Age       Metric  Pref
-------------------------------------------------------------------------------
10.10.0.4/32       10.10.34.4      Remote  OSPF        3523      1001    10
-------------------------------------------------------------------------------
No. of Routes: 1
===============================================================================
+ : indicates that the route matches on a longer prefix
*A:ALA-12#
```

```
*A:ALA-12# show router 3 route-table protocol ospf
===============================================================================
Route Table
===============================================================================
Dest Address       Next Hop        Type    Protocol    Age       Metric  Pref
-------------------------------------------------------------------------------
10.10.0.1/32       10.10.13.1      Remote  OSPF        65844     1001    10
10.10.0.2/32       10.10.13.1      Remote  OSPF        65844     2001    10
10.10.0.4/32       10.10.34.4      Remote  OSPF        3523      1001    10
10.10.0.5/32       10.10.35.5      Remote  OSPF        1084022   1001    10
10.10.12.0/24      10.10.13.1      Remote  OSPF        65844     2000    10
10.10.15.0/24      10.10.13.1      Remote  OSPF        58836     2000    10
```

```
10.10.24.0/24      10.10.34.4      Remote  OSPF        3523      2000    10
10.10.25.0/24      10.10.35.5      Remote  OSPF        399059    2000    10
10.10.45.0/24      10.10.34.4      Remote  OSPF        3523      2000    10
-------------------------------------------------------------------------------
*A:ALA-12#


*A:ALA-12# show router 3 route-table summary
===============================================================================
Route Table Summary
===============================================================================
                            Active                  Available
-------------------------------------------------------------------------------
Static                      1                       1
Direct                      6                       6
BGP                         0                       0
OSPF                        9                       9
ISIS                        0                       0
RIP                         0                       0
Aggregate                   0                       0
-------------------------------------------------------------------------------
Total                       15                      15
===============================================================================
*A:ALA-12#
```

static-arp

Syntax **static-arp** [*ip-address* | *ip-int-name* | **mac** *ieee-mac-addr*]

Context show>router

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays the router static ARP table sorted by IP address.

If no options are present, all ARP entries are displayed.

Parameters *ip-address* — Only displays static ARP entries associated with the specified IP address.

ip-int-name — Only displays static ARP entries associated with the specified IP interface name.

mac *ieee-mac-addr* **—** Only displays static ARP entries associated with the specified MAC address.

Output **Static ARP Table Output —** The following table describes the output fields for the ARP table.

Sample Output

```
*A:ALA-12# show router 3 static-arp
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Age      Type Interface
-------------------------------------------------------------------------------
10.200.0.253    00:00:5a:40:00:01 00:00:00 Sta  to-ser1
12.200.1.1      00:00:5a:01:00:33 00:00:00 Inv  to-ser1a
-------------------------------------------------------------------------------
No. of ARP Entries: 2
===============================================================================
*A:ALA-12#


*A:ALA-12# show router 3 static-arp 12.200.1.1
===============================================================================
ARP Table
===============================================================================
IP Address    MAC Address        Age     Type Interface
-------------------------------------------------------------------------------
12.200.1.1    00:00:5a:01:00:33  00:00:00 Inv  to-ser1 a
===============================================================================
*A:ALA-12#

*A:ALA-12# show router 3 static-arp to-ser1
===============================================================================
ARP Table
===============================================================================
IP Address    MAC Address       Age      Type Interface
-------------------------------------------------------------------------------
10.200.0.253  00:00:5a:40:00:01 00:00:00 Sta  to-ser1
===============================================================================
S*A:ALA-12#

*A:ALA-12# show router 3 static-arp mac 00:00:5a:40:00:01
===============================================================================
ARP Table
===============================================================================
```

```
IP Address    MAC Address       Age      Type Interface
-------------------------------------------------------------------------------
10.200.0.253  00:00:5a:40:00:01 00:00:00 Sta  to-ser1
===============================================================================
*A:ALA-12#
```

static-route

Syntax **static-route** [*ip-prefix /mask*] | [**preference** *preference*] | [**next-hop** *ip-addr*] [**detail**]

Context show>router

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays the static entries in the routing table.

If no options are present. all static routes are displayed sorted by prefix.

Parameters *ip-prefix /mask* — Displays static routes only matching the specified *ip-prefix* and *mask*.

preference *preference* — Only displays static routes with the specified route preference.

Values 0 — 65535

next-hop *ip-addr* — Only displays static routes with the specified next hop IP address.

detail — Displays detailed information about the static route.

Output **Show Static Route Output —** The following table describes the output fields for the static route table:

Sample Output

```
*A:ALA-12# show router 3 static-route
===============================================================================
Route Table
===============================================================================
IP Addr/mask       Pref Metric Type Nexthop              Interface       Active
-------------------------------------------------------------------------------
192.168.250.0/24   5    1      ID   10.200.10.1          to-ser1         Y
192.168.252.0/24   5    1      NH   10.10.0.254          n/a             N
192.168.253.0/24   5    1      NH   to-ser1              n/a             N
192.168.253.0/24   5    1      NH   10.10.0.254          n/a             N
192.168.254.0/24   4    1      BH   black-hole           n/a             Y
===============================================================================
*A:ALA-12#
```

```
*A:ALA-12# show router 3 static-route 192.168.250.0/24
===============================================================================
Route Table
===============================================================================
IP Addr/mask       Pref Metric Type Nexthop              Interface       Active
-------------------------------------------------------------------------------
192.168.250.0/24   5    1      ID   10.200.10.1          to-ser1         Y
===============================================================================
```

```
*A:ALA-12#


*A:ALA-12# show router 3 static-route preference 4
===============================================================================
Route Table
===============================================================================
IP Addr/mask       Pref Metric Type Nexthop              Interface       Active
-------------------------------------------------------------------------------
192.168.254.0/24   4    1      BH   black-hole           n/a             Y
===============================================================================
*A:ALA-12#


*A:ALA-12# show router 3 static-route next-hop 10.10.0.254
===============================================================================
Route Table
===============================================================================
IP Addr/mask       Pref Metric Type Nexthop              Interface       Active
-------------------------------------------------------------------------------
192.168.253.0/24   5    1      NH   10.10.0.254          n/a             N
===============================================================================
*A:ALA-12#

*A:Dut-B# show router static-route

===============================================================================
Static Route Table (Router: Base)  Family: IPv4
===============================================================================
Prefix                                    Tag         Met    Pref Type Act
   Next Hop                                Interface
-------------------------------------------------------------------------------
1.2.3.4/32                                0           1      5    NH   Y
   10.11.25.6
ip-10.11.25.5_base_to_cpe_static
10.11.15.0/24                             0           1      5    NH   Y
   10.11.25.6
ip-10.11.25.5_base_to_cpe_static
-------------------------------------------------------------------------------
No. of Static Routes: 2
===============================================================================


*A:Dut-B# show router static-route detail
===============================================================================
Static Route Table (Router: Base)  Family: IPv4
===============================================================================
Network          : 1.2.3.4/32
Nexthop          : 10.11.25.6
Type             : Nexthop                         Nexthop Type      : IP
Interface        : ip-10.11.25.5_base_to_cpe_stat* Active            : Y
Metric           : 1                               Preference        : 5
Admin  State     : Up                              Tag               : 0
BFD              : disabled
CPE-check        : enabled                         State             : n/a
Target           : 10.11.18.6
Interval         : 1                               Drop Count        : 3
Log              : N
CPE Host Up Time : 0d 00:00:02
CPE Echo Req Tx  : 3                               CPE Echo Reply Rx : 3
CPE Up Trans     : 1                               CPE Down Trans    : 0
```

```
CPE TTL          : 2
-------------------------------------------------------------------------------
Network          : 10.11.15.0/24
Nexthop          : 10.11.25.6
Type             : Nexthop                          Nexthop Type      : IP
Interface        : ip-10.11.25.5_base_to_cpe_stat* Active            : Y
Metric           : 1                                Preference        : 5
Admin  State     : Up                               Tag               : 0
BFD              : disabled
CPE-check        : disabled
-------------------------------------------------------------------------------
No. of Static Routes: 2

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

tunnel-table

Syntax **tunnel-table** [*ip-address*[*/mask*] [**protocol** *protocol* | **sdp** *sdp-id*]
tunnel-table [**summary**]

Context show>router

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays tunnel table information.

When the **auto-bind** command is used when configuring a VPRN service, it means the MP-BGP NH resolution is referring to core routing instance for IP reachability. For a VPRN service this object specifies the lookup to be used by the routing instance if no SDP to the destination exists.

Parameters *ip-address[/mask*] — Displays the specified tunnel table's destination IP address and mask.

protocol *protocol* **—** Displays LDP protocol information.

sdp *sdp-id* **—** Displays information pertaining to the specified SDP.

summary — Displays summary tunnel table information.

Output **Show Tunnel Table Output —** The following table describes tunnel table output fields:

Sample Output

```
*A:ALA-12>config>service#  show router 3 tunnel-table summary
===============================================================================
Tunnel Table Summary (Router: Base)
===============================================================================
                              Active                  Available
-------------------------------------------------------------------------------
LDP                           1                       1
SDP                           1                       1
===============================================================================
*A:ALA-12>config>service#
```

VPRN Clear Commands

Platforms Supported: 7210 SAS-K 2F4T6C.

Platforms Supported: 7210 SAS-K 2F4T6C.

interface

Syntax **interface** [*ip-int-name* | *ip-addr*] [**icmp**]

Context clear>router

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command clears IP interface statistics.

If no IP interface is specified either by IP interface name or IP address, the command will perform the clear operation on all IP interfaces.

Parameters *ip-int-name* | *ip-addr* — The IP interface name or IP interface address.

Default All IP interfaces.

icmp — Specifies to reset the ICMP statistics for the IP interface(s) used for ICMP rate limit.

damping

Syntax **damping** [[*ip-prefix*/*mask*] [**neighbor** *ip-address*]] | [**group** *name*]

Context clear>router>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command clears or resets the route damping information for received routes.

Parameters *ip-prefix/mask* — Clears damping information for entries that match the IP prefix and mask length.

neighbor *ip-address* **—** Clears damping information for entries received from the BGP neighbor.

group *name* **—** Clears damping information for entries received from any BGP neighbors in the peer group.

flap-statistics

Syntax **flap-statistics** [[*ip-prefixlmask*] [**neighbor** *ip-addr*]] | [**group** *group-name*] | [**regex** *reg-exp*] | [**policy** *policy-name*]

Context clear>router>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command clears route flap statistics.

Parameters *ip-prefix/mask* — Clears route flap statistics for entries that match the specified IP prefix and mask length.

neighbor *ip-addr* **—** Clears route flap statistics for entries received from the specified BGP neighbor.

group *group-name* **—** Clears route flap statistics for entries received from any BGP neighbors in the specified peer group.

regex *reg-exp* **—** Clears route flap statistics for all entries which have the regular expression and the AS path that matches the regular expression.

policy *policy-name* **—** Clears route flap statistics for entries that match the specified route policy.

neighbor

Syntax **neighbor** {*ip-addr* | **as** *as-number* | **external** | **all**} [**soft** | **soft-inbound** | **statistics**]

Context clear>router>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command resets the specified BGP peer or peers. This can cause existing BGP connections to be shutdown and restarted.

Parameters *ip-addr* — Resets the BGP neighbor with the specified IP address.

as *as-number* **—** Resets all BGP neighbors with the specified peer AS.

external — Resets all EBGP neighbors.

all — Resets all BGP neighbors.

soft — The specified BGP neighbor(s) re-evaluates all routes in the Local-RIB against the configured export policies.

soft-inbound — The specified BGP neighbor(s) re-evaluates all routes in the RIB-In against the configured import policies.

statistics — The BGP neighbor statistics.

protocol

Syntax **protocol**

Context clear>router>bgp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command resets the entire BGP protocol. If the AS number was previously changed, the BGP AS number does not inherit the new value.

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command clears commands for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service.

Values 1 — 2147483648

sap

Syntax **sap** *sap-id* {**all** | **counters** | **stp**}

Context clear>service>statistics

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Clears SAP statistics for a SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 for command syntax.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*:*vc-id* **ingress-vc-label**

Context clear>service>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

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

sdp

Syntax **sdp** *sdp-id* **keep-alive**

Context clear>service>statistics

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command clears keepalive statistics associated with the SDP ID.

Parameters *sdp-id* — The SDP ID for which to clear keepalive statistics.

Values 1 — 17407

counters

Syntax **counters**

Context clear>service>statistics>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Clears all traffic queue counters associated with the service ID.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] {**all** | **counters** | **stp**}

Context clear>service>statistics>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command clears statistics for the spoke SDP bound to the service.

Parameters *sdp-id* — The spoke SDP ID for which to clear statistics.

Values 1 — 17407

vc-id — The virtual circuit ID on the SDP ID to be reset.

Values 1 — 4294967295

all — Clears all queue statistics and STP statistics associated with the SDP.

counters — Clears all queue statistics associated with the SDP.

stp — Clears all STP statistics associated with the SDP.

stp

Syntax **stp**

Context clear>service>statistics>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

Clears all spanning tree statistics for the service ID.

VPRN Debug Commands

id

Syntax [**no**] **id** *service-id*

Context debug>service

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command debugs commands for a specific service.

The **no** form of the command disables debugging.

Parameters *service-id* — The ID that uniquely identifies a service.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 for command syntax.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id
debug>service>stp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables STP debugging for a specific SAP.

The **no** form of the command disables debugging.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 for command syntax.

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables STP debugging for a specific SDP.

The **no** form of the command disables debugging.

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 **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables debugging for a particular event type.

The **no** form of the command disables debugging.

event-type

Syntax [**no**] **event-type** {**config-change | oper-status-change**}

Context debug>service>id>sap

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables debugging for a particular event type.

The **no** form of the command disables debugging.

stp

Syntax [**no**] **stp**

Context debug>service>id

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables the context for debugging STP.

The **no** form of the command disables debugging.

all-events

Syntax **all-events**

Context debug>service>id>stp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables STP debugging for all events.

The **no** form of the command disables debugging.

bpdu

Syntax [**no**] **bpdu**

Context debug>service>stp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables STP debugging for received and transmitted BPDUs.

The **no** form of the command disables debugging.

core-connectivity

Syntax [**no**] **core-connectivity**

Context debug>service>stp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables STP debugging for core connectivity.

The **no** form of the command disables debugging.

exception

Syntax [**no**] **exception**

Context debug>service>stp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables STP debugging for exceptions.

The **no** form of the command disables debugging.

fsm-state-changes

Syntax [**no**] **fsm-state-changes**

Context debug>service>stp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables STP debugging for FSM state changes.

The **no** form of the command disables debugging.

fsm-timers

Syntax [**no**] **fsm-timers**

Context debug>service>stp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables STP debugging for FSM timer changes.

The **no** form of the command disables debugging.

port-role

Syntax [**no**] **port-role**

Context debug>service>stp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables STP debugging for changes in port roles.

The **no** form of the command disables debugging.

port-state

Syntax [**no**] **port-state**

Context debug>service>stp

Description **Platforms Supported**: 7210 SAS-K 2F4T6C.

This command enables STP debugging for port states.

The **no** form of the command disables debugging.

Service Global Commands

In This Chapter

This section provides show command descriptions and output.

- Services Show Commands on page 826
  - → Service Commands on page 826
  - → VLL
    - – VLL Show Commands on page 835
    - – VLL Clear Commands on page 855
  - → VPLS
    - – VPLS Show Commands on page 865
    - – VPLS Clear Commands on page 913
    - – VPLS Debug Commands on page 916
  - → IES
    - – IES Show Commands on page 919
  - → VPRN
    - – VPRN Show Commands on page 743

Show, Clear, Debug, Commands

Services Show Commands

Service Commands

customer

Syntax **customer** [*customer-id*] [**site** *customer-site-name*]]

Context show>service

Description This command displays service customer information.

Parameters *customer-id* — Displays only information for the specified customer ID.

Default All customer IDs display.

Values 1 — 2147483647

site *customer-site-name* — Specifies the customer site which is an anchor point for an ingress and egress virtual scheduler hierarchy.

Output **Show Customer Command Output —** The following table describes show customer command output fields:

Label	Description
Customer-ID	The ID that uniquely identifies a customer.
Contact	The name of the primary contact person.
Description	Generic information about the customer.
Phone	The phone/pager number to reach the primary contact person.
Total Customers	The total number of customers configured.
Site	Multi-service site name. A multi-service customer site is a group of SAPs with common origination and termination points.
Description	Displays information about a specific customer's multi-service site.
Assignment	The port ID, MDA, or card number, where the SAPs members of this multi-service site are defined.
I. Sched Pol	The ingress QoS scheduler policy assigned to this multi-service site.
E. Sched Pol	The egress QoS scheduler policy assigned to this multi-service site.
Service-ID	The ID that uniquely identifies a service.

SAP — Specifies the SAP assigned to the service.

Sample Output

```
*A:ALA-12# show service customer
============================================================
Customers
============================================================
Customer-ID : 1
Contact     : Manager
Description : Default customer
Phone       : (123) 555-1212

Customer-ID : 2
Contact     : Tech Support
Description : TiMetra Networks
Phone       : (234) 555-1212

Customer-ID : 3
Contact     : Test
Description : TiMetra Networks
Phone       : (345) 555-1212

Customer-ID : 6
Contact     : Test1
Description : Epipe Customer
Phone       : (456) 555-1212

Customer-ID : 7
Contact     : Test2
Description : VPLS Customer
Phone       : (567) 555-1212

Customer-ID : 8
Contact     : Customer Service
Description : IES Customer
Phone       : (678) 555-1212

Customer-ID : 274
Contact     : TestA
Description : ABC Company
Phone       : 650 123-4567

Customer-ID : 94043
Contact     : Test Engineer on Duty
Description : TEST Customer
Phone       : (789) 555-1212
------------------------------------------------------
Total Customers : 8
------------------------------------------------------------
*A:ALA-12#
*A:ALA-12# show service customer 274
===============================================================================
Customer  274
===============================================================================
Customer-ID : 274
Contact     : Mssrs. Beaucoup
Description : ABC Company
```

```
Phone       : 650 123-4567
-------------------------------------------------------------------------------
Multi Service Site
-------------------------------------------------------------------------------
Site        : west
Description : (Not Specified)
===============================================================================
*A:ALA-12#
```

fdb-mac

Syntax **fdb-mac** [*ieee-address*] [**expiry**]

Context show>service

Description This command displays the FDB entry for a given MAC address.

Parameters *ieee-address* — 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.

expiry — shows amount of time until MAC is aged out.

Sample Output

```
*A:ALA-48# show service fdb-mac
===============================================================================
Service Forwarding Database
===============================================================================
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
103       12:34:56:78:90:0f sap:1/1/7:0           Static    02/02/2009 09:27:57
700       90:30:ff:ff:ff:8f cpm                   Host      02/02/2009 09:27:57
-------------------------------------------------------------------------------
No. of Entries: 2
===============================================================================
*A:ALA-48#


*A:ALA-48# show service fdb-mac expiry
===============================================================================
Service Forwarding Database
===============================================================================
ServId    MAC               Source-Identifier     Type/     Last Change
                                                  Expiry
-------------------------------------------------------------------------------
103       12:34:56:78:90:0f sap:1/1/7:0           Static    02/02/2009 09:27:57
700       90:30:ff:ff:ff:8f cpm                   Host      02/02/2009 09:27:57
-------------------------------------------------------------------------------
No. of Entries: 2
===============================================================================
*A:ALA-48#
```

sdp

sdp [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]

Context show>service

Description This command displays SDP information.

If no optional parameters are specified, a summary SDP output for all SDPs is displayed.

Parameters *sdp-id* — The SDP ID for which to display information.

Default All SDPs.

Values 1 — 17407

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.

Output **Show Service SDP —** The following table describes show service SDP output fields.

Sample Output

```
*A:ALA-7210M# show service sdp
===============================================================================
Services: Service Destination Points
===============================================================================
SdpId    Adm MTU   Opr MTU   IP address       Adm  Opr         Del LSP Signal
-------------------------------------------------------------------------------
10       4462      4462      10.20.1.3        Up   Dn NotReady MPLS B TLDP
40       4462      1534      10.20.1.20       Up   Up          MPLS B TLDP
60       4462      1514      10.20.1.21       Up   Up          MPLS B TLDP
100      4462      4462      180.0.0.2        Down Down        MPLS B TLDP
500      4462      4462      10.20.1.50       Up   Dn NotReady MPLS B TLDP
-------------------------------------------------------------------------------
Number of SDPs : 5
===============================================================================
*A:ALA-7210M#


*7210SAS>show>service# sdp 1 detail

===============================================================================
Service Destination Point (Sdp Id : 1) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 1  -0.0.0.0
-------------------------------------------------------------------------------
Description           : (Not Specified)
SDP Id                : 1                     SDP Source         : manual
Admin Path MTU        : 0                     Oper Path MTU      : 0
Far End               : 0.0.0.0               Delivery           : MPLS
Tunnel Far End        : n/a                   LSP Types          : None
```

```
Admin State          : Down                  Oper State         : Down
Signaling            : TLDP                  Metric             : 0
Acct. Pol            : None                  Collect Stats      : Disabled
Last Status Change   : 11/04/2099 22:56:41   Adv. MTU Over.     : No
Last Mgmt Change     : 11/10/2099 15:56:44   VLAN VC Etype      : 0x8100
Bw BookingFactor     : 100                   PBB Etype          : 0x88e7
Oper Max BW(Kbps)    : 0                     Avail BW(Kbps)     : 0
Net-Domain           : default               Egr Interfaces     : n/a
Flags                : SdpAdminDown NoSysIPAddr
                       TranspTunnDown

Mixed LSP Mode Information :
Mixed LSP Mode       : Enabled             Active LSP Type    : RSVP....also be LDP,
BGP
Revert Time          : 200                   Revert Count Down  : n/a

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

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

===============================================================================
*7210SAS>show>service#
```

sdp-using

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

Context show>service

Description This command displays services using SDP or far-end address options.

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

Values 1 — 17407

vc-id — The virtual circuit identifier.

Values 1 — 4294967295

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

Default Services with any far-end IP address.

Output **Show Service SDP Using X —** 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
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-7210# 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               Spok 10.0.0.13      Up        131071   131071
2           300:2               Spok 10.0.0.13      Up        131070   131070
100         300:100             Spok 10.0.0.13      Up        131069   131069
101         300:101             Spok 10.0.0.13      Up        131068   131068
-------------------------------------------------------------------------------
Number of SDPs : 4
===============================================================================
*A:ALA-7210#
```

service-using

Syntax **service-using** [**epipe**][**ies**] [**vpls**] [**mirror**] [**i-vpls**] [**m-vpls**] [**customer** *customer-id*]

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.

customer *customer-id* **—** Displays services only associated with the specified customer ID.

Default Services associated with a customer.

Values 1 — 2147483647

Output **Show Service Service-Using —** The following table describes show command output fields.

Label	Description
Service Id	The service identifier.
Type	Specifies the service type configured for the service ID.
Adm	The desired state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.
Service name	The name of the service.

Sample Output

```
*7210SAS>show>service# service-using customer 1

===============================================================================
Services  Customer 1
===============================================================================
ServiceId    Type      Adm  Opr  CustomerId Service Name
-------------------------------------------------------------------------------
1            VPLS      Up   Up   1
2            VPLS      Up   Up   1
3            VPLS      Up   Up   1
4            VPLS      Up   Up   1
2147483648   IES       Up   Down 1          _tmnx_InternalIesService
2147483649   intVpls   Up   Down 1          _tmnx_InternalVplsService
-------------------------------------------------------------------------------
Matching Services : 6
-------------------------------------------------------------------------------
===============================================================================
*7210SAS>show>service#
```

id

Syntax **id** <*service-id*>

Context show>service

Description This command displays vpls-template used to instantiate this service and m-vpls that controls this service.

eth-ring

Syntax ***eth-ring*** *[status]*
eth-ring *[ring-index] hierarchy*
eth-ring *ring-index [path {a|b}]*

Context show

Description This command displays the Ethernet rings information.

Parameters *status* — Displays the status information of the Ethernet rings configured on the system.

hierarchy — Displays eth-ring hierarical relationships.

path {a|b} — Displays information related to the configured Ethernet rings.

Output **Show Ethernet Ring Status —** The following table describes show command output fields.

Label	Description
Ring Id	The ring identifier
Admin State	Displays the administrative state
Oper State	Displays the operational state
Path Information	
Path	Displays the path information
Tag	Displays the tag information
State	Displays the state of the path
MEP Information	
Ctrl-MEP	Displays the Ctrl-MEP information
CC-Intvl	Displays the Ctrl-Interval information
Defects	Displays the defects

```
*A:NS1015C0821>show# eth-ring status
```

```
===============================================================================
Ethernet Ring (Status information)
===============================================================================
Ring   Admin  Oper       Path Information                MEP Information
ID     State  State  Path         Tag       State     Ctrl-MEP CC-Intvl Defects
-------------------------------------------------------------------------------
1      Up     Up     a - 1/1/1    100       Up        Yes      100ms     -----
                     b - 1/1/2    100       Up        Yes      100ms     -----
10     Down   Down   a - N/A                -         -        -         -----
                     b - N/A                -         -        -         -----
===============================================================================
Ethernet Tunnel MEP Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM
*A:NS1015C0821>show#
```

Output **Show Ethernet Ring —** The following table describes show command output fields.

Label	Description
Description	The ring description
Admin State	Displays the administrative state
Oper State	Displays the operational state
Node ID	Displays the node identifier
Guard Time	Displays the configured guard time
Max Revert time	Displays the configured maximum revert time
CCM Hold down time	Displays the configured CCM Hold down time
APS TX PDU	Displays the APS TX PDU information
Defect Status	Displays the defect status
RPL Node	Displays the RPL node information
Time to revert	Displays the configured time to revert
CCM Hold Up Time	Displays the configured CCM Hold up time
Sub-Ring Type	Displays the sub-ring type information, the sub-ring type can be virtual link or on-virtual link.
Interconnect-ID	Displays the interconnect ID. The ID can be a ring-index ID or VPLS service ID.
Compatible Ver-sion	Displays the Ethernet ring version information.

```
*A:NS1015C0821>show# eth-ring 10
```

```
===============================================================================
Ethernet Ring 10 Information
===============================================================================
Description        : (Not Specified)
Admin State        : Down               Oper State       : Down
Node ID            : 00:25:ba:03:48:04
Guard Time         :    5 deciseconds   RPL Node         : rplNone
Max Revert Time    :  300 seconds       Time to Revert   : N/A
CCM Hold Down Time :    0 centiseconds  CCM Hold Up Time :   20 deciseconds
APS Tx PDU         : N/A
Defect Status      :

-------------------------------------------------------------------------------
Ethernet Ring Path Summary
-------------------------------------------------------------------------------
Path Port     Raps-Tag     Admin/Oper      Type            Fwd State
-------------------------------------------------------------------------------
  a  -        -                -/-         -               -
  b  -        -                -/-         -               -
===============================================================================
*A:NS1015C0821>show#
```

pw-routing

Syntax **pw-routing {local-prefix|static-route|paths|all}**
pw-routing route-table [all-routes]
pw-routing route-table summary

Context show>service

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command displays PW routing information at this 7210 node.

Parameters **local-prefix|static-route|paths|all —** Shows details of the T-PE prefixes configured on this node, static routes from this node, explicit PW paths configured on this node, or all of these.

route-table [**all-routes**] **—** Displays the PW routing table on this node. If all-routes is specified, then the full routing table is displayed.

route-table summary — Displays a summary of the PW routing table for this node.

Sample Output

```
*A:Dut-C# show service pw-routing local-prefix
===============================================================================
Service PW Routing Information
===============================================================================
===============================================================================
Service PW Routing Local-Prefix RD Information
===============================================================================
Local-Prefix              Route-Dist                Community       Adv-Bgp
-------------------------------------------------------------------------------
3:10.20.1.3               100:3                     100:3           enabled
```

```
                           100:4                       100:4           enabled
-------------------------------------------------------------------------------
Local-Prefix Entries found: 1
===============================================================================
===============================================================================
*A:Dut-C# show service pw-routing static-route
===============================================================================
Service PW Routing Information
===============================================================================
===========================================================
Service PW Routing Static-Route Information
===========================================================
Prefix                        Next-Hop
-----------------------------------------------------------
6:10.20.1.6/64                10.20.1.5
-----------------------------------------------------------
Static Route Entries found: 1
===========================================================
===============================================================================


*A:Dut-C# show service pw-routing paths
===============================================================================
Service PW Routing Information
===============================================================================
================================================================
Service PW Routing Path Information
================================================================
Path                              Adm     Hop IP Address
----------------------------------------------------------------
path1_to_F                        up      1   10.20.1.5
                                          2   10.20.1.2
path1_to_F2                       up      1   10.20.1.2
                                          2   10.20.1.5
----------------------------------------------------------------
Path Entries found: 2
================================================================
===============================================================================


*A:Dut-C# show service pw-routing all
===============================================================================
Service PW Routing Information
===============================================================================
SPE-Address       : 3:10.20.1.3
Boot Timer        : 10 secs
Boot Timer Remain : 0 secs
Retry Timer       : 30 secs
Retry Count       : 30

===============================================================================
Service PW Routing Local-Prefix RD Information
===============================================================================
Local-Prefix              Route-Dist                  Community       Adv-Bgp
-------------------------------------------------------------------------------
3:10.20.1.3               100:3                       100:3           enabled
                          100:4                       100:4           enabled
-------------------------------------------------------------------------------
Local-Prefix Entries found: 1
===============================================================================
===========================================================
```

```
Service PW Routing Static-Route Information
==========================================================
Prefix                        Next-Hop
----------------------------------------------------------
6:10.20.1.6/64                10.20.1.5
----------------------------------------------------------
Static Route Entries found: 1
==========================================================
==============================================================
Service PW Routing Path Information
==============================================================
Path                              Adm     Hop IP Address
--------------------------------------------------------------
path1_to_F                        up      1   10.20.1.5
                                          2   10.20.1.2
path1_to_F2                       up      1   10.20.1.2
                                          2   10.20.1.5
--------------------------------------------------------------
Path Entries found: 2
==============================================================
==============================================================================


*A:Dut-C# show service pw-routing route-table all-routes
==============================================================================
Service PW L2 Routing Information
==============================================================================
AII-Type2/Prefix-Len                          Next-Hop       Owner  Age
 Route-Distinguisher                           Community      Best
------------------------------------------------------------------------------
3:10.20.1.3:0/64                              10.20.1.3      local  00h32m08s
 0:0                                           0:0            yes
3:10.20.1.3:1/96                              10.20.1.3      host   00h32m08s
 0:0                                           0:0            yes
3:10.20.1.3:2/96                              10.20.1.3      host   00h32m08s
 0:0                                           0:0            yes
3:10.20.1.3:3/96                              10.20.1.3      host   00h32m08s
 0:0                                           0:0            yes
3:10.20.1.3:4/96                              10.20.1.3      host   00h32m08s
 0:0                                           0:0            yes
3:10.20.1.3:5/96                              10.20.1.3      host   00h32m08s
 0:0                                           0:0            yes
3:10.20.1.3:6/96                              10.20.1.3      host   00h32m08s
 0:0                                           0:0            yes
3:10.20.1.3:7/96                              10.20.1.3      host   00h32m08s
 0:0                                           0:0            yes
3:10.20.1.3:8/96                              10.20.1.3      host   00h32m08s
 0:0                                           0:0            yes
3:10.20.1.3:9/96                              10.20.1.3      host   00h32m08s
 0:0                                           0:0            yes
3:10.20.1.3:10/96                             10.20.1.3      host   00h32m07s
 0:0                                           0:0            yes
6:10.20.1.6:0/64                              10.20.1.5      static 00h07m33s
 0:0                                           0:0            yes
6:10.20.1.6:0/64                              10.20.1.5      bgp    00h31m34s
 100:6                                         100:6          no
------------------------------------------------------------------------------
Entries found: 13
==============================================================================
```

```
*A:Dut-C# show service pw-routing route-table summary
=======================================
Service PW L2 Routing Summary
=======================================
Source              Active
---------------------------------------
BGP                 1
Static              1
Host                10
Local               3
---------------------------------------
Total               15
=======================================
```

pw-template

Syntax **pw-template**

Context show>service

Sample Output

```
*A:Dut-B#    show service  pw-template 1
===============================================================================
PW Template Information
===============================================================================
PW Tmpl Id           : 1
Use Provisioned Sdp  : enabled              VcType           : vlan
Acctg Policy         : default              Collect Stats    : disabled
Mac-Learning         : enabled              Mac-Ageing       : enabled
Discard Unkn Src     : disabled             Limit MacMove    : blockable
Mac-Pinning          : disabled             Vlan VcTag       : 4095
MAC Address Limit    : no limit             Rest Prot Src Mac: disabled
Auto Learn Mac Prot  : disabled             RestProtSrcMacAct: disable
Block On Peer Fault  : disabled

SHG
Name                 :
Description          : (Not Specified)
Rest Prot Src Mac    : disabled             Rest Unprot Dst  : disabled
Auto Learn Mac Prot  : disabled             RestProtSrcMacAct: disable

Egress
Mac FilterId         : none                 Ip FilterId      : none
Ipv6 FilterId        : none                 QoS NetPlcyId    : none
Port RedirectQGrp    : none                 Instance Id      : none

Ingress
Mac FilterId         : none                 Ip FilterId      : none
Ipv6 FilterId        : none                 QoS NetPlcyId    : none
Fp RedirectQGrp      : none                 Instance Id      : none

IGMP
Fast Leave           : disabled             Import Plcy      : none
Last Memb Intvl      : 10 deci-secs         Max Nbr Grps     : 0
Send Queries         : disabled
Version              : 3
```

```
Force VlanVc Fwd     : disabled             Control Word     : disabled
Hash Label           : disabled             Hash Lbl Sig Cap : disabled
Last Changed         : 02/12/2013 22:11:49
-------------------------------------------------------------------------
Included SDP-Groups
-------------------------------------------------------------------------
red
-------------------------------------------------------------------------
```

pw-template-using

Syntax **pw-template-using**

Context show>service

```
*A:dut-a>show>service# pw-template-using

===========================================
Services using PW Template
===========================================
PW Template Id        SvcId
-------------------------------------------
5                     5
===========================================
*A:dut-a>show>service#
```

sdp

Syntax **sdp** *sdp-id* **pw-port** [*pw-port-id*]
sdp [**consistent**|**inconsistent**|**na**] **egressifs**
sdp *sdp-id* **keep-alive-history**
sdp far-end *ip-address* **keep-alive**-history
sdp [*sdp-id*] [**detail**]
sdp far-end *ip-address* [**detail**]

Context show>service

Description **Platforms Supported:** 7210 SAS-K 2F4T6C.

This command displays SDP information.

If no optional parameters are specified, a summary SDP output for all SDPs is displayed.

Parameters *sdp-id* — The SDP ID for which to display information.

Default All SDPs.

Values 1 — 17407

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.

pw-port *pw-port-id* **—** Displays the SAP identifier for PW-SAPs.

Output **Show Service SDP —** The following table describes show service SDP output fields.

Label	Description
SDP Id	The SDP identifier.
Description	Displays a text string describing the SDP.
Admin Path MTU	Displays 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. The default value of zero indicates that the path MTU should be computed dynamically from the corresponding MTU of the tunnel.
Opr Path MTU	Displays the actual 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. In order to be able to bind this SDP to a given service, the value of this object minus the control word size (if applicable) must be equal to or larger than the MTU of the service, as defined by its service MTU.
Far End	Displays the far end IP address.
Delivery	The type of delivery used by the SDP: MPLS.
IP address	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Adm Admin State	The desired state of the SDP.
Opr Oper State	The operating state of the SDP.
Flags	Specifies all the conditions that affect the operating status of this SDP.
Signal Signaling	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
Last Status Change	The time of the most recent operating status change to this SDP.
Adv. NTU Over	Specifies whether the advertised MTU of a VLL spoke SDP bind includes the 14-byte L2 header, so that it is backward compatible with pre-2.0 software.
Last Mgmt Change	The time of the most recent management-initiated change to this SDP.
KeepAlive Information	This section displays Keepalive information.

Label	Description (Continued)
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	The length of the SDP echo request messages transmitted on this SDP.
Hello Timeout	The number of seconds to wait for an SDP echo response message before declaring a timeout.
Unmatched Replies	The number of SDP unmatched message replies timer expired.
Max Drop Count	The maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	The amount of time to wait before the keepalive operating status is eligible to enter the alive state.
TX Hello Msgs	The number of SDP echo request messages transmitted since the keepalive was administratively enabled or the counter was cleared.
Rx Hello Msgs	The number of SDP echo request messages received since the keepalive was administratively enabled or the counter was cleared.
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.
Lsp Name	Displays the LSP name.
Time Since Last Transaction	Displays the time of the last transaction.
Signaling	Specifies the signaling type.
Collect Stats	Specifies whether the agent collects accounting statistics for this SDP. When the value is true the agent collects accounting statistics on this SDP.
VLAN VC Etype	Displays the VLAN VC type.
BW Booking Factor	Specifies the value used to calculate the max SDP available bandwidth. The value specifies the percentage of the SDP max available bandwidth for VLL call admission. When the value of is set to zero (0), no new VLL spoke-sdp bindings with non-zero bandwidth are permitted with this SDP. Overbooking, >100% is allowed.
PBB Etype	Displays the Ethertype used in frames sent out on this SDP when specfied as **vlan** for Provider Backbone Bridging frames.
Oper Max BW (Kbps)	Indicates the operational bandwidth in kilo-bits per seconds (Kbps) available for this SDP. The value is determined by the sum of the bandwidth of all the RSVP LSPs used by the SDP.
Avail BW (Kbps)	Indicates the bandwidth that is still free for booking by the SDP bindings on the SDP.

Label	Description (Continued)
Net-Domain	Specifies the network-domain name configured on this SDP. The default value of this object is the default'network-domain.
Egr Interface	Indicates whether all the egress network interfaces that can carry traffic on this SDP are associated with the network-domain configured on this SDP. `not applicable.` indicates that there is no egress network interface that can carry traffic on this SDP. `consistent.` Indicates that the network-domains for all the egress network interfaces that can carry traffic on this SDP are consistent. `inconsistent.` indicates that the network-domain for one or more egress network interfaces that can carry traffic on this SDP are inconsistent.
Revert Time	Specifies the time to wait before reverting back from LDP to the configured LSPs, after having failed over to LDP.
Revert Count Down	Indicates the timer countdown before reverting back from LDP on this SDP. The timer countdown begins after the first configured LSP becomes active.
Flags	Displays all the conditions that affect the operating status of this SDP.
Class Forwarding	Indicates the admin state of class-based forwarding on this SDP. When the value is true, class-based forwarding is enabled.
EnforceDSTELspFc	Specifies whether service manager must validate with RSVP the support of the FC by the LSP.
Default LSP	Specifies the LSP ID that is used as a default when class-based forwarding is enabled on this SDP. This object must be set when enabling class-based forwarding.
Multicast LSP	Displays the LSP ID that all multicast traffic will be forwarded on when class-based forwarding is enabled on this SDP. When this object has its default value, multicast traffic will be forwarded on an LSP according to its forwarding class mapping.
Number of SDPs	Displays the metric to be used within the Tunnel Table Manager for decision making purposes. When multiple SDPs going to the same destination exist, this value is used as a tie-breaker by Tunnel Table Manager users like MP-BGP to select route with lower value.

Sample Output

```
*A:Dut-D# show service id 1 sdp 17407:4294967294 detail
===============================================================================
Service Destination Point (Sdp Id : 17407:4294967294) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 17407:4294967294  -(not applicable)
-------------------------------------------------------------------------------
Description     : (Not Specified)
```

```
SDP Id             : 17407:4294967294        Type              : VplsPmsi
Split Horiz Grp    : (Not Specified)
VC Type            : Ether                   VC Tag            : n/a
Admin Path MTU     : 9194                    Oper Path MTU     : 9194
Delivery           : MPLS
Far End            : not applicable
Tunnel Far End     : n/a                     LSP Types         : None
Hash Label         : Disabled                Hash Lbl Sig Cap  : Disabled
Oper Hash Label    : Disabled

Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 0                       Egress Label      : 3
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred           Oper ControlWord  : False
Last Status Change : 12/14/2012 12:42:22     Signaling         : None
Last Mgmt Change   : 12/14/2012 12:42:19     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
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : 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.      : 2979761                 E. Fwd. Octets    : 476761760

-------------------------------------------------------------------------
Control Channel Status
-------------------------------------------------------------------------
PW Status          : disabled                Refresh Timer     : <none>
Peer Status Expire : false                   Clear On Timeout  : true

MCAC Policy Name   :
MCAC Max Unconst BW: no limit                MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                       MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                       MCAC Avail Opnl BW: unlimited

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

-------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------
Class forwarding   : Disabled                EnforceDSTELspFc  : Disabled
Default LSP        : Uknwn                   Multicast LSP     : None

=========================================================================
FC Mapping Table
=========================================================================
FC Name             LSP Name
-------------------------------------------------------------------------
No FC Mappings


-------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------
Stp Admin State    : Down                    Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Forwarding
Port Number        : 0                       Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                     Designated Port Id: N/A

Fwd Transitions    : 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                       TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
-------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------
=========================================================================



*A:Dut-B#  show service sdp 204 detail
```

```
===============================================================================
Service Destination Point (Sdp Id : 204) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 204  -10.20.1.4
-------------------------------------------------------------------------------
Description           : (Not Specified)
SDP Id                : 204                    SDP Source         : manual
Admin Path MTU        : 0                      Oper Path MTU      : 1492
Delivery              : MPLS
Far End               : 10.20.1.4
Tunnel Far End        : n/a                    LSP Types          : RSVP

Admin State           : Up                     Oper State         : Up
Signaling             : TLDP                   Metric             : 0
Acct. Pol             : None                   Collect Stats      : Disabled
Last Status Change    : 02/12/2013 22:10:43    Adv. MTU Over.     : No
Last Mgmt Change      : 02/12/2013 22:09:55    VLAN VC Etype      : 0x8100
Bw BookingFactor      : 100                    PBB Etype          : 0x88e7
Oper Max BW(Kbps)     : 0                      Avail BW(Kbps)     : 0
Net-Domain            : default                Egr Interfaces     : Consistent
Flags                 : None

Mixed LSP Mode Information :
Mixed LSP Mode        : Disabled               Active LSP Type    : RSVP

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
-------------------------------------------------------------------------------
SDP-Groups
-------------------------------------------------------------------------------
red
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
Lsp Name              : lsp-b2d
Admin State           : Up                     Oper State         : Up
Time Since Last Tran*: 00h17m33s

-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding      : Disabled               EnforceDSTELspFc   : Disabled
Default LSP           : Uknwn                  Multicast LSP      : None

===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings

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

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


*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:Dut-A# show service sdp 1 detail
===============================================================================
Service Destination Point (Sdp Id : 1) Details
-------------------------------------------------------------------------------
 Sdp Id 1  -(10.20.1.3)
-------------------------------------------------------------------------------
Description           : epipe sdp 1 for lspId 00:00:00:01:00:00:00:00
SDP Id                : 1                     SDP Source          : manual
Admin Path MTU        : 0                     Oper Path MTU       : 1492
Far End               : 10.20.1.3             Delivery            : MPLS
Admin State           : Up                    Oper State          : Up
Signaling             : TLDP                  Metric              : 0
Acct. Pol             : None                  Collect Stats       : Disabled
Last Status Change    : 12/08/2008 22:54:30   Adv. MTU Over.      : No
Last Mgmt Change      : 12/08/2008 22:54:01   VLAN VC Etype       : 0x8100
Bw BookingFactor      : 100                   PBB Etype           : 0x88e7
Oper Max BW(Kbps)     : 1000                  Avail BW(Kbps)      : 1000
Flags                 : None

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              : tof1
Admin State           : Up                    Oper State          : Up
Time Since Last Tran* : 00h04m01s
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding      : disabled              EnforceDSTELspFc    : disabled
Default LSP           : Uknwn                 Multicast LSP       : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
```

```
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:Dut-A>config>service#


*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#
===============================================================================
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#

*A:MV-SR12>config>service>vprn# show service sdp 10 detail
===============================================================================
Service Destination Point (Sdp Id : 10) Details
===============================================================================
Sdp Id 10  -(200.20.1.201)
-------------------------------------------------------------------------------
Description           : (Not Specified)
SDP Id                : 10                    SDP Source         : manual
Admin Path MTU        : 0                     Oper Path MTU      : 9182
Far End               : 200.20.1.201          Delivery           : MPLS/LDP
Admin State           : Up                    Oper State         : Up
Signaling             : TLDP                  Metric             : 0
Acct. Pol             : None                  Collect Stats      : Disabled
Last Status Change    : 02/12/2010 22:37:08   Adv. MTU Over.     : No
Last Mgmt Change      : 02/12/2010 22:37:03   VLAN VC Etype      : 0x8100
Bw BookingFactor      : 100                   PBB Etype          : 0x88e7
Oper Max BW(Kbps)     : 0                     Avail BW(Kbps)     : 0
```

```
Net-Domain           : default              Egr Interfaces     : Consistent
Mixed LSP Mode       : Enabled
Revert Time          : 0                    Revert Count Down  : n/a
Flags                : None

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
-------------------------------------------------------------------------------
LDP Information :
-------------------------------------------------------------------------------
LDP LSP Id           : 65539                LDP Active         : No
-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP LIST :
Lsp Name             : To_7710
Admin State          : Up                   Oper State         : Up
Time Since Last Tran*: 01h20m56s
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding     : Disabled             EnforceDSTELspFc   : Disabled
Default LSP          : Uknwn                Multicast LSP      : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name              LSP Name
-------------------------------------------------------------------------------
No FC Mappings
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:MV-SR12>config>service>vprn#
```

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#
*A:Dut-C># show service sdp 1 pw-port
===============================================================================
Service Destination Point (Sdp Id 1 Pw-Port )
===============================================================================
SDP Binding port     : 1/1/3

SDP: 1 Pw-port: 11
-------------------------------------------------------------------------------
VC-Id                : 11                   Admin Status       : up
```

```
Encap               : dot1q                 Oper Status        : up
VC Type             : vlan                  Vlan VC Tag        : 0
Oper Flags          : (Not Specified)


SDP: 1 Pw-port: 44
-------------------------------------------------------------------------------
VC-Id               : 2                     Admin Status       : up
Encap               : dot1q                 Oper Status        : up
VC Type             : ether
Oper Flags          : (Not Specified)

-------------------------------------------------------------------------------
Entries found: 2
-------------------------------------------------------------------------------
*A:Dut-C> #


*A:Dut-C> # show service sdp 1 pw-port 44
===============================================================================
Service Destination Point (Sdp Id 1 Pw-Port 44)
===============================================================================
SDP Binding port    : 1/1/3
VC-Id               : 2                     Admin Status       : up
Encap               : dot1q                 Oper Status        : up
VC Type             : ether
Oper Flags          : (Not Specified)
===============================================================================
*A:Dut-C> #
```

ETH-CFM Show Commands

eth-cfm

Syntax **eth-cfm**

Context show

Description This command enables the context to display eth-cfm information.

association

Syntax **association** [*ma-index*] [**detail**]

Context show>eth-cfm

Description This command displays eth-cfm association information.

Parameters *ma-index* — Specifies the maintenance association (MA) index.

Values 1— 4294967295

detail — Displays detailed information for the eth-cfm association.

Output **Show eth-cfm Association Command Output —** The following table describes show eth-cfm association command output fields:

Label	Description
Md-index	Displays the the maintenance domain (MD) index.
Ma-index	Displays the the maintenance association (MA) index.
Name	Displays the part of the maintenance association identifier which is unique within the maintenance domain name.
CCM-interval	Displays the CCM transmission interval for all MEPs in the association.
Bridge-id	Displays the bridge-identifier value for the domain association.
MHF Creation	Displays the MIP half function (MHF) for the association.
Primary VLAN	Displays the primary bridge-identifier VLAN ID.
Num Vids	Displays the number of VIDs associated with the VLAN.
Remote Mep Id	Displays the remote maintenance association end point (MEP) identifier

Sample Output

```
*A:ALU-IPD#  show eth-cfm association
```

```
===============================================================================
CFM Association Table
===============================================================================
Md-index   Ma-index   Name                      CCM-interval Bridge-id
-------------------------------------------------------------------------------
1          1          abcabcabcabc1             1            1
1          2          abcabcabcabc2             1            2
1          3          abcabcabcabc3             1            3
1          4          abcabcabcabc4             1            4
===============================================================================
*A:ALU-IPD#
```

cfm-stack-table

Syntax **cfm-stack-table [{all-ports}] [level** <0..7>**] [direction <down>]**

Context show>eth-cfm

Description This command displays stack-table information. This stack-table is used to display the various management points MEPs and MIPs configured on the system. This can be Service based. The various options allow the operator to be specific. If no parameters are include then the entire stack-table will be displayed.

Parameters **port** *port-id* **—** Displays the bridge port or aggregated port on which MEPs or MHFs are configured.

vlan *vlan-id* **—** Displays the associated VLAN ID.

level — Display the MD level of the maintenance point.

Values 0 — 7

direction down — Displays the direction in which the MP faces on the bridge port.

Output **Show eth-cfm CFM Stack Table Command Output —** The following table describes show eth-cfm CFM stack table command output fields:

Label	Description
Sap	Displays associated SAP IDs.
Level Dir	Displays the MD level of the maintenance point.
Md-index	Displays the the maintenance domain (MD) index.
Ma-index	Displays the the maintenance association (MA) index.
Mep-id	Displays the integer that is unique among all the MEPs in the same MA.
Mac-address	Displays the MAC address of the MP.

Sample Output

```
*A:ALU-IPD#   show eth-cfm cfm-stack-table
```

```
===============================================================================
CFM SAP Stack Table
===============================================================================
Sap            Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
lag-1:1.1      0     Down 2          1          10     00:f3:f0:98:97:1b
lag-1:1.1      6     Down 1          1          1      00:f3:f0:98:97:1b
lag-1:2.2      0     Down 2          2          20     00:f3:f0:98:97:1b
lag-1:2.2      6     Down 1          2          2      00:f3:f0:98:97:1b
===============================================================================
===============================================================================
CFM Ethernet Tunnel Stack Table
===============================================================================
Eth-tunnel     Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
CFM SDP Stack Table
===============================================================================
Sdp            Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
===============================================================================
CFM Virtual Stack Table
===============================================================================
Service        Level Dir  Md-index   Ma-index   Mep-id Mac-address
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:ALU-IPD#
```

domain

Syntax **domain** [*md-index*] [**association** *ma-index* | **all-associations**] [**detail**]

Context show>eth-cfm

Description This command displays domain information.

Parameters *md-index* — Displays the index of the MD to which the MP is associated, or 0, if none.

association *ma-index* — Displays the index to which the MP is associated, or 0, if none.

all-associations — Displays all associations to the MD.

detail — Displays detailed domain information.

Output **Show eth-cfm Domain Command Output —** The following table describes show eth-cfm domain command output fields:

Label	Description
Md-index	Displays the Maintenance Domain (MD) index value.
Level	Displays an integer identifying the Maintenance Domain Level (MD Level). Higher numbers correspond to higher Maintenance Domains, those with the greatest physical reach, with the highest values for customers' CFM PDUs. Lower numbers correspond to lower Maintenance Domains, those with more limited physical reach, with the lowest values for CFM PDUs protecting single bridges or physical links.
Name	Displays a generic Maintenance Domain (MD) name.
Format	Displays the type of the Maintenance Domain (MD) name. Values include **dns**, **mac**, and *string*.

Sample Output

```
*A:7210-2#  show eth-cfm domain
===============================================================================
CFM Domain Table
===============================================================================
Md-index   Level Name                                          Format
-------------------------------------------------------------------------------
1          6                                                   none
2          0                                                   none
===============================================================================
*A:7210-2#
```

mep

Syntax **mep** *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**]
mep *mep-id* **domain** *md-index* **association** *ma-index* **remote-mepid** *mep-id* | **all-remote-mepids**
mep *mep-id* **domain** *md-index* **association** *ma-index* **eth-test-results** [**remote-peer** *mac-address*]
mep *mep-id* **domain** *md-index* **association** *ma-index* **one-way-delay-test** [**remote-peer** *mac-address*]
mep*mep-id* **domain** *md-index* **association** *ma-index* **two-way-delay-test** [**remote-peer** *mac-address*]
mep *mep-id* **domain** *md-index* **association** *ma-index* **two-way-slm-test [remote-peer** *macaddress***]**

Context show>eth-cfm

Description This command displays Maintenance Endpoint (MEP) information.

NOTES:

- The show eth-cfm mep mep-id domain md-id association ma-id command does not display CCM ERROR, CCM XCON frames in the output.
- The show eth-cfm mep mep-id domain md-id association ma-id remote-mep rmep-id command does not display some TLVs details.

Parameters *mep-id* — Displays the integer that is unique among all the MEPs in the same MA.

domain *md-index* **—** Displays the index of the MD to which the MP is associated, or 0, if none.

association *ma-index* — Displays the index to which the MP is associated, or 0, if none.

loopback — Displays loopback information for the specified MEP.

linktrace — Displays linktrace information for the specified MEP.

remote-mepid *mep-id* **—** Includes specified remote mep-id information for specified the MEP.

all-remote-mepids — Includes all remote mep-id information for the specified MEP.

eth-test-results — Includes eth-test-result information for the specified MEP.

one-way-delay-test — Includes one-way-delay-test information for the specified MEP.

two-way-delay-test — Includes two-way-delay-test information for the specified MEP.

two-way-slm-test — Includes two-way-slm-test information for the specified MEP.

remote-peer *mac-address* — Includes specified remote mep-id information for the specified MEP.

Sample Output

```
*A:7210-2#  show eth-cfm mep 1 domain 1 association 1 linktrace
-------------------------------------------------------------------------------
Mep Information
-------------------------------------------------------------------------------
Md-index           : 1                         Direction         : Down
Ma-index           : 1                         Admin             : Enabled
MepId              : 1                         CCM-Enable        : Enabled
```

```
IfIndex            : 1342177281              PrimaryVid        : 65537
FngState           : fngReset                ControlMep        : False
LowestDefectPri    : macRemErrXcon           HighestDefect     : none
Defect Flags       : None
Mac Address        : 00:f3:f0:98:97:1b       CcmLtmPriority    : 7
CcmTx              : 531                     CcmSequenceErr    : 0
Eth-1Dm Threshold  : 3(sec)
Eth-Ais:           : Disabled
Eth-Tst:           : Enabled                 Eth-Tst Pattern:  : allZerosNoCrc
Eth-Tst dataLength : 64                      Eth-Tst Priority: : 7
Eth-Tst Threshold  : 1(bitError)
CcmLastFailure Frame:
    None
XconCcmFailure Frame:
    None
-------------------------------------------------------------------------------
Mep Linktrace Message Information
-------------------------------------------------------------------------------
LtRxUnexplained    : 0                       LtNextSequence    : 2
LtStatus           : False                   LtResult          : False
TargIsMepId        : False                   TargMepId         : 0
TargMac            : 00:00:00:00:00:00       TTL               : 64
EgressId           : 00:00:00:f3:f0:98:97:1b SequenceNum       : 1
LtFlags            : useFDBonly
-------------------------------------------------------------------------------
Mep Linktrace Replies
-------------------------------------------------------------------------------
SequenceNum        : 1                       ReceiveOrder      : 1
Ttl                : 63                      Forwarded         : False
LastEgressId       : 00:00:00:f3:f0:98:97:1b TerminalMep       : True
NextEgressId       : 00:00:00:e0:b1:99:cb:46 Relay             : n/a
ChassisIdSubType   : unknown value (0)
ChassisId:
    None
ManAddressDomain:
    None
ManAddress:
    None
IngressMac         : 00:e0:b1:99:cb:46       Ingress Action    : ingOk
IngrPortIdSubType  : unknown value (0)
IngressPortId:
    None
EgressMac          : 00:00:00:00:00:00       Egress Action     : egrNoTlv
EgrPortIdSubType   : unknown value (0)
EgressPortId:
    None
Org Specific TLV:
    None
-------------------------------------------------------------------------------
*A:7210-2#


*A:7210-2#  show eth-cfm mep 1 domain 1 association 1 loopback
-------------------------------------------------------------------------------
Mep Information
-------------------------------------------------------------------------------
Md-index           : 1                       Direction         : Down
Ma-index           : 1                       Admin             : Enabled
MepId              : 1                       CCM-Enable        : Enabled
IfIndex            : 1342177281              PrimaryVid        : 65537
FngState           : fngReset                ControlMep        : False
```

```
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : None
Mac Address        : 00:f3:f0:98:97:1b        CcmLtmPriority    : 7
CcmTx              : 566                      CcmSequenceErr    : 0
Eth-1Dm Threshold  : 3(sec)
Eth-Ais:           : Disabled
Eth-Tst:           : Enabled                  Eth-Tst Pattern:  : allZerosNoCrc
Eth-Tst dataLength : 64                       Eth-Tst Priority: : 7
Eth-Tst Threshold  : 1(bitError)
CcmLastFailure Frame:
    None
XconCcmFailure Frame:
    None
-------------------------------------------------------------------------------
Mep Loopback Information
-------------------------------------------------------------------------------
LbRxReply          : 1                        LbRxBadOrder      : 0
LbRxBadMsdu        : 0                        LbTxReply         : 0
LbSequence         : 2                        LbNextSequence    : 2
LbStatus           : False                    LbResultOk        : True
DestIsMepId        : False                    DestMepId         : 0
DestMac            : 00:00:00:00:00:00        SendCount         : 0
VlanDropEnable     : True                     VlanPriority      : 7
Data TLV:
    None
-------------------------------------------------------------------------------
*A:7210-2#


*A:7210-2#  show eth-cfm mep 1 domain 1 association 1 remote-mepid 10
===============================================================================
Eth-CFM Remote-Mep Table
===============================================================================
R-mepId Rx CC  Rx Rdi Port-Tlv If-Tlv Peer Mac Addr     CCM status since
-------------------------------------------------------------------------------
10      True   False  Absent   Absent 00:e0:b1:99:cb:46 05/20/2010 12:59:55
===============================================================================
*A:7210-2#


*A:7210-2#  show eth-cfm mep 1 domain 1 association 1 eth-test-results
===============================================================
Eth CFM ETH-Test Result Table
===============================================================
                                  Current        Accumulate
                  FrameCount      ErrBits        ErrBits
Peer Mac Addr     ByteCount       CrcErrs        CrcErrs
---------------------------------------------------------------
00:e0:b1:99:cb:46 1               0              0
                  64              0              0
===============================================================


*A:7210-2#  show eth-cfm mep 1 domain 1 association 1 eth-test-results remote-peer
00:e0:b1:99:cb:46
===============================================================
Eth CFM ETH-Test Result Table
===============================================================
                                  Current        Accumulate
                  FrameCount      ErrBits        ErrBits
Peer Mac Addr     ByteCount       CrcErrs        CrcErrs
```

```
-------------------------------------------------------------
00:e0:b1:99:cb:46 1              0              0
                  64             0              0
=============================================================
*A:7210-2#


*A:7210-2#  show eth-cfm mep 1 domain 1 association 1 one-way-delay-test
=============================================================================
Eth CFM One-way Delay Test Result Table
=============================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
-----------------------------------------------------------------------------
00:e0:b1:99:cb:46     10000               10000
=============================================================================


*A:7210-2#  show eth-cfm mep 1 domain 1 association 1 one-way-delay-test remote-peer
00:e0:b1:99:cb:46
=============================================================================
Eth CFM One-way Delay Test Result Table
=============================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
-----------------------------------------------------------------------------
00:e0:b1:99:cb:46     10000               10000
=============================================================================
*A:7210-2#


*A:7210-2#  show eth-cfm mep 1 domain 1 association 1 two-way-delay-test
=============================================================================
Eth CFM Two-way Delay Test Result Table
=============================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
-----------------------------------------------------------------------------
00:e0:b1:99:cb:46     10000               10000
=============================================================================
*A:7210-2#


*A:7210-2# # show eth-cfm mep 1 domain 1 association 1 two-way-delay-test remote-peer
00:e0:b1:99:cb:46
=============================================================================
Eth CFM Two-way Delay Test Result Table
=============================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
-----------------------------------------------------------------------------
00:e0:b1:99:cb:46     10000               10000
=============================================================================
*A:7210-2#
A:dut-b# show eth-cfm mep 1 domain 1 association 1 linktrace
-------------------------------------------------------------------------------
Mep Information
-------------------------------------------------------------------------------
Md-index           : 1                        Direction         : Down
Ma-index           : 1                        Admin             : Enabled
MepId              : 1                        CCM-Enable        : Enabled
IfIndex            : 35946496                 PrimaryVid        : 1
FngState           : fngReset                 ControlMep        : False
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : None
```

```
Mac Address        : 00:25:ba:01:c3:6a        CcmLtmPriority    : 7
CcmTx              : 0                        CcmSequenceErr    : 0
Eth-1Dm Threshold  : 3(sec)
Eth-Ais:           : Disabled
Eth-Tst:           : Disabled
CcmLastFailure Frame:
    None
XconCcmFailure Frame:
    None
-------------------------------------------------------------------------------
Mep Linktrace Message Information
-------------------------------------------------------------------------------
LtRxUnexplained    : 0                        LtNextSequence    : 2
LtStatus           : False                    LtResult          : False
TargIsMepId        : False                    TargMepId         : 0
TargMac            : 00:00:00:00:00:00        TTL               : 64
EgressId           : 00:00:00:25:ba:01:c3:6a  SequenceNum       : 1
LtFlags            : useFDBonly
-------------------------------------------------------------------------------
Mep Linktrace Replies
-------------------------------------------------------------------------------
SequenceNum        : 1                        ReceiveOrder      : 1
Ttl                : 63                       Forwarded         : False
LastEgressId       : 00:00:00:25:ba:01:c3:6a  TerminalMep       : True
NextEgressId       : 00:00:00:25:ba:00:5e:bf  Relay             : rlyHit
ChassisIdSubType   : unknown value (0)
ChassisId:
    None
ManAddressDomain:
    None
ManAddress:
    None
IngressMac         : 00:25:ba:00:5e:bf        Ingress Action    : ingOk
IngrPortIdSubType  : unknown value (0)
IngressPortId:
    None
EgressMac          : 00:00:00:00:00:00        Egress Action     : egrNoTlv
EgrPortIdSubType   : unknown value (0)
EgressPortId:
    None
Org Specific TLV:
    None
A:dut-b#
A:dut-b#

A:dut-b# show eth-cfm mep 1 domain 1 association 1 loopback
-------------------------------------------------------------------------------
Mep Information
-------------------------------------------------------------------------------
Md-index           : 1                        Direction         : Down
Ma-index           : 1                        Admin             : Enabled
MepId              : 1                        CCM-Enable        : Enabled
IfIndex            : 35946496                 PrimaryVid        : 1
FngState           : fngReset                 ControlMep        : False
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : None
Mac Address        : 00:25:ba:01:c3:6a        CcmLtmPriority    : 7
CcmTx              : 0                        CcmSequenceErr    : 0
Eth-1Dm Threshold  : 3(sec)
Eth-Ais:           : Disabled
Eth-Tst:           : Disabled
```

```
CcmLastFailure Frame:
    None
XconCcmFailure Frame:
    None
-------------------------------------------------------------------------------
Mep Loopback Information
-------------------------------------------------------------------------------
LbRxReply          : 1                       LbRxBadOrder      : 0
LbRxBadMsdu        : 0                       LbTxReply         : 0
LbSequence         : 2                       LbNextSequence    : 2
LbStatus           : False                   LbResultOk        : True
DestIsMepId        : False                   DestMepId         : 0
DestMac            : 00:00:00:00:00:00       SendCount         : 0
VlanDropEnable     : True                    VlanPriority      : 7
Data TLV:
    None
A:dut-b#


*A:dut-b# show eth-cfm mep 1 domain 4 association 4 two-way-delay-test remote-peer
00:25:ba:00:5e:bf

===============================================================
Eth CFM Two-way Delay Test Result Table
===============================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
---------------------------------------------------------------
00:25:ba:00:5e:bf     507                 507
===============================================================
*A:dut-b#
*A:dut-b# show eth-cfm mep 1 domain 4 association 4 two-way-delay-test

===============================================================
Eth CFM Two-way Delay Test Result Table
===============================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
---------------------------------------------------------------
00:25:ba:00:5e:bf     507                 507
===============================================================
*A:dut-b#
*A:dut-a# show eth-cfm mep 2 domain 4 association 4  eth-test-results remote-peer
00:25:ba:01:c3:6a

===============================================================
Eth CFM ETH-Test Result Table
===============================================================
                                  Current        Accumulate
                  FrameCount      ErrBits        ErrBits
Peer Mac Addr     ByteCount       CrcErrs        CrcErrs
---------------------------------------------------------------
00:25:ba:01:c3:6a 6               0              0
                  384             0              0
===============================================================
*A:dut-a#
*A:dut-a# show eth-cfm mep 2 domain 4 association 4  eth-test-results

===============================================================
Eth CFM ETH-Test Result Table
===============================================================
                                  Current        Accumulate
                  FrameCount      ErrBits        ErrBits
```

```
Peer Mac Addr     ByteCount      CrcErrs        CrcErrs
------------------------------------------------------------
00:25:ba:01:c3:6a 6              0              0
                  384            0              0
============================================================
*A:dut-a# show eth-cfm mep 2 domain 4 association 4 one-way-delay-test remote-peer
00:25:ba:01:c3:6a

===================================================================
Eth CFM One-way Delay Test Result Table
===================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
-------------------------------------------------------------------
00:25:ba:01:c3:6a     402                 402
===================================================================
*A:dut-a#

*A:dut-a# show eth-cfm mep 2 domain 4 association 4 one-way-delay-test

===================================================================
Eth CFM One-way Delay Test Result Table
===================================================================
Peer Mac Addr         Delay (us)          Delay Variation (us)
-------------------------------------------------------------------
00:25:ba:01:c3:6a     402                 402
===================================================================
*A:dut-a#
```

Show output for two-way-slm-test

```
*A:7210SAS# show eth-cfm mep 1 domain 7 association 100 two-way-slm-test

================================================================================

Eth CFM Two-way SLM Test Result Table (Test-id: 1)

================================================================================

Peer Mac Addr      Remote MEP       Count     In Loss    Out Loss         Unack

--------------------------------------------------------------------------------

00:25:ba:0d:1e:12           2           1           0           0            0

================================================================================

*A:7210SAS#
```

connection-profile

Syntax **connection-profile [conn-prof-id] [associations]**

Context show

Description **Platforms Supported:** 7210 SAS-D, 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

This command displays connection profile information.

Parameters *conn-prof-id* — Specifies the connection profile ID.

Values 1 — 1000 - 7210 SAS-D
1 — 8000 - 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C

associations — Displays the SAP and the service ID that use this connection profile.

Output **The following table describes show connection-profile command output fields**

Label	Description
CP Index	Identifies the connection-profile.
Number of Members	Indicates the number of ATM connection profile members not applicable for 7210.
HasRange	Indicates whether VLAN range is configured or not

Sample Output

Show output for connection-profile

```
*7210SAS>show# connection-profile

===============================================================================
Connection Profile Summary Information
===============================================================================
CP Index  Number of HasRange
          Members
-------------------------------------------------------------------------------
1         0         Yes
2         0         Yes
3         0         Yes
5         0         Yes
6         0         Yes
100       0         Yes
200       0         Yes
300       0         Yes
400       0         Yes
500       0         Yes
600       0         Yes
700       0         Yes
800       0         Yes
900       0         Yes
===============================================================================
```

```
*7210SAS>show#
```

Show output for connection-profile associations

```
*A:7210SAS>show# connection-profile associations

===============================================================================
Connection Profile Summary Information
===============================================================================
CP Index  Number of HasRange
          Members
-------------------------------------------------------------------------------
1         0         No
===============================================================================
*A:7210SAS>show#
```

```
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------
--------
Number of SDPs : 1
=========================================================================
========
*A:ALA-B#
```

The following is an example when both sides have control word disabled:

```
*A:ALA-Dut-B>config>service>epipe# show service id 2100 sdp detail
=========================================================================
========
Services: Service Destination Points Details
=========================================================================
========
Sdp Id 1:2001  -(1.1.1.1)
-------------------------------------------------------------------------
--------
Description     : Default sdp description
SDP Id           : 1:2001                Type            : Spoke
VC Type          : Ether                 VC Tag          : n/a
Admin Path MTU    : 1600                 Oper Path MTU    : 1600
Far End           : 1.1.1.1                   Delivery         : MPLS
Admin State       : Up                   Oper State       : Up
Acct. Pol        : None                 Collect Stats    : Disabled
Ingress Label     : 115066               Egress Label     : 119068
Ing mac Fltr      : n/a                  Egr mac Fltr     : n/a
Ing ip Fltr       : n/a                  Egr ip Fltr      : n/a
Ing ipv6 Fltr     : n/a                  Egr ipv6 Fltr    : n/a
Admin ControlWord  : Not Preferred       Oper ControlWord  : False
Last Status Change : 02/05/2007 16:49:05    Signaling       : TLDP
Last Mgmt Change   : 02/05/2007 16:47:54
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip     : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr: No Limit             Total MAC Addr    : 0
Learned MAC Addr   : 0                         Static MAC Addr   : 0
MAC Learning      : Enabled              Discard Unkwn Srce: Disabled
MAC Aging         : Enabled
L2PT Termination   : Disabled                BPDU Translation  : Dis-
abled
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#
```

VPLS Show Commands

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.

Label	Description (Continued)
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.
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:7210-SASE# show service fdb-info
===============================================================================
Forwarding Database(FDB) Information
===============================================================================================
============ Service Id : 1 Mac Move : Disabled
Primary Factor : 3 Secondary Factor : 2
Mac Move Rate : 2 Mac Move Timeout : 10
Mac Move Retries : 3
Table Size : 250 Total Count : 1
Learned Count : 0 Static Count : 0
Host-learned Count: 1
Remote Age : 900 Local Age : 300
High Watermark : 95% Low Watermark : 90%
Mac Learning : Enabled Discard Unknown : Disabled
Mac Aging : Enabled Relearn Only : False
Mac Subnet Len : 48
-------------------------------------------------------------------------------
Total Service FDBs : 1
Total FDB Configured Size : 250
Total FDB Entries In Use : 1
PBB MAC Address Indices In Use : 0
-------------------------------------------------------------------------------
===============================================================================
A:7210-SASE#
```

fdb-mac

Syntax **fdb-mac** *ieee-address* [**expiry**]

Context show>service

Description This command displays the FDB entry for a given MAC address.

Parameters *ieee-address* — The 48-bit MAC address for which to display the FDB entry in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

expiry — Shows the time until the MAC is aged out.

Output **Show FDB-MAC Command Output —** The following table describes the show FDB MAC command output fields:

Label	Description
Service ID	The service ID number.
MAC	The specified MAC address
Source-Identifier	The location where the MAC is defined.
Type/Age	Static – FDB entries created by management. Learned – Dynamic entries created by the learning process. 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.

Sample Output

```
*A:ALA-12# show service fdb-mac 00:99:00:00:00:00
===============================================================================
Services Using Forwarding Database Mac 00:99:00:00:00:00
===============================================================================
ServId  MAC                    Source-Identifier       Type/Age Last Change
-------------------------------------------------------------------------------
1       00:99:00:00:00:00             sap:1/2/7:0               Static
===============================================================================
*A:ALA-12#
```

sap-using

Syntax **sap-using** [**ingress** | **egress**] **filter** *filter-id*
sap-using [**sap** *sap-id*]

Context show>service

Description This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters **ingress** — Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

filter *filter-id* — The ingress or egress filter policy ID for which to display matching SAPs.

Values 1 — 65535

sap-id — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 for command syntax.

Output **Show Service SAP —** The following table describes show service SAP output fields:

Label	Description
Port ID	The ID of the access port where the SAP is defined.
Svc ID	The service identifier.
I.QoS	The SAP ingress QoS policy number specified on the ingress SAP.
I.MAC/IP	The MAC or IP filter policy ID applied to the ingress SAP.
Egr. Fltr	The filter policy ID applied to the egress SAP.
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The administrative state of the SAP.
Opr	The actual state of the SAP.

Sample Output

```
*A:ALU_SIM2>config>service>vpls# show service sap-using
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.   Ing.    Egr.    Adm  Opr
                                           QoS    Fltr    Fltr
-------------------------------------------------------------------------------
1/1/1:10                        1          1      none    none    Up   Up
1/1/3:500.*                     1          1      none    none    Up   Up
1/1/1:200                       200        1      none    none    Up   Up
1/1/3:100.200                   200        1      none    none    Up   Up
1/1/1:300                       300        1      none    none    Up   Up
```

```
-------------------------------------------------------------------------------
Number of SAPs : 5
-------------------------------------------------------------------------------
*A:ALU_SIM2>config>service>vpls#
```

service-using

Syntax **service-using** [**epipe**] [**ies**] [**vpls**] [**mirror**] [**customer** *customer-id*]

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.

customer *customer-id* **—** Displays services only associated with the specified customer ID.

Default Services associated with a customer.

Values 1 — 2147483647

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.
Adm	The administrative state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

Sample Output

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

```
-------------------------------------------------------------------------------
Matching Services : 3
===============================================================================
*A:ALA-12#
```

```
*A:ALA-12# show service service-using epipe
===============================================================================
Services [epipe]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
6            Epipe     Up     Up         6                 09/22/2006 23:05:58
7            Epipe     Up     Up         6                 09/22/2006 23:05:58
8            Epipe     Up     Up         3                 09/22/2006 23:05:58
103          Epipe     Up     Up         6                 09/22/2006 23:05:58
-------------------------------------------------------------------------------
Matching Services : 4
===============================================================================
*A:ALA-12#
```

```
*A:ALA-14# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
10           mVPLS     Down   Down       1                 10/26/2006 15:44:57
11           mVPLS     Down   Down       1                 10/26/2006 15:44:57
100          mVPLS     Up     Up         1                 10/26/2006 15:44:57
101          mVPLS     Up     Up         1                 10/26/2006 15:44:57
102          mVPLS     Up     Up         1                 10/26/2006 15:44:57
-------------------------------------------------------------------------------
Matching Services : 5
-------------------------------------------------------------------------------
*A:ALA-14#
```

```
*A:SetupCLI# show service service-using
===============================================================================
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
...
-------------------------------------------------------------------------------
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
...
-------------------------------------------------------------------------------
Matching Services : 27
-------------------------------------------------------------------------------
*A:SetupCLI#
```

id

Syntax **id** *service-id*

Context show>service

Description This command displays information for a particular service-id.

Parameters *service-id* — The unique service identification number that identifies the service in the service domain.

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

all — Display detailed information about the service.

arp — Display ARP entries for the service.

base — Display basic service information.

fdb — Display FDB entries.

igmp-snooping — Display IGMP snooping information.

interface — Display service interfaces.

mstp-configuration — - Display MSTP information.

sap — Display SAPs associated to the service.

stp — Display STP information.

all

Syntax **all**

Context show>service>id

Description This command displays detailed information for all aspects of the service.

Output **Show service ID all output —** The following table describes the command output fields.

Label	Description
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.
Service Type	Specifies the type of service.
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.

Label	Description (Continued)
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: MPLS.
Admin State	The administrative state of this SDPSAP.
Oper State	The operational state of this SDPSAP.
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 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.

Label	Description (Continued)
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.
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.
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.

Label	Description (Continued)
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.
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.

Label	Description (Continued)
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 SAP.
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.

Sample Output

```
*A:7210-E> show service id 1 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 05/31/2002 19:26:08
Last Mgmt Change  : 06/07/2002 18:14:58
Admin State       : Down                Oper State        : Down
SAP Count         : 1
Uplink Type:      : L2
SAP Type:         : Null-star           Customer vlan:    : n/a
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
SAP 1/1/1
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/1/1                   Encap             : null
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Last Status Change : 05/31/2002 19:26:08
Last Mgmt Change   : 06/07/2002 18:15:31
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
```

```
Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
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
tod-suite          : None
Ing Agg Rate Limit : max                     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

Acct. Pol          : None                    Collect Stats     : Disabled
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Discarding
Port Number        : 2048                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : N/A
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                     Designated Port   : N/A

Forward transitions: 0                       Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                       Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                       TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                       RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                       Min Auth Interval : 15 minutes
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
-------------------------------------------------------------------------------
Ingress QoS Classifier Usage
-------------------------------------------------------------------------------
Classifiers Allocated: 16                   Meters Allocated   : 8
Classifiers Used     : 2                    Meters Used        : 2
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Stats:          0                       0
Egress Stats:           0                       0
Ingress Drop Stats:     0                       0
Extra-Tag Drop Stats:   n/a                     n/a
-------------------------------------------------------------------------------
Sap per Meter stats (forward/drop counter mode)
-------------------------------------------------------------------------------
```

```
                        Packets                 Octets
Ingress Meter 1 (Unicast)
Total Forwarded       : 0                       0
Total Dropped         : 0                       0

Ingress Meter 11 (Multipoint)
Total Forwarded       : 0                       0
Total Dropped         : 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          : 80:00.00:a0:ff:43:21:cb  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 : 1 Mac Move : Disabled
Primary Factor : 3 Secondary Factor : 2
Mac Move Rate : 2 Mac Move Timeout : 10
Mac Move Retries : 3
Table Size : 250 Total Count : 1
Learned Count : 0 Static Count : 0
Host-learned Count: 1
Remote Age : 900 Local Age : 300
High Watermark : 95% Low Watermark : 90%
Mac Learning : EnabledDiscard Unknown : Disabled
Mac Aging : Enabled Relearn Only : False
Mac Subnet Len : 48
-------------------------------------------------------------------------------
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Sap                    Oper    MRtr Send    Max  Max  Num
Id                     State   Port Queries Grps Srcs Grps
-------------------------------------------------------------------------------
sap:1/1/1              Down    No   No      None None 0
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:7210-E>
```

```
Ignore Standby Sig : False                    Block On Mesh Fail: False
```

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 learningprocess.
	Other — Local entries for the IP interfaces created.
Expiry	The age of the ARP entry.
Interface	The interface applied to the service.
SAP	The SAP ID.

base

Syntax **base** [**msap**]

Context show>service>id

Description This command displays basic information about the service ID including service type, description, and SAPs.

Output **Show Service-ID Base —** The following table describes show service-id base output fields:

Label	Description
Service Id	The service identifier.
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 port can handle.
Adm	The largest frame size (in octets) that the SAP can handle.
SAP Count	The number of SAPs defined on the service.
SAP Type	The type of SAPs allowed in the service. It also describes the applied processing by the node to the packets received on these SAPs.
Identifier	Specifies the service access (SAP).
OprMTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this port, without requiring the packet to be fragmented.
Opr	The operating state of the SAP

Sample Output

```
A:ALU-7210>config>service>vpls# show service id 700 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 700                Vpn Id            : 0
```

```
Service Type      : VPLS
Customer Id       : 1
Last Status Change: 09/19/2005 16:25:28
Last Mgmt Change  : 09/21/2005 14:07:07
Admin State       : Down                Oper State        : Down
SAP Count         : 1
Uplink Type:      : L2
SAP Type:         : Null-star           Customer vlan:    : n/a
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/2                                null         1514    1514    Up   Down
===============================================================================
A:ALU-7210>config>service>vpls#
*A:ALU_SIM2>config>service>vpls# show service  id 600 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 600                 Vpn Id            : 0
Service Type      : uVPLS
Customer Id       : 1
Last Status Change: 11/17/2008 00:12:16
Last Mgmt Change  : 11/14/2008 03:36:21
Admin State       : Up                  Oper State        : Down
SAP Count         : 3
Uplink Type:      : L2
SAP Type:         : Dot1q               Customer vlan:    : n/a
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap:1/1/1:950                            q-tag        1518    1518    Down Down
sap:1/1/2:1450                           q-tag        1518    1518    Down Down
sap:1/1/3:3000.*                         qinq         1522    1522    Down Down
===============================================================================
*A:ALU_SIM2>config>service>vpls#

*A:7210SAS-D>show>service# id 15 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 15                  Vpn Id            : 0
Service Type      : Epipe
Description       : (Not Specified)
Customer Id       : 40
Last Status Change: 01/01/1970 00:00:14
Last Mgmt Change  : 01/11/1970 23:34:24
Admin State       : Down                Oper State        : Down
SAP Count         : 1
Uplink Type:      : L2
SAP Type:         : Any                 Customer vlan:    : n/a

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
```

```
sap:1/1/1                                 null         1514    1514    Up   Down
===============================================================================
*A:SAS-D>show>service#
*A:7210SAS# show service id 10 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 10                  Vpn Id            : 0
Service Type      : VPLS
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 02/06/2106 06:28:12
Last Mgmt Change  : 01/10/1970 01:55:31
Admin State       : Down                Oper State        : Down
MTU               : Not Applicable      Def. Mesh VC Id   : 10
SAP Count         : 0
Uplink Type:      : L2
SAP Type:         : Dot1q Range         Customer vlan:    : n/a

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                                Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:7210SAS# show service id 10 base
```

bgp-ad

Syntax **bgp-ad**

Context show>service>id

Description This command shows BGP AD related information for the service

Output **Show Service-ID BGP-AD —** The following table describes show service-id bgp-ad output fields:

Label	Description
Admin state	
VPLS ID	
Prefix	

```
*A:dut-a>show>service# id 5 bgp-ad

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

```
BGP Auto-discovery Information
-------------------------------------------------------------------------------
Admin State       : Up
Vpls Id           : 100:5
Prefix            : 1.1.1.1
-------------------------------------------------------------------------------
*A:dut-a>show>service#
```

fdb

Syntax **fdb** [**sap** *sap-id* [**expiry**]] | [**mac** *ieee-address* [**expiry**]] | [**detail**] [**expiry**]

Context show>service>id
show>service>fdb-mac

Description This command displays FDB entries for a given MAC address.

Parameters **sap** *sap-id* — Specifies the physical port identifier portion of the SAP. See Common CLI Command Descriptions on page 945 for command syntax.

detail — Displays detailed information.

expiry — Displays time until MAC is aged out.

Show FDB Information — The following table describes service FDB output fields.

Label	Description
ServID	Displays the service ID.
MAC	Displays the associated MAC address.
Mac Move	Displays the administrative state of the MAC movement feature associated with this service.
Primary Factor	Displays a factor for the primary ports defining how many MAC-relearn periods should be used to measure the MAC-relearn rate.
Secondary Factor	Displays a factor for the secondary ports defining how many MAC-relearn periods should be used to measure the MAC-relearn rate.
Mac Move Rate	Displays the maximum rate at which MAC's can be re-learned in this service, before the SAP where the moving MAC was last seen is automatically disabled in order to protect the system against undetected loops or duplicate MAs. The rate is computed as the maximum number of re-learns allowed in a 5 second interval: for example, the default rate of 2 re-learns per second corresponds to 10 re-learns in a 5 second period.
Mac Move Timeout	Displays the time in seconds to wait before a SAP that has been disabled after exceeding the maximum re-learn rate is re-enabled. A value of zero indicates that the SAP will not be automatically re-enabled after being disabled. If after the SAP is re-enabled it is disabled again, the effective retry timeout is doubled in order to avoid thrashing.
Mac Move Retries	Displays the number of times retries are performed for reenabling the SAP.

Label	Description
Table Size	Specifies the maximum number of learned and static entries allowed in the FDB of this service.
Total Count	Displays the total number of learned entries in the FDB of this service.
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 SAP. 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
	Age – Specifies the number of seconds used to age out TLS FDB entrieslearned 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.

```
A:ALU-7210# show service id 1 fdb
===============================================================================
Forwarding Database, Service 1
===============================================================================
Service Id : 1     Mac Move : Disabled
Primary Factor : 3 Secondary Factor : 2
Mac Move Rate : 2 Mac Move Timeout : 10
Mac Move Retries : 3
Table Size : 250 Total Count : 1
Learned Count : 0 Static Count : 0
Host-learned Count: 1
Remote Age : 900 Local Age : 300
High Watermark : 95% Low Watermark : 90%
Mac Learning : Enabled Discard Unknown : Disabled
Mac Aging : Enabled Relearn Only : False
Mac Subnet Len : 48
===============================================================================
A:ALU-7210#
```

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.
SAP Id	Specifies the SAP ID.

Sample:

```
*A:7210SAS>show>service# id 1 l2pt detail

===============================================================================
L2pt details, Service id 1
===============================================================================

Service Access Points
-------------------------------------------------------------------------------
SapId                       L2pt-                      Admin Bpdu-  Oper Bpdu-
                            termination                translation  translation
-------------------------------------------------------------------------------
1/1/1                       stp cdp vtp dtp pagp udld  disabled     disabled
-------------------------------------------------------------------------------
Number of SAPs : 1
```

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

===============================================================================
L2pt summary, Service id 1
===============================================================================
       L2pt-term   L2pt-term   Bpdu-trans  Bpdu-trans  Bpdu-trans  Bpdu-trans
       enabled     disabled    auto        disabled    pvst        stp
-------------------------------------------------------------------------------
SAP's  1           0           0           1           0           0
SDP's  0           0           0           0           0           0
-------------------------------------------------------------------------------
Total  1           0           0           1           0           0
===============================================================================
*A:7210SAS>show>service#
```

mac-move

Syntax **mac-move**

Context show>service>id

Description This command displays MAC move related information about the service.

mrouters

Syntax **mrouters** [**detail**]

Context show>service>id>mld-snooping

Description This command displays all multicast routers.

sap

Syntax **sap** *sap-id* [*sap-id* [**detail**|**stp**]]

Context show>service>id

Description This command displays information for the SAPs associated with the service.

If no optional parameters are specified, a summary of all associated SAPs is displayed.

Parameters **sap** *sap-id* — The ID that displays SAPs for the service in the *slot*/*mda*/*port*[**.***channel*] form. See Common CLI Command Descriptions on page 945 for command syntax.

detail — Displays detailed information for the SAP.

stp - — Displays the stp information of the SAP.

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, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch,RelearnLimitExceeded, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
Last Status Change	Specifies the time of the most recent operating status change to this SAP
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SAP.
Ingress qos-policy	The ingress 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.
SAP per Meter stats	
Ingress Meter	Specifies the meter ID.
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 and below PIR) forwarded by the ingress meter.
Ingress TD Profile	The profile ID applied to the ingress SAP.
Egress TD Profile	The profile ID applied to the egress SAP.

Label	Description (Continued)
Alarm Cell Handling	The indication that OAM cells are being processed.
AAL-5 Encap	The AAL-5 encapsulation type.
The labels and description listed below are not supported on 7210 SAS-E devices:	
Loopback Mode	Displays the Ethernet port loopback mode
Loopback Src Addr	Displays the configured loopback source address
Loopback Dst Addr	Displays the configured loopback destination address
No-svc-port used	Displays the port ID of the port on which no service is configured. This port is used for the port loop back with MAC swap functionality.

Sample Output

```
*A:ALU_SIM2>config>service# show service  id 1  sap 1/1/3:500.*  detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/1/3:500.*              Encap             : qinq
QinQ Dot1p         : Default
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Admin State        : Up                       Oper State        : Down
Flags              : PortOperDown
Last Status Change : 11/17/2008 00:26:56
Last Mgmt Change   : 11/14/2008 02:45:15
Loopback Mode : Internal    No-svc-port used : 1/1/13
Loopback Src Addr : 00:00:00:22:22:22
Loopback Dst Addr : 00:00:00:11:11:11
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
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          : TodResourceUnavail       qinq-pbit-marking : both
Egr Agg Rate Limit : max
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Disabled
L2PT Termination   : Disabled                 BPDU Translation  : Disabled
Vlan-translation   : None

Acct. Pol          : None                     Collect Stats     : Disabled

Anti Spoofing      : None                     Nbr Static Hosts  : 0
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Mac Move           : Blockable
Stp Admin State    : Up                       Stp Oper State    : Up
Core Connectivity  : Down
Port Role          : Disabled                 Port State        : Discarding
```

```
Port Number        : 2049                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : False
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Rstp
Last BPDU from     : N/A
CIST Desig Bridge  : This Bridge              Designated Port   : 34817

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      : 124267
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
                      Packets
Ingress Packets rcvd:   0
-------------------------------------------------------------------------------
Sap per Meter stats
-------------------------------------------------------------------------------
                      Packets                 Octets
Ingress Meter 1 (Unicast)
For. InProf          : 0                      0
For. OutProf         : 0                      0

Ingress Meter 11 (Multipoint)
For. InProf          : 0                      0
For. OutProf         : 0                      0
===============================================================================
*A:ALU_SIM2>config>service#


*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#
```

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.

Output **Show Service-ID STP Output —** The following table describes show service-id STP output fields:

Label	Description
Bridge-id	Specifies the MAC address used to identify this bridge in the network.
Bridge fwd delay	Specifies how fast a bridge changes its state when moving toward the forwarding state.
Bridge Hello time	Specifies the amount of time between the transmission of Configuration BPDUs.
Bridge max age	Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded. This is the actual value that this bridge is currently using.
Bridge priority	Defines the priority of the Spanning Tree Protocol instance associated with this service.
Topology change	Specifies whether a topology change is currently in progress.
Last Top. change	Specifies the time (in hundredths of a second) since the last time a topology change was detected by the Spanning Tree Protocol instance associated with this service.

Label	Description (Continued)
Top. change count	Specifies the total number of topology changes detected by the Spanning Tree Protocol instance associated with this service since the management entity was last reset or initialized.
Root bridge-id	Specifies the bridge identifier of the root of the spanning tree as determined by the Spanning Tree Protocol instance associated with this service. This value is used as the Root Identifier parameter in all Configuration BPDUs originated by this node.
Root path cost	Specifies the cost of the path to the root bridge as seen from this bridge.
Root forward delay	Specifies how fast the root changes its state when moving toward the forwarding state.
Rcvd hello time	Specifies the amount of time between the transmission of configuration BPDUs.
Root max age	Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded.
Root priority	This object specifies the priority of the bridge that is currently selected as root-bridge for the network.
Root port	Specifies the port number of the port which offers the lowest cost path from this bridge to the root bridge.
SAP Identifier	The ID of the access port where this SAP is defined.
BPDU encap	Specifies the type of encapsulation used on BPDUs sent out and received on this SAP.
Port Number	Specifies the value of the port number field which is contained in the least significant 12 bits of the 16-bit port ID associated with this SAP.
Port 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.
Port Path Cost	Specifies the contribution of this port to the path cost of paths towards the spanning tree root which include this port.
Designated Port	Specifies the port identifier of the port on the designated bridge for this port's segment.
Designated Bridge	Specifies the bridge identifier of the bridge which this port considers to be the designated bridge for this port's segment.

Sample Output

```
*A:ALU_SIM2>config>service>vpls>sap# show service id 1 stp
===============================================================================
Stp info, Service 1
===============================================================================
```

```
Bridge Id          : 00:00.00:45:67:32:10:ab  Top. Change Count : 2
Root Bridge        : This Bridge              Stp Oper State    : Up
Primary Bridge     : N/A                      Topology Change   : Inactive
Mode               : Rstp                     Last Top. Change  : 1d 18:34:36
Vcp Active Prot.   : N/A
Root Port          : N/A                      External RPC      : 0
===============================================================================
Stp port info
===============================================================================
Sap Id             Oper-     Port-      Port-       Port-  Oper-  Link-  Active
                   State     Role       State       Num    Edge   Type   Prot.
-------------------------------------------------------------------------------
1/1/21:1           Up        Designated Forward     2048   True   Pt-pt  Rstp
lag-1:1.1          Up        Designated Forward     2049   False  Pt-pt  Rstp
lag-2:1.1          Up        Designated Forward     2050   False  Pt-pt  Rstp
===============================================================================
*A:ALU_SIM2>config>service>vpls>sap#


*A:ALU_SIM2>config>service>vpls>sap#  show service id 1 stp detail
===============================================================================
Spanning Tree Information
===============================================================================
-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Up
Stp Admin State    : Up                       Stp Oper State    : Up
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 00:00.00:45:67:32:10:ab  Bridge Instance Id: 0
Bridge Priority    : 0                        Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 1d 18:35:18              Bridge Max Age    : 20
Top. Change Count  : 2                        Bridge Fwd Delay  : 15

Root Bridge        : This Bridge
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 0                        Root Port         : N/A
-------------------------------------------------------------------------------
Spanning Tree Sap Specifics
-------------------------------------------------------------------------------
SAP Identifier     : 1/1/21:1                 Stp Admin State   : Up
Port Role          : Designated               Port State        : Forwarding
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : True
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : Rstp
Last BPDU from     : N/A
CIST Desig Bridge  : This Bridge              Designated Port   : 34816
Forward transitions: 2                        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      : 79126
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0

SAP Identifier     : lag-1:1.1                Stp Admin State   : Up
```

```
Port Role          : Designated              Port State        : Forwarding
Port Number        : 2049                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : False
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : Rstp
Last BPDU from     : 10:00.00:f3:f0:98:97:00
CIST Desig Bridge  : This Bridge             Designated Port   : 34817
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     : 5                       RST BPDUs tx      : 79128
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 0

SAP Identifier     : lag-2:1.1               Stp Admin State   : Up
Port Role          : Designated              Port State        : Forwarding
Port Number        : 2050                    Port Priority     : 128
Port Path Cost     : 10                      Auto Edge         : Enabled
Admin Edge         : Disabled                Oper Edge         : False
Link Type          : Pt-pt                   BPDU Encap        : Dot1d
Root Guard         : Disabled                Active Protocol   : Rstp
Last BPDU from     : 20:00.00:e0:b1:99:cb:2a
CIST Desig Bridge  : This Bridge             Designated Port   : 34818
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     : 6                       RST BPDUs tx      : 78760
MST BPDUs rcvd     : 0                       MST BPDUs tx      : 0
===============================================================================
*A:ALU_SIM2>config>service>vpls>sap#
```

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.

Output **Show Output —** The following table describes the show all command output fields:

Label	Description
Region Name	Displays the MSTP region name.
Region Revision	Displays the MSTP region revision.
MST Max Hops	Displays the MSTP maximum hops specified.
Instance	Displays the MSTP instance number.
Priority	Displays the MSTP priority.
Vlans mapped	Displays the VLAN range of the MSTP instance.

Sample output with MSTP information for 7210 SAS-E:

```
A:7210-SASE>show>service>id# stp mst-instance 1

===============================================================================
MSTP specific info for service 1 MSTI 1
===============================================================================
Regional Root      : 80:01.00:25:ba:02:8a:30  Root Port         : 2048
Internal RPC       : 20                       Remaining Hopcount: 18

===================================================================
MSTP port info for MSTI 1
===================================================================
Sap Id             Oper-     Port-      Port-       Port-  Same
                   State     Role       State       Num    Region
-------------------------------------------------------------------
1/1/1:0.*          Up        Root       Forward     2048   True
1/1/5:0.*          Up        Designated Forward     2049   True
1/1/10:0.*         Up        Designated Forward     2050   True
1/1/17:0.*         Up        Designated Forward     2051   True
===================================================================
A:7210-SASE>show>service>id#
```

Sample output with MSTP information for 7210 SAS-D:

```
A:SASD1>show>service>id# stp mst-instance 1

===============================================================================
MSTP specific info for service 1 MSTI 1
===============================================================================
Regional Root      : This Bridge              Root Port         : N/A
Internal RPC       : 0                        Remaining Hopcount: 20

===================================================================
MSTP port info for MSTI 1
===================================================================
Sap Id             Oper-     Port-      Port-       Port-  Same
                   State     Role       State       Num    Region
-------------------------------------------------------------------
1/1/2:0.*          Up        Designated Forward     2048   True
1/1/5:0.*          Up        Alternate  Discard     2049   False
1/1/6:0.*          Up        Master     Forward     2050   False
===================================================================
A:SASD1>show>service>id#
```

dhcp

Syntax **dhcp**

Context show>service>id

Description This command enables the context to display DHCP information for the specified service.

statistics

Syntax **statistics** [**sap** *sap-id*]
statistics [sdp *sdp-id:vc-id***]**
statistics [**interface** *interface-name*]

Context show>service>id>dhcp

Description Displays DHCP statistics information.

Parameters **sap** *sap-id* — Specifies the physical port identifier portion of the SAP definition.

sdp-id — The SDP identifier.

Values 1 — 17407

interface *interface-name* — Displays information for the specified IP interface.

Output **Show DHCP Statistics Output —** The following table describes the output fields for DHCP statistics.

Label	Description
Received Packets	The number of packets received from the DHCP clients. Includes DHCP packets received from both DHCP client and DHCP server.
Transmitted Packets	The number of packets transmitted to the DHCP clients. Includes DHCP packets transmitted from both DHCP client and DHCP server.
Received Malformed Packets	The number of corrupted/invalid packets received from the DHCP clients. Includes DHCP packets received from both DHCP client and DHCP server
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients. In this case, a frame is dropped due to the client sending a DHCP packet with Option 82 filled in before “trust” is set under the DHCP interface command.
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded.
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded.

Label	Description
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Server Packets Discarded	The number of packets received from the DHCP server that were discarded.
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded.
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.

```
*A:7210SAS>show>service>id>dhcp# statistics

===============================================================================
DHCP Global Statistics, service 1
===============================================================================
Rx Packets                          : 416554
Tx Packets                          : 206405
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 0
Client Packets Discarded            : 0
Client Packets Relayed              : 221099
Client Packets Snooped              : 0
Client Packets Proxied (RADIUS)     : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded            : 0
Server Packets Relayed              : 195455
Server Packets Snooped              : 0
DHCP RELEASEs Spoofed               : 0
DHCP FORCERENEWs Spoofed            : 0
===============================================================================
*A:7210SAS>show>service>id>dhcp#
```

summary

Syntax **summary [interface** *interface-name]*

Context show>service>id>dhcp

Description Displays DHCP configuration summary information.

Parameters *interface interface-name* — Displays information for the specified IP interface.

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. 7210 SAS does not support ARP populate.

Label	Description (Continued)
Used/Provided	7210 SAS does not maintain lease state.
Info Option	Indicates whether Option 82 processing is enabled on the interface.
Admin State	Indicates the administrative state.

Sample Output

```
A:7210SAS# show service id 1 dhcp summary
DHCP Summary, service 1
===============================================================================

Interface Name                   Arp      Used/                Info    Admin
  SapId/Sdp                      Populate Provided             Option  State
-------------------------------------------------------------------------------
egr_1                            No       0/0                  Replace Up
i_1                              No       0/0                  Replace Up
-------------------------------------------------------------------------------
Interfaces: 2
===============================================================================

*A:7210SAS>show>service>id>dhcp#
```

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 MVR Groups	Specifies the actual number of multicast groups that can be joined on this SAP.
MVR From VPLS Cfg Drops	Displays the from VPLS drop count.
MVR To SAP Cfg Drops	Displays the to SAP drop count.
MVR Admin State	Displays the administrative state of MVR.

Label	Description (Continued)
MVR Policy	The MVR policy name.

Sample Output

```
*A:7210-2>show>service>id>snooping# all
===============================================================================
IGMP Snooping info for service 900
===============================================================================
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Up
Querier     : 5.5.5.5 on SAP 1/1/14:100
-------------------------------------------------------------------------------
Sap                      Oper    MRtr  Send      Max Num   Num
Id                       State   Port  Queries   Groups    Groups
-------------------------------------------------------------------------------
sap:1/1/13:100           Up      No    Disabled  No Limit  0
sap:1/1/14:100           Up      No    Disabled  No Limit  0
sap:1/1/17:100           Up      No    Disabled  No Limit  1
sap:1/1/18:100           Up      No    Disabled  No Limit  1
-------------------------------------------------------------------------------
IGMP Snooping Querier info
-------------------------------------------------------------------------------
Sap Id                  : 1/1/14:100
IP Address              : 10.10.10.2
Expires                 : 254s
Up Time                 : 0d 00:02:42
Version                 : 2

General Query Interval  : 125s
Query Response Interval : 10.0s
Robust Count            : 2
-------------------------------------------------------------------------------
IGMP Snooping Multicast Routers
-------------------------------------------------------------------------------
MRouter          Sap Id                 Up Time        Expires   Version
-------------------------------------------------------------------------------
10.10.10.2       1/1/14:100             0d 00:02:44    254s      2
-------------------------------------------------------------------------------
Number of mrouters: 1
-------------------------------------------------------------------------------
IGMP Snooping Proxy-reporting DB
-------------------------------------------------------------------------------
Group Address    Up Time
-------------------------------------------------------------------------------
226.6.6.6        0d 00:02:17
229.9.9.9        0d 00:02:56
-------------------------------------------------------------------------------
Number of groups: 2
-------------------------------------------------------------------------------
IGMP Snooping SAP 1/1/13:100 Port-DB
-------------------------------------------------------------------------------
Group Address  Type  From-VPLS Up Time    Expires
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
```

```
IGMP Snooping SAP 1/1/14:100 Port-DB
-------------------------------------------------------------------------------
Group Address   Type    From-VPLS  Up Time       Expires
===============================================================================
*A:7210-2>show>service>id>snooping#
```

mfib

Syntax **mfib** [**brief**]
mfib [**group** *grp-address*]

Context show>service>id

Description This command displays the multicast FIB on the VPLS service.

Parameters **brief** — Displays a brief output.

group grp *grp-address* — Displays the multicast FIB for a specific multicast group address.

Output **Show Output —** The following table describes the command output fields:

Label	Description
Group Address	IPv4 multicast group address.
SAP ID	Indicates the SAP to which the corresponding multicast stream will be forwarded/blocked.
Forwarding/Blocking	Indicates whether the corresponding multicast stream will be blocked/forwarded.
Number of Entries	Specifies the number of entries in the MFIB.
Forwarded Packets	Indicates the number of multicast packets forwarded for the corresponding source/group.
Forwarded Octets	Indicates the number of octets forwarded for the corresponding source/group.
Svc ID	Indicates the service to which the corresponding multicast stream will forwarded/blocked. Local means that the multicast stream will be forwarded/blocked to a SAP local to the service.

Sample Output

```
*A:SAS# show service id 1 mfib

===============================================================================
Multicast FIB, Service 1
===============================================================================
Group Address         Sap/Sdp Id               Svc Id   Fwd/Blk
-------------------------------------------------------------------------------
224.4.4.4             sap:1/1/1                Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 1
===============================================================================
*A:7210-2# show service id 5 mfib
===============================================================================
Multicast FIB, Service 5
===============================================================================
Source Address  Group Address         Sap Id                   Svc Id   Fwd/Blk
```

```
-------------------------------------------------------------------------------
*               234.5.34.1            sap:lag-1:5.*            Local    Fwd
5.5.5.1         234.5.34.1            sap:lag-1:5.*            Local    Blk
                                      sap:lag-2:5.*            Local    Fwd
*               234.5.34.2            sap:lag-1:5.*            Local    Fwd
5.5.5.1         234.5.34.2            sap:lag-1:5.*            Local    Blk
                                      sap:lag-2:5.*            Local    Fwd
*               234.5.34.3            sap:lag-1:5.*            Local    Fwd
5.5.5.1         234.5.34.3            sap:lag-1:5.*            Local    Blk
                                      sap:lag-2:5.*            Local    Fwd
*               234.5.34.4            sap:lag-1:5.*            Local    Fwd
5.5.5.1         234.5.34.4            sap:lag-1:5.*            Local    Blk
                                      sap:lag-2:5.*            Local    Fwd
*               234.5.34.5            sap:lag-1:5.*            Local    Fwd
5.5.5.1         234.5.34.5            sap:lag-1:5.*            Local    Blk
                                      sap:lag-2:5.*            Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 10
===============================================================================
*A:7210-2
```

mrouters

Syntax **mrouters [detail]**

Context show>service>id>igmp-snooping

Description This command displays all multicast routers.

Parameters **detail —** Displays detailed information.

Sample Output

```
*A:7210-2 show service id 3 igmp-snooping mrouters
===============================================================================
IGMP Snooping Multicast Routers for service 3
===============================================================================
MRouter          Sap Id                Up Time        Expires   Version
-------------------------------------------------------------------------------
10.20.1.2        lag-2:3.*             0d 02:30:39    197s      3
-------------------------------------------------------------------------------
Number of mrouters: 1
===============================================================================
*A:SNPG-1#


*A:7210-2# show service id 3 igmp-snooping mrouters detail
===============================================================================
IGMP Snooping Multicast Routers for service 3
===============================================================================
-------------------------------------------------------------------------------
MRouter 10.20.1.2
-------------------------------------------------------------------------------
Sap Id                  : lag-2:3.*
Expires                 : 181s
Up Time                 : 0d 02:30:55
Version                 : 3
```

```
General Query Interval  : 125s
Query Response Interval : 10.0s
Robust Count            : 2
-------------------------------------------------------------------------------
Number of mrouters: 1
===============================================================================
*A:7210-2#
```

mvr

Syntax **mvr**

Context show>service>id>igmp-snooping

Description This command displays Multicast VPLS Registration (MVR) information.

port-db

Syntax **port-db sap** *sap-id* [**detail**]
port-db sap *sap-id* **group** *grp-address*

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 945 for command syntax.

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.

Output **Show Output —** The following table describes the show output fields:

Label	Description
Group Address	The IP multicast group address for which this entry contains information.

Label	Description
Mode	Specifies the type of membership report(s) received on the interface for the group. In the include mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter of the IGMP membership report. In exclude' mode, reception of packets sent to the given multicast address is requested from all IP source addresses except those listed in the source-list parameter.
Type	Indicates how this group entry was learned. If this group entry was learned by IGMP, the value is set to dynamic. For statically configured groups, the value is set to static.
Compatibility mode	Specifies the IGMP mode. This is used in order for routers to be compatible with older version routers. IGMPv3 hosts must operate in Version 1 and Version 2 compatibility modes. IGMPv3 hosts must keep state per local interface regarding the compatibility mode of each attached network. A host's compatibility mode is determined from the host compatibility mode variable which can be in one of three states: IGMPv1, IGMPv2 or IGMPv3. This variable is kept per interface and is dependent on the version of general queries heard on that interface as well as the older version querier present timers for the interface.
V1 host expires	The time remaining until the local router will assume that there are no longer any IGMP Version 1 members on the IP subnet attached to this interface. Upon hearing any IGMPv1 membership report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 leave messages for this group that it receives on this interface.
V2 host expires	The time remaining until the local router will assume that there are no longer any IGMP Version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 membership report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv3 leave messages for this group that it receives on this interface.
Source address	The source address for which this entry contains information.
Up Time	The time since the source group entry was created.
Expires	The amount of time remaining before this entry will be aged out.
Number of sources	Indicates the number of IGMP group and source specific queries received on this SAP.
Forwarding/Blocking	Indicates whether this entry is on the forward list or block list.
Number of groups	Indicates the number of groups configured for this SAP.

Sample Output

```
*A:7210-2>show>service>id>snooping# port-db sap 1/1/18:100
===============================================================================
IGMP Snooping SAP 1/1/18:100 Port-DB for service 900
===============================================================================
Group Address   Type    From-VPLS  Up Time        Expires

-------------------------------------------------------------------------------
226.6.6.6       dynamic local      0d 00:12:35    260s
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
*A:7210-2>show>service>id>snooping#
*A:7210-2# show service id 5 igmp-snooping port-db sap lag-1:5.*
===============================================================================
IGMP Snooping SAP lag-1:5.* Port-DB for service 5
===============================================================================
Group Address   Mode    Type    Up Time        Expires  Num
                                                        Src
-------------------------------------------------------------------------------
234.5.34.1      exclude dynamic 0d 00:23:38    216s     1
234.5.34.2      exclude dynamic 0d 00:23:38    216s     1
234.5.34.3      exclude dynamic 0d 00:23:38    216s     1
234.5.34.4      exclude dynamic 0d 00:23:38    216s     1
234.5.34.5      exclude dynamic 0d 00:23:38    216s     1
-------------------------------------------------------------------------------
Number of groups: 5
===============================================================================
*A:7210-2#

*A:7210-2>show>service>id>snooping# port-db sap 1/1/18:100 detail
===============================================================================
IGMP Snooping SAP 1/1/18:100 Port-DB for service 900
-------------------------------------------------------------------------------
IGMP Group 226.6.6.6
-------------------------------------------------------------------------------
Type             : dynamic
Up Time          : 0d 00:12:39          Expires          : 259s
Compat Mode      : IGMP Version 2
V1 Host Expires  : 0s                   V2 Host Expires  : 259s
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
*A:7210-2>show>service>id>snooping#
```

proxy-db

Syntax **proxy-db** [**detail**]
proxy-db group *grp-address*

Context show>service>id>igmp-snooping

Description This command displays information on the IGMP snooping proxy reporting database for the VPLS service.

Parameters **group** *grp-ip-address* — Displays the IGMP snooping proxy reporting database for a specific multicast group address.

Output **Show Output —** The following table describes the show output fields:

Label	Description
Group Address	The IP multicast group address for which this entry contains information.
Mode	Specifies the type of membership report(s) received on the interface for the group. In the include mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter of the IGMP membership report. In the "exclude" mode, reception of packets sent to the given multicast address is requested from all IP source addresses except those listed in the source-list parameter.
Up Time	The total operational time in seconds.
Num Sources	Indicates the number of IGMP group and source specific queries received on this interface.
Number of groups	Number of IGMP groups.
Source Address	The source address for which this entry contains information.

Sample Output

```
*A:7210-2 show service id 5 igmp-snooping proxy-db
===============================================================================
IGMP Snooping Proxy-reporting DB for service 5
===============================================================================
Group Address    Mode     Up Time        Num Sources
-------------------------------------------------------------------------------
234.5.34.1       exclude  0d 00:25:54    0
234.5.34.2       exclude  0d 00:25:54    0
234.5.34.3       exclude  0d 00:25:54    0
234.5.34.4       exclude  0d 00:25:54    0
234.5.34.5       exclude  0d 00:25:54    0
-------------------------------------------------------------------------------
Number of groups: 5
===============================================================================
*A:7210-2

*A:7210-2>show>service>id>snooping# proxy-db detail
===============================================================================
IGMP Snooping Proxy-reporting DB for service 900
-------------------------------------------------------------------------------
IGMP Group 226.6.6.6
-------------------------------------------------------------------------------
Up Time : 0d 00:15:17
------------------------------
Source Address  Up Time
```

```
------------------------------
No sources.
-------------------------------------------------------------------------------
IGMP Group 229.9.9.9
-------------------------------------------------------------------------------
Up Time : 0d 00:15:56
------------------------------
Source Address  Up Time
------------------------------
No sources.
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================
*A:7210-2>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:7210-2# show service id 3 igmp-snooping querier
===============================================================================
IGMP Snooping Querier info for service 3
===============================================================================
Sap Id                  : lag-2:3.*
IP Address              : 10.20.1.2
```

```
Expires                 : 226s
Up Time                 : 0d 02:30:11
Version                 : 3

General Query Interval  : 125s
Query Response Interval : 10.0s
Robust Count            : 2
===============================================================================
*A:7210-2#
```

static

Syntax **static** [**sap** *sap-id*]

Context show>service>id>igmp-snooping

Description This command displays information on static IGMP snooping source groups for the VPLS service.

Parameters **sap** *sap-id* — Displays static IGMP snooping source groups for a specific SAP. See Common CLI Command Descriptions on page 945 for command syntax.

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:7210-2# show service id 4093 igmp-snooping static
===============================================================================
IGMP Snooping Static Source Groups for service 4093
===============================================================================
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups for SAP 1/1/5:4093
-------------------------------------------------------------------------------
Source           Group
-------------------------------------------------------------------------------
93.93.93.1       239.93.39.1
93.93.93.1       239.93.39.2
93.93.93.1       239.93.39.3
93.93.93.1       239.93.39.4
93.93.93.1       239.93.39.5
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 5
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups for SAP lag-3:4093
-------------------------------------------------------------------------------
Source           Group
-------------------------------------------------------------------------------
93.93.93.1       239.93.39.1
93.93.93.1       239.93.39.2
93.93.93.1       239.93.39.3
93.93.93.1       239.93.39.4
93.93.93.1       239.93.39.5
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 5
===============================================================================
*A:7210-2#
```

statistics

Syntax **statistics** [**sap** *sap-id*]

Context show>service>id>igmp-snooping

Description This command displays IGMP snooping statistics for the VPLS service.

Parameters **sap** *sap-id* — Displays IGMP snooping statistics for a specific SAP. See Common CLI Command Descriptions on page 945 for command syntax.

```
*A:7210-2# show service id 5 igmp-snooping statistics
===============================================================================
IGMP Snooping Statistics for service 5
===============================================================================
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         0             228           0
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V3 Reports              282           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

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
Exceeded Max Num Sources : 0
===============================================================================
*A:7210-2#
```

VPLS Clear Commands

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears commands for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service.

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

fdb

Syntax **fdb** {**all** | **mac** *ieee-address* | **sap** *sap-id*]

Context clear>service>id

Description This command clears FDB entries for the service.

Parameters **all —** Clears all FDB entries.

mac *ieee-address* **—** Clears only FDB entries in the FDB table with the specified 48-bit MAC address. The MAC address can be expressed in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

sap-id — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 for command syntax.

sap

Syntax **sap** *sap-id* {**all** | **counters** | **stp**}

Context clear>service>statistics

Description This command clears statistics for the SAP bound to the service.

Parameters *sap-id* — See Common CLI Command Descriptions on page 945 for command syntax.

all — Clears all queue statistics and STP statistics associated with the SAP.

counters — Clears all queue statistics associated with the SAP.

counters

Syntax **counters**

Context clear>service>statistics>id

Description This command clears all traffic queue counters associated with the service ID.

stp

Syntax **stp**

Context clear>service>statistics>id

Description Clears all spanning tree statistics for the service ID.

detected-protocols

Syntax **detected-protocols** {**all** | **sap** *sap-id*}

Context clear>service>id>stp

Description RSTP automatically falls back to STP mode when it receives an STP BPDU. The **clear detected-protocols** command forces the system to revert to the default RSTP mode on the SAP.

Parameters **all** — Clears all detected protocol statistics.

sap-id — Clears the specified lease state SAP information. See Common CLI Command Descriptions on page 945 for command syntax.

igmp-snooping

Syntax **igmp-snooping**

Context clear>service>id

Description This command enables the context to clear IGMP snooping data.

port-db

Syntax **port-db** [**sap** *sap-id*] [**group** *grp-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 945 for command syntax.

group *grp-address* — Clears IGMP snooping statistics matching the specified group address.

querier

Syntax **querier**

Context clear>service>id>igmp-snooping

Description This command clears the information on the IGMP snooping queriers for the VPLS service.

statistics

Syntax **statistics** { **all** | **sap** *sap-id*}

Context clear>service>id>snooping

Description This command clears IGMP snooping statistics.

Parameters **all** — Clears all statistics for the service ID.

sap *sap-id* — Clears statistics for the specified SAP ID.

VPLS Debug Commands

id

Syntax **id** *service-id*

Context debug>service

Description This command debugs commands for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service.

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

event-type

Syntax [**no**] **event-type** {**config-change** | **svc-oper-status-change** | **sap-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.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id

Description This command enables debugging for a particular SAP.

Parameters *sap-id* — Specifies the SAP ID.

stp

Syntax **stp**

Context debug>service>id

Description This command enables the context for debugging STP.

all-events

Syntax **all-events**

Context debug>service>id>stp

Description This command enables STP debugging for all events.

bpdu

Syntax [**no**] **bpdu**

Context debug>service>id>stp

Description This command enables STP debugging for received and transmitted BPDUs.

exception

Syntax [**no**] **exception**

Context debug>service>id>stp

Description This command enables STP debugging for exceptions.

fsm-state-changes

Syntax [**no**] **fsm-state-changes**

Context debug>service>id>stp

Description This command enables STP debugging for FSM state changes.

fsm-timers

Syntax [**no**] **fsm-timers**

Context debug>service>id>stp

Description This command enables STP debugging for FSM timer changes.

port-role

Syntax [**no**] **port-role**

Context debug>service>id>stp

Description This command enables STP debugging for changes in port roles.

port-state

Syntax [**no**] **port-state**

Context debug>service>id>stp

Description This command enables STP debugging for port states.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>stp

Description This command enables STP debugging for a specific SAP.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 for command syntax.

IES Show Commands

customer

Syntax **customer** [*customer-id*] [**site** *customer-site-name*]

Context show>service

Description This command displays service customer information.

Parameters *customer-id* — Displays only information for the specified customer ID.

Default All customer IDs display

Values 1 — 2147483647

site *customer-site-name* — Specifies the customer site which is an anchor point for an ingress and egress virtual scheduler hierarchy.

Output **Show Customer Command Output —** The following table describes show customer command output fields:

Label	Description
Customer-ID	The ID that uniquely identifies a customer.
Contact	The name of the primary contact person.
Description	Generic information about the customer.
Phone	The phone/pager number to reach the primary contact person.
Total Customers	The total number of customers configured.
Multi-service site	
Site	Multi-service site name. A multi-service customer site is a group of SAPs with common origination and termination points.
Description	Information about a specific customer's multi-service site.
Assignment	The port ID, MDA, or card number, where the SAP's that are members of this multi- service site are defined.
I. Sched Pol	The ingress QoS scheduler policy assigned to this multi-service site.
E. Sched Pol	The egress QoS scheduler policy assigned to this multi-service site.
Service Association	
Service-ID	The ID that uniquely identifies a service.
SAP	Specifies the SAP assigned to the service.

Sample Output

```
*A:ALA-12# show service customer
===============================================================
Customers
===============================================================
Customer-ID : 1
Contact     : Manager
Description : Default customer
Phone       : (123) 555-1212

Customer-ID : 2
Contact     : Tech Support
Description : TiMetra Networks
Phone       : (234) 555-1212

Customer-ID : 3
Contact     : Fred
Description : TiMetra Networks
Phone       : (345) 555-1212

Customer-ID : 6
Contact     : Ethel
Description : Epipe Customer
Phone       : (456) 555-1212

Customer-ID : 7
Contact     : Lucy
Description : ABC Customer
Phone       : (567) 555-1212

Customer-ID : 8
Contact     : Customer Service
Description : IES Customer
Phone       : (678) 555-1212

Customer-ID : 274
Contact     : Mssrs. Beaucoup
Description : ABC Company
Phone       : 650 123-4567

Customer-ID : 94043
Contact     : Test Engineer on Duty
Description : TEST Customer
Phone       : (789) 555-1212

---------------------------------------------------------
Total Customers : 8
---------------------------------------------------------------
*A:ALA-12#


*A:ALA-12# show service customer 274
===============================================================================
Customer  274
===============================================================================
Customer-ID : 274
Contact     : Mssrs. Beaucoup
Description : ABC Company
Phone       : 650 123-4567
-------------------------------------------------------------------------------
```

```
Multi Service Site
-------------------------------------------------------------------------------
Site        : west
Description : (Not Specified)
===============================================================================
*A:ALA-12#


*A:ALA-12# show service customer 274 site west
===============================================================================
Customer  274
===============================================================================
Customer-ID : 274
Contact     : Mssrs. Beaucoup
Description : ABC Company
Phone       : 650 123-4567
-------------------------------------------------------------------------------
Multi Service Site
-------------------------------------------------------------------------------
Site        : west
Description : (Not Specified)
Assignment  : Card 5
I. Sched Pol: SLA1
E. Sched Pol: (Not Specified)
-------------------------------------------------------------------------------
Service Association
-------------------------------------------------------------------------------
No Service Association Found.
===============================================================================
*A:ALA-12#
```

sap-using

Syntax **sap-using** [**sap** *sap-id*]
sap-using interface [*ip-address* | *ip-int-name*]
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**ingress**] **qos-policy** *qos-policy-id*

Context show>service

Description 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 **sap *sap-id* —** Specifies the physical port identifier portion of the SAP definition. See Common CLI Command Descriptions on page 945 for command syntax.

ingress — Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

filter *filter-id* **—** The ingress or egress filter policy ID for which to display matching SAPs.

Values 1 — 65535

interface — Specifies matching SAPs with the specified IP interface.

ip-addr — 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.

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 value that identifies the service.
SapMTU	The SAP MTU value.
Igr.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.
E.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.
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The administrative state of the SAP.
Opr	The actual state of the SAP.

Sample Output

```
*A:DUT-B# show service sap-using sap 1/1/3:100.*
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.    Adm  Opr
                                           QoS   Fltr    Fltr
-------------------------------------------------------------------------------
1/1/1                           6          1     none    none    Up   Down
1/1/2                           700        1     none    none    Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 2
===============================================================================
*A:DUT-B#
```

service-using

Syntax **service-using** [**ies**] [**customer** *customer-id*]

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 **ies** — Displays matching IES services.

customer *customer-id* — Displays services only associated with the specified customer ID.

Default Services associated with an customer.

Values 1 — 2147483647

Output **Show Service Service-Using —** The following table describes show service service-using output fields:

Label	Description
Service Id	The value that identifies the service.
Type	Specifies the service type configured for the service ID.
Adm	The administrative state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

Sample Output

```
A:ALA-48# show service service-using ies
===============================================================================
Services [ies]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
88           IES       Up     Down       8                 07/25/2006 15:46:28
89           IES       Up     Down       8                 07/25/2006 15:46:28
104          IES       Up     Down       1                 07/25/2006 15:46:28
200          IES       Up     Down       1                 07/25/2006 15:46:28
214          IES       Up     Down       1                 07/25/2006 15:46:28
321          IES       Up     Down       1                 07/25/2006 15:46:28
322          IES       Down   Down       1                 07/25/2006 15:46:28
1001         IES       Up     Down       1730              07/25/2006 15:46:28
-------------------------------------------------------------------------------
Matching Services : 8
-------------------------------------------------------------------------------
A:ALA-48#
```

id

Syntax **id** *service-id* {**all** | **arp** | **base** | **sap** }

Context show>service

Description This command displays information for a particular service-id.

Parameters *service-id* — The unique service identification number to identify the service in the service domain.

all — Display detailed information about the service.

arp — Display ARP entries for the service.

base — Display basic service information.

sap — Display SAPs associated to the service.

all

Syntax **all**

Context show>service>id

Description This command displays detailed information for all aspects of the service.

Output **Show All Service-ID Output —** The following table describes the show all service-id command output fields:

Label	Description
Service Detailed Information	
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.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.

Label	Description (Continued)
Admin Path MTU	The 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: 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 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 MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the service.
Oper State	The current status of the service.
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.
Number of SDPs	The total number SDPs applied to this service ID.
Service Access Points	

Label	Description (Continued)
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	The value of the label used to identify this SAP on the access port.
Admin State	The desired 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 largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The SAP ingress QoS policy ID.
Egress qos-policy	The SAP egress QoS policy ID.
Ingress Filter-Id	The SAP ingress filter policy ID.
Egress Filter-Id	The SAP egress filter policy ID.
Multi Svc Site	Indicates the multi-service site that the SAP is a member.
Ingress sched-policy	Indicates the ingress QoS scheduler for the SAP.
Egress sched-policy	Indicates the egress QoS scheduler for the SAP.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
SAP Statistics	
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.
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.

arp

Syntax **arp** [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*]

Context show>service>id

Description Displays the ARP table for the IES instance. The ARP entries for a subscriber interface are displayed uniquely. Each MAC associated with the subscriber interface child group-interfaces are displayed with each subscriber interface ARP entry. They do not reflect actual ARP entries but are displayed along the interfaces ARP entry for easy lookup.

Parameters *ip-address* — Displays only ARP entries in the ARP table with the specified IP address.

Default All IP addresses.

mac *ieee-address* **—** Displays only ARP entries in the ARP table with the specified 48-bit MAC address. The MAC address can be expressed in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

Default All MAC addresses.

sap *sap-id* **—** Displays SAP information for the specified SAP ID. See Common CLI Command Descriptions on page 945 for command syntax.

port-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.

Sample Output

```
*A:DUT-B# show service id 100 arp
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address         Type       Expiry    Interface         SAP
```

```
-------------------------------------------------------------------------------
192.168.1.2     00:00:01:00:00:01  Other     00h00m00s   HW                1/1/1:10*
195.168.1.1     32:67:01:01:00:03  Other     00h00m00s   to7x              1/1/3:10*
195.168.1.2     32:68:01:01:00:02  Dynamic   03h59m58s   to7x              1/1/3:10*
===============================================================================
*A:DUT-B#
```

base

Syntax **base**

Context show>service>id

Description This command displays basic information about this IES service.

Sample Output

```
*A:ALA-A# show service id 100 base
-------------------------------------------------------------
Service Basic Information
-------------------------------------------------------------
Service Id        : 100                Vpn Id            : 100
Service Type      : IES
Description       : Default Ies description for service id 100
Customer Id       : 1
Last Status Change: 08/29/2006 17:44:28
Last Mgmt Change  : 08/29/2006 17:44:28
Admin State       : Up                 Oper State        : Up
SAP Count         : 2
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                      Type         AdmMTU  OprMTU  Adm     Opr
-------------------------------------------------------------------------------
sap:1/1/3                       null         1514    1514    Up      Up
sap:1/1/4                       null         1514    1514    Up      Up
===============================================================================
*A:ALA-A#
```

interface

Syntax **interface** [*ip-address* | *ip-int-name*] [**detail | summary**]

Context show>service>id

Description This command displays information for the IP interfaces associated with the IES service. If no optional parameters are specified, a summary of all IP interfaces associated to the service are displayed.

Parameters *ip-address* — The IP address of the interface for which to display information.

Values ipv4-address: a.b.c.d (host bits must be 0)

ip-int-name — Specifies the IP interface name for which to display information.

Values 32 characters maximum

detail — Displays detailed IP interface information.

Default IP interface summary output.

summary — Displays the summary of IP interface information.

Output **Show Service-ID —** The following table describes show service-id output fields.

Label	Description
If Name	The name used to refer to the IES interface.
Type	Specifies the interface type.
IP-Address	Specifies the IP address/IP subnet/broadcast address of the interface.
Adm	The administrative state of the interface.
Opr	The operational state of the interface.
Admin State	The administrative state of the interface.
Oper State	The operational state of the interface.
IP Addr/mask	Specifies the IP address/IP subnet/broadcast address of the interface.
If Index	The index corresponding to this IES interface. The primary index is 1; all IES interfaces are defined in the base virtual router context.
If Type	Specifies the interface type.
SAP Id	Specifies the SAP's port ID.
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled.
Arp Timeout	Specifies the timeout for an ARP entry learned on the interface.

Label	Description (Continued)
MAC Address	Specifies the 48-bit IEEE 802.3 MAC address.
ICMP Mask Reply	Specifies whether ICMP mask reply is enabled or disabled.
Cflowd	Specifies whether Cflowd collection and analysis on the interface is enabled or disabled.
Redirects	Specifies the rate for ICMP redirect messages.
Unreachables	Specifies the rate for ICMP unreachable messages.
TTL Expired	Specifies the rate for ICMP TTL messages.

Sample Output

```
A:ALA-49# show service id 88 interface
===============================================================================
Interface Table
===============================================================================
Interface-Name                   Adm          Opr          Type    Port/SapId
   IP-Address
-------------------------------------------------------------------------------
Sector A                         Up           Down/Down    IES     1/1/1.2.2
   -                                                                -
test                             Up           Down/Down    IES     1/1/2:0
   1.1.1.1/31                                                       n/a
   1.1.1.1/31                                                       n/a
   1.1.2.1/31                                                       n/a
test27                           Up           Up/--        IES Sub subscriber
   192.168.10.21/24                                                 n/a
grp-if                           Up           Down/--      IES Grp 1/2/2
Interfaces : 4
===============================================================================
A:ALA-49#
A:ALA-49# show service id 88 interface
===============================================================================
Interface Table
===============================================================================
Interface-Name Adm Opr(v4/v6) Type Port/SapId
IP-Address PfxState
-------------------------------------------------------------------------------
Sector A Up Down/Down IES 1/1/1.2.2
- -
test Up Down/Down IES 1/1/2:0
1.1.1.1/31 n/a
1.1.1.1/31 n/a
1.1.2.1/31 n/a
test27 Up Up/-- IES Sub subscriber
192.168.10.21/24 n/a
grp-if Up Down/-- IES Grp 1/2/2
Interfaces : 4
===============================================================================
A:ALA-49#
```

Appendix: Port-Based Split Horizon

In This Chapter

This section provides Port-Based Split Horizon configuration information.

- Overview on page 932
- Configuration Guidelines on page 934

Overview

NOTE: Port-based HSG is supported on 7210 SAS-E and 7210 SAS-D. Port-based SHG is not supported on 7210 SAS-K 2F2T1C and 7210 SAS-K 2F4T6C.

The port-based split horizon feature can be used to disable local switching on the 7210 SAS. A loop-free topology can be achieved using split horizon on 7210 SAS switches.

Traffic arriving on an access or an access-uplink port within a split horizon group will not be copied to other access and an access-uplink ports in the same split horizon group, but will be copied to an access-uplink ports in other split horizon groups.

Since split horizon is a per port feature in 7210 SAS, all SAPs associated with the port becomes part of split horizon group configured on that port.

Topology

Figure illustrates an example of split horizon groups used to prevent communication between two access SAPs and between two access-uplink SAPs.

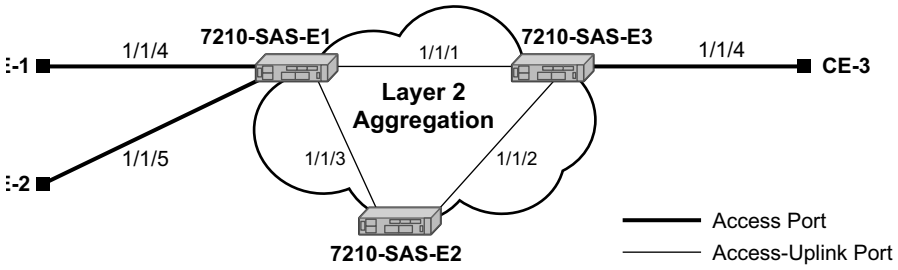

Figure 76: Split Horizon Group Example

Using 7210-SAS-1 as an example:

1. Split horizon group “access” is created to prevent any communication between the SAP’s part of port 1/1/4 and port 1/1/5 (configured as access port) within the same VPLS.
2. Split horizon group “access uplink” is created to prevent any communication between SAP’s part of port 1/1/1 and port 1/1/3 (configured as an access-uplink) within the same VPLS.
3. VPLS 100 is created on 7210 SAS-1 with SAPs 1/1/1, 1/1/3, and SAPs on 1/1/4 and 1/1/5 as part of this VPLS. CE1, CE2 and CE3 are the customer sites.
4. With this configuration, any communication between ports 1/1/4 and 1/1/5 gets blocked, similarly communication between ports 1/1/1 and 1/1/3 gets blocked but any traffic received on ports (SAPs) part of split horizon group “access uplink” will

be switched to ports (SAPs) part of split horizon group “access” and vice versa based on the FDB entries for VPLS 100.

Configuration Guidelines

The following configuration guidelines must be followed to configure a split horizon group.

1. Create a split horizon group in the config prompt. The group name must be unique across the system.

```
7210-SAS1>config#info
#------------------------------------------------
echo "Split-horizon-group Configuration"
#------------------------------------------------
    split-horizon-group access create
        description "Block access between access Ports"
    split-horizon-group access-uplink create
        description "Block access between access-uplink Ports"
    exit
#------------------------------------------------
7210-SAS1>config#
```

2. Associate ports 1/1/4 and 1/1/5 with split horizon group "access" By default all ports are access ports. The default Ethernet encapsulation for access port is "null".

```
7210-SAS1>config#info
#------------------------------------------------
echo "Port Configuration"
#------------------------------------------------
    port 1/1/4
        split-horizon-group access
        ethernet
        exit
        no shutdown
exit
    port 1/1/5
        split-horizon-group access
        ethernet
        exit
        no shutdown
  exit
#------------------------------------------------
7210-SAS1>config#
```

3. Configure ports 1/1/1 and 1/1/3 as access uplink and associate these ports with split horizon group "access-uplink" default Ethernet encapsulation for access uplink port is "qinq".

```
7210-SAS1>config# info
#------------------------------------------------
echo "Port Configuration"
#------------------------------------------------
    port 1/1/1
```

```
        split-horizon-group access-uplink
        ethernet
               mode access uplink
        exit
        no shutdown
exit
    port 1/1/3
        split-horizon-group access-uplink
        ethernet
               mode access uplink
        exit
        no shutdown
  exit
#--------------------------------------------------
7210-SAS1>config#
```

4. Create a VPLS instance 100.

```
#--------------------------------------------------
echo "Service Configuration"
#--------------------------------------------------
    service
          customer 2 create
          exit
         vpls 100 customer 2 create
            stp
                shutdown
            exit
            sap 1/1/1:100.* create
            exit
            sap 1/1/3:100.* create
            exit
            sap 1/1/4 create
            exit
            sap 1/1/5 create
            exit
            no shutdown
    exit
...
#--------------------------------------------------
```

Note: A split horizon on a port must be configured before creating any SAPs associated with that port.

Verification

The following output verifies the split horizon configuration on a 7210 SAS:

```
7210-SAS1# show split-horizon-group
===============================================================================
Port: Split Horizon Group
===============================================================================
Name                                Description
-------------------------------------------------------------------------------
access                              Block access between access Ports
access-uplink                       Block access between access-uplink Ports

No. of Split Horizon Groups: 2
===============================================================================
7210-SAS1#
```

Execute the below mentioned command to verify the port association with split horizon groups:

```
7210-SAS1# show split-horizon-group access
===============================================================================
Port: Split Horizon Group
===============================================================================
Name                                Description
-------------------------------------------------------------------------------
access                              Block access between access Ports
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Port1/1/4                           10/100/Gig Ethernet SFP
Port1/1/5                           10/100/Gig Ethernet SFP

Ports Associated : 2
===============================================================================
7210-SAS1#


7210-SAS1# show split-horizon-group access-uplink
===============================================================================
Port: Split Horizon Group
===============================================================================
Name                                Description
-------------------------------------------------------------------------------
Access-uplink                  Block access between access-uplink Ports
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
Port1/1/1                           10/100/Gig Ethernet SFP
Port1/1/3                           10/100/Gig Ethernet SFP

Ports Associated : 2
===============================================================================
7210-SAS1#
```

Appendix: DHCP Management

In This Chapter

This chapter provides information about using DHCP, including theory, supported features and configuration process overview.

The topics in this chapter include:

- DHCP Principles on page 938
- DHCP Features on page 940
- Common Configuration Guidelines on page 943

DHCP Principles

NOTE: The reference to spoke-SDP or mesh-SDP in this section is applicable to only 7210 SAS-K 2F4T6C. It is not applicable to 7210 SAS-D, E, 7210 SAS-K 2F2T1C.

In a Triple Play network, client devices (such as a routed home gateway, a session initiation protocol (SIP) phone or a set-top box) use Dynamic Host Configuration Protocol (DHCP) to dynamically obtain their IP address and other network configuration information. 7210 autoinit procedure also uses DHCP to dynamically obtain the BOF file used for first-time booting of the system (along with IP address required to retrieve the BOF file, the configuration file and the Timos software image from the network). DHCP is defined and shaped by several RFCs and drafts in the IETF DHC working group including the following

- RFC 2131, Dynamic Host Configuration Protocol
- RFC 3046, DHCP Relay Agent Information Option

The DHCP operation is illustrated in Figure 77.

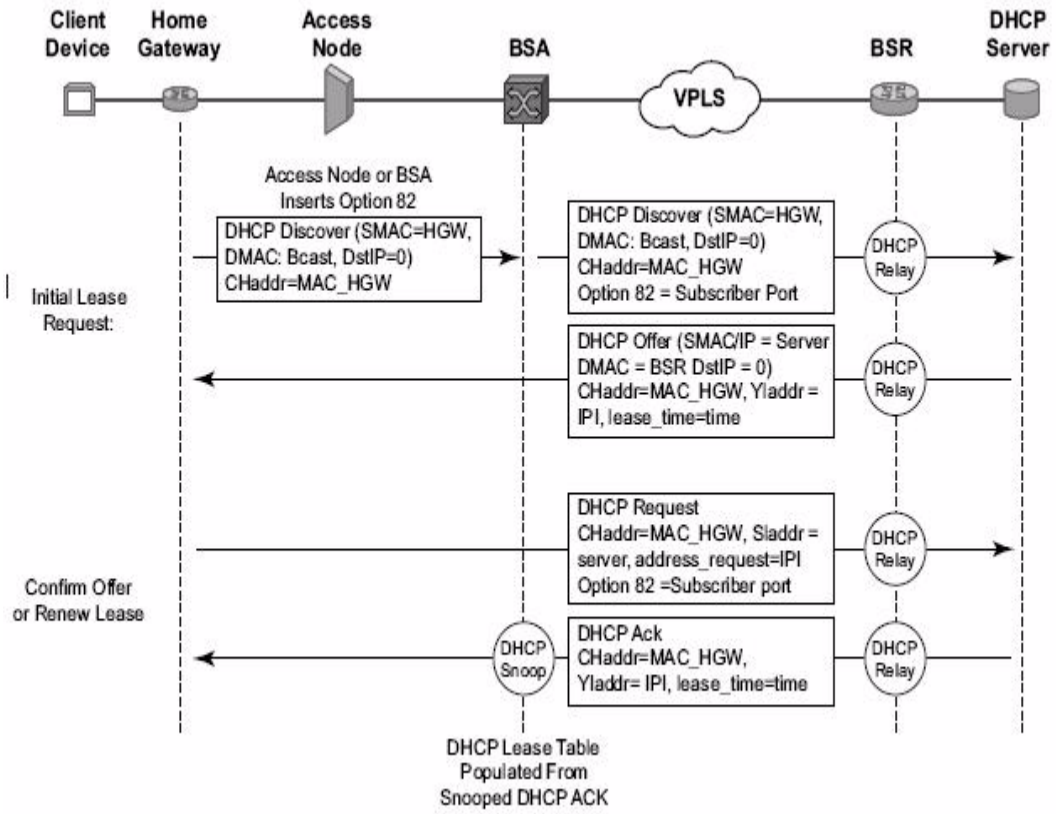

Figure 77: IP Address Assignment with DHCP

1. During boot-up, the client device sends a DHCP discover message to get an IP address from the DHCP Server. The message contains:
   - → Destination MAC address — broadcast
   - → Source MAC address — MAC of client device
   - → Client hardware address — MAC of client device

If this message passes through a DSLAM or other access node (possibly a 7210 SAS device), typically the Relay information option (Option 82) field is added, indicating shelf, slot, port, VPI, VCI and other fields, to identify the subscriber.

DHCP relay is enabled on the first IP interface in the upstream direction. Depending on the scenario, the DSLAM, BSA or the BSR will relay the discover message as a unicast packet towards the configured DHCP server. DHCP relay is configured to insert the giaddr in order to indicate to the DHCP server in which subnet an address should be allocated.

2. The DHCP server will lookup the client MAC address and Option 82 information in its database. If the client is recognized and authorized to access the network, an IP address will be assigned and a DHCP offer message returned. The BSA or BSR will relay this back to the client device.
3. It is possible that the discover reached more than one DHCP server, and thus that more than one offer was returned. The client selects one of the offered IP addresses and confirms it wants to use this in a DHCP request message, sent as unicast to the DHCP server that offered it.
4. The DHCP server confirms that the IP address is still available, updates its database to indicate it is now in use, and replies with a DHCP ACK message back to the client. The ACK also contains the Lease Time of the IP address.

DHCP Features

- Using Option 82 Field on page 940
- Trusted and Untrusted on page 941
- DHCP Snooping on page 941

Using Option 82 Field

Option 82, or the relay information option is specified in RFC 3046, DHCP Relay Agent Information Option, allows the router to append some information to the DHCP request that identifies where the original DHCP request arrives from.

There are two sub-options under Option 82:

- Agent Circuit ID Sub-option (RFC 3046, section 3.1): This sub-option specifies data which must be unique to the box that is relaying the circuit.
- Remote ID Sub-option (RFC 3046 section 3.2): This sub-option identifies the host at the other end of the circuit. This value must be globally unique.

Both sub-options are supported by the Alcatel-Lucent 7210 SAS and can be used separately or together.

Inserting Option 82 information is supported independently of DHCP relay.

When the circuit id sub-option field is inserted by the 7210 SAS, it can take following values:

- *sap-id*: The SAP index (only under a IES service)
- *ifindex*: The index of the IP interface (only under a IES service)
- *ascii-tuple*: An ASCII-encoded concatenated tuple, consisting of [system-name|serviceid| interface-name] (for IES) or [system-name|service-id|sap-id] (for VPLS).
- *vlan-ascii-tuple*: An ASCII-encoded concatenated tuple, consisting of the ascii-tuple followed by Dot1p bits and Dot1q tags.

When a DHCP packet is received with Option 82 information already present, the system can do one of three things. The available actions are:

- *Replace* — On ingress the existing information-option is replaced with the information-option parameter configured on the 7210 SAS. On egress (towards the customer) the information-option is stripped (per the RFC).
- *Drop* — The DHCP packet is dropped and a counter is incremented.

- *Keep* — The existing information is kept on the packet and the router does not add any additional information. On egress the information option is not stripped and is sent on to the downstream node.

In accordance with the RFC, the default behavior is to keep the existing information; except if the giaddr of the packet received is identical to a local IP address on the router, then the packet is dropped and an error incremented regardless of the configured action.

The maximum packet size for a DHCP relay packet is 1500 bytes. If adding the Option 82 information would cause the packet to exceed this size, the DHCP relay request will be forwarded without the Option 82 information. This packet size limitation exists to ensure that there will be no fragmentation on the end Ethernet segment where the DHCP server attaches.

In the downstream direction, the inserted Option 82 information should not be passed back towards the client (as per RFC 3046, DHCP Relay Agent Information Option). To enable downstream stripping of the option 82 field, DHCP snooping should be enabled on the SDP or SAP connected to the DHCP server.

Trusted and Untrusted

There is a case where the relay agent could receive a request where the downstream node added Option 82 information without also adding a giaddr (giaddr of 0). In this case the default behavior is for the router to drop the DHCP request. This behavior is in line with the RFC.

The 7210 SAS supports a command trusted, which allows the router to forward the DHCP request even if it receives one with a giaddr of 0 and Option 82 information attached. This could occur with older access equipment. In this case the relay agent would modify the request's giaddr to be equal to the ingress interface. This only makes sense when the action in the information option is keep, and the service is IES . In the case where the Option 82 information gets replaced by the relay agent, either through explicit configuration or the VPLS DHCP Relay case, the original Option 82 information is lost, and the reason for enabling the trusted option is lost.

DHCP Snooping

To support DHCP based address assignment in L2 aggregation network, 7210 supports DHCP snooping. 7210 can copy packets designated to the standard UDP port for DHCP (port 67) to its control plane for inspection, this process is called DHCP snooping.

DHCP snooping can be performed in two directions:

1. From the client to the DHCP server (Discover or Request messages) to insert Option 82 information; For these applications, DHCP snooping must be enabled on the SAP towards the subscriber.
2. From the DHCP server (ACK messages), to remove the Option 82 field towards the client. For these applications, DHCP snooping must be enabled on both the SAP towards the network and the SAP towards the subscriber.

Common Configuration Guidelines

The topic in this section are:

- Configuration Guidelines for DHCP relay and snooping on page 943
- Configuring Option 82 Handling on page 943

Configuration Guidelines for DHCP relay and snooping

The following configuration guidelines must be followed to configure DHCP relay and snooping.

- 7210 SAS devices does not support the ARP populate based on the DHCP lease, assigned to the DHCP client
- 7210 SAS devices does not maintain the DHCP lease assigned to the client
- 7210 SAS devices do not perform IP spoofing checks and MAC spoofing checks based on the DHCP parameters assigned to the client
- MAC learning must be enabled in the VPLS service, for DHCP snooping.
- DHCP snooping is not supported for B-SAPs in B-VPLS services and I-SAPs in I-VPLS services.
- Ingress ACLs cannot be used to drop DHCP control packet.
- On 7210 SAS-K 2F4T6C DHCP packets received over a SDP are identified and option-82 inserted by the node can be removed by the node, in the downstream direction. SAP or a SDP, as applicable. DHCP snooping configuration on a SDP is supported only on 7210 SAS-K 2F4T6C.

Configuring Option 82 Handling

Option 82, or “Relay Information Option” is a field in DHCP messages used to identify the subscriber. The Option 82 field can already be filled in when a DHCP message is received at the router, or it can be empty. If the field is empty, the router shoulMAC learning must be enabled in the VPLS service, for DHCP snooping.d add identifying information (circuit ID, remote ID or both). If the field is not empty, the router can decide to replace it.

The following example displays an example of a partial BSA configuration with Option 82 adding on a VPLS service. Note that snooping must be enabled explicitly on a SAP or a SDP, as applicable. DHCP snooping configuration on a SDP is supported only on 7210 SAS-K 2F4T6C.

```
*A:7210SAS>config>service#
------------------------------------------------
```

```
vpls 2 customer 1 create
            shutdown
            stp
                shutdown
            exit
sap 1/1/12:100 create
                dhcp                    //Configuration example to add option 82
                    option
                        action replace
                        circuit-id
                        no remote-id
                    exit
                    no shutdown
                exit
            exit
            no shutdown
        exit
----------------------------------------------

*A:7210SAS>config>service#
```

The following example displays an example of a partial BSA configuration to remove the Option 82 on a VPLS service.

```
vpls 2 customer 1 create
            stp
                shutdown
            exit
            sap 1/1/14:100 create      //Configuration example to remove option 82
                dhcp
                    snoop
                    no shutdown
                exit
            exit
```

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 946

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	*slot*/*mda*/*port*[.*channel*]	1/1/5
null	[*port-id* | *lag-id*]	*port-id*: 1/1/3 *lag-id*: lag-3
dot1q	[*port-id* | *lag-id*]:qtag1	*port-id*:qtag1: 1/1/3:100 *lag-id*:qtag1:lag-3:102 *cp*.conn-prof-id: 1/2/1:cp.2
qinq	[*port-id* | *lag-id*]:*qtag1.qtag2*	*port-id*:qtag1.qtag2: 1/1/3:100.10 *lag-id*:qtag1.qtag2: lag-10:

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.

Standards and Protocol Support

Note: The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

M(A,N) means 7210 SAS-M in both Access-uplink mode and Network mode; Similarly M(N) means 7210 SAS-M in network mode only

T(A,N) means 7210 SAS-M in both Access-uplink mode and Network mode; Similarly T(N) means 7210 SAS-T in network mode only

K5 means 7210 SAS-K 2F2T1C

K12 means 7210 SAS-K 2F4T6C

Sx/S-1/10GE means all variants of 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE platforms

Sx-1/10GE means only the variants of 7210 SAS-Sx 1/10G

R6 means 7210 SAS-R6

R12 means 7210 SAS-R12

D means 7210 SAS-D and 7210 SAS-D ETR; if a line item applies to 7210 SAS-D ETR, then it is indicated as D-ETR

E means 7210 SAS-E

X means 7210 SAS-X

BGP

draft-ietf-idr-add-paths-04, Advertisement of Multiple Paths in BGP is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-sidr-origin-validation-signaling-04, BGP Prefix Origin Validation State Extended Community is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1772, Application of the Border Gateway Protocol in the Internet is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1997, BGP Communities Attribute is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2385, Protection of BGP Sessions via the TCP MD5 Signature Option is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2439, BGP Route Flap Damping is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2545, Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2858, Multiprotocol Extensions for BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2918, Route Refresh Capability for BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx-10/100GE, R6, and R12

RFC 3107, Carrying Label Information in BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3392, Capabilities Advertisement with BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4271, A Border Gateway Protocol 4 (BGP-4) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4360, BGP Extended Communities Attribute is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4456, BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4659, BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4724, Graceful Restart Mechanism for BGP (Helper Mode) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4760, Multiprotocol Extensions for BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4798, Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4893, BGP Support for Four-octet AS Number Space is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5004, Avoid BGP Best Path Transitions from One External to Another is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5291, Outbound Route Filtering Capability for BGP-4 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5668, 4-Octet AS Specific BGP Extended Community is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6811, Prefix Origin Validation is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

Circuit Emulation

RFC 4553, Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP) is supported on M(N)

RFC 5086, Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN) is supported on M(N)

RFC 5287, Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks is supported on M(N)

Ethernet

IEEE 802.1AB, Station and Media Access Control Connectivity Discovery is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1ad, Provider Bridges is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, R6, and R12

IEEE 802.1ag, Connectivity Fault Management is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1ah, Provider Backbone Bridges is supported on M(N), X, and T(N)

IEEE 802.1ax, Link Aggregation is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1D, MAC Bridges is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1p, Traffic Class Expediting is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1Q, Virtual LANs is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1s, Multiple Spanning Trees is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1w, Rapid Reconfiguration of Spanning Tree is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.1X, Port Based Network Access Control is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.3ab, 1000BASE-T is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.3ac, VLAN Tag is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.3ad, Link Aggregation is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.3ae, 10 Gb/s Ethernet is supported on M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.3ah, Ethernet in the First Mile is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.3ba, 40 Gb/s and 100 Gb/s Ethernet is supported on R6 and R12

IEEE 802.3i, Ethernet is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.3u, Fast Ethernet is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 802.3z, Gigabit Ethernet is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T G.8032, Ethernet Ring Protection Switching is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T Y.1731, OAM functions and mechanisms for Ethernet based networks is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, Operational management of Loop Free Alternates is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5286, Basic Specification for IP Fast Reroute: Loop-Free Alternates is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IP — General

draft-grant-tacacs-02, The TACACS+ Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 768, User Datagram Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 793, Transmission Control Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 854, TELNET Protocol Specifications is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 951, Bootstrap Protocol (BOOTP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1034, Domain Names - Concepts and Facilities is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1035, Domain Names - Implementation and Specification is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1350, The TFTP Protocol (revision 2) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1534, Interoperation between DHCP and BOOTP is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1542, Clarifications and Extensions for the Bootstrap Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2131, Dynamic Host Configuration Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2347, TFTP Option Extension is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2348, TFTP Blocksize Option is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2349, TFTP Timeout Interval and Transfer Size Options is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2428, FTP Extensions for IPv6 and NATs is supported on D, E, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2865, Remote Authentication Dial In User Service (RADIUS) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2866, RADIUS Accounting is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3046, DHCP Relay Agent Information Option (Option 82) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3596, DNS Extensions to Support IP version 6 is supported on D, E, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3768, Virtual Router Redundancy Protocol (VRRP) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4250, The Secure Shell (SSH) Protocol Assigned Numbers is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4251, The Secure Shell (SSH) Protocol Architecture is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4254, The Secure Shell (SSH) Connection Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4632, Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5880, Bidirectional Forwarding Detection (BFD) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5881, Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6528, Defending against Sequence Number Attacks is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IP — Multicast

RFC 1112, Host Extensions for IP Multicasting is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2236, Internet Group Management Protocol, Version 2 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3306, Unicast-Prefix-based IPv6 Multicast Addresses is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3376, Internet Group Management Protocol, Version 3 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3446, Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4601, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4604, Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4607, Source-Specific Multicast for IP is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4608, Source-Specific Protocol Independent Multicast in 232/8 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4610, Anycast-RP Using Protocol Independent Multicast (PIM) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5059, Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5384, The Protocol Independent Multicast (PIM) Join Attribute Format is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6513, Multicast in MPLS/BGP IP VPNs is supported on T(N), Mxp, R6, and R12

RFC 6514, BGP Encodings and Procedures for Multicast in MPLS/IP VPNs is supported on T(N), Mxp, R6, and R12

RFC 6515, IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs is supported on T(N), Mxp, R6, and R12

RFC 6625, Wildcards in Multicast VPN Auto-Discover Routes is supported on T(N), Mxp, R6, and R12

RFC 6826, Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path is supported on T(N), Mxp, R6, and R12

RFC 7385, IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points is supported on T(N), Mxp, R6, and R12

IP — Version 4

RFC 791, Internet Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 792, Internet Control Message Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 826, An Ethernet Address Resolution Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1519, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1812, Requirements for IPv4 Routers is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1981, Path MTU Discovery for IP version 6 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2401, Security Architecture for Internet Protocol is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2460, Internet Protocol, Version 6 (IPv6) Specification is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IP — Version 6

RFC 2464, Transmission of IPv6 Packets over Ethernet Networks is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3021, Using 31-Bit Prefixes on IPv4 Point-to-Point Links is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3122, Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3587, IPv6 Global Unicast Address Format is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4007, IPv6 Scoped Address Architecture is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4193, Unique Local IPv6 Unicast Addresses is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4291, Internet Protocol Version 6 (IPv6) Addressing Architecture is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4443, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4861, Neighbor Discovery for IP version 6 (IPv6) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4862, IPv6 Stateless Address Autoconfiguration (Router Only) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5095, Deprecation of Type 0 Routing Headers in IPv6 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5952, A Recommendation for IPv6 Address Text Representation is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6106, IPv6 Router Advertisement Options for DNS Configuration is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6164, Using 127-Bit IPv6 Prefixes on Inter-Router Links is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IPsec

RFC 2401, Security Architecture for the Internet Protocol is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2406, IP Encapsulating Security Payload (ESP) is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IS-IS

draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-kaplan-isis-ext-eth-02, Extended Ethernet Frame Size Support is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ISO/IEC 10589:2002, Second Edition, Nov. 2002, Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3359, Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3719, Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3787, Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4971, Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5120, M-ISIS: Multi Topology (MT) Routing in IS-IS is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5130, A Policy Control Mechanism in IS-IS Using Administrative Tags is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5301, Dynamic Hostname Exchange Mechanism for IS-IS is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5302, Domain-wide Prefix Distribution with Two-Level IS-IS is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5303, Three-Way Handshake for IS-IS Point-to-Point Adjacencies is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5304, IS-IS Cryptographic Authentication is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5305, IS-IS Extensions for Traffic Engineering TE is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5306, Restart Signaling for IS-IS (Helper Mode) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5308, Routing IPv6 with IS-IS is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5310, IS-IS Generic Cryptographic Authentication is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6232, Purge Originator Identification TLV for IS-IS is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6233, IS-IS Registry Extension for Purges is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

Management

draft-ieft-snmpv3-update-mib-05, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-idr-bgp4-mib-05, Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-isis-wg-mib-06, Management Information Base for Intermediate System to Intermediate System (IS-IS) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-mpls-ldp-mib-07, Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-mpls-lsr-mib-06, Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-mpls-te-mib-04, Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-ospf-mib-update-08, OSPF Version 2 Management Information Base is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ianaaddressfamilynumbers-mib, IANA-ADDRESS-FAMILY-NUMBERS-MIB is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ianaiftype-mib, IANAifType-MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ianaiprouteprotocol-mib, IANA-RTPROTO-MIB is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE8021-CFM-MIB, IEEE P802.1ag(TM) CFM MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE8021-PAE-MIB, IEEE 802.1X MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE8023-LAG-MIB, IEEE 802.3ad MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

LLDP-MIB, IEEE P802.1AB(TM) LLDP MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1157, A Simple Network Management Protocol (SNMP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1215, A Convention for Defining Traps for use with the SNMP is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1724, RIP Version 2 MIB Extension is supported on Mxp

RFC 2021, Remote Network Monitoring Management Information Base Version 2 using SMIv2 is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2115, Management Information Base for Frame Relay DTEs Using SMIv2 is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2138, Remote Authentication Dial In User Service (RADIUS) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2206, RSVP Management Information Base using SMIv2 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2213, Integrated Services Management Information Base using SMIv2 is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2494, Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type is supported on M(N)

RFC 2571, An Architecture for Describing SNMP Management Frameworks is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2573, SNMP Applications is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2574, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2575, View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2578, Structure of Management Information Version 2 (SMIv2) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2579, Textual Conventions for SMIv2 is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2787, Definitions of Managed Objects for the Virtual Router Redundancy Protocol is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2819, Remote Network Monitoring Management Information Base is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2856, Textual Conventions for Additional High Capacity Data Types is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2863, The Interfaces Group MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2864, The Inverted Stack Table Extension to the Interfaces Group MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2933, Internet Group Management Protocol MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3014, Notification Log MIB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3164, The BSD syslog Protocol is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3165, Definitions of Managed Objects for the Delegation of Management Scripts is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3231, Definitions of Managed Objects for Scheduling Management Operations is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3273, Remote Network Monitoring Management Information Base for High Capacity Networks is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3416. Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3417, Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3419, Textual Conventions for Transport Addresses is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3584, Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3593, Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals is supported on K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3635, Definitions of Managed Objects for the Ethernet-like Interface Types is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3826, The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3877, Alarm Management Information Base (MIB) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3895, Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types is supported on M(N)

RFC 4001, Textual Conventions for Internet Network Addresses is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4022, Management Information Base for the Transmission Control Protocol (TCP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4113, Management Information Base for the User Datagram Protocol (UDP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4220, Traffic Engineering Link Management Information Base is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4292, IP Forwarding Table MIB is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4293, Management Information Base for the Internet Protocol (IP) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6241, Network Configuration Protocol (NETCONF) is supported on K5, K12, R6, and R12

RFC 6242, Using the NETCONF Protocol over Secure Shell (SSH) is supported on K5, K12, R6, and R12

MPLS — General

RFC 3031, Multiprotocol Label Switching Architecture is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3032, MPLS Label Stack Encoding is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3443, Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4182, Removing a Restriction on the use of MPLS Explicit NULL is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5332, MPLS Multicast Encapsulations is supported on T(N), Mxp, R6, and R12

MPLS — GMPLS

draft-ietf-ccamp-rsvp-te-srlg-collect-04, RSVP-TE Extensions for Collecting SRLG Information is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

MPLS — LDP

draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-pdutta-mpls-ldp-v2-00, LDP Version 2 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-pdutta-mpls-tldp-hello-reduce-04, Targeted LDP Hello Reduction is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3037, LDP Applicability is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (Helper Mode) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5036, LDP Specification is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5443, LDP IGP Synchronization is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5561, LDP Capabilities is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

MPLS — MPLS-TP

RFC 5586, MPLS Generic Associated Channel is supported on T(N), R6, and R12

RFC 5921, A Framework for MPLS in Transport Networks is supported on T(N), R6, and R12

RFC 5960, MPLS Transport Profile Data Plane Architecture is supported on T(N), R6, and R12

RFC 6370, MPLS Transport Profile (MPLS-TP) Identifiers is supported on T(N), R6, and R12

RFC 6378, MPLS Transport Profile (MPLS-TP) Linear Protection is supported on T(N), R6, and R12

RFC 6426, MPLS On-Demand Connectivity and Route Tracing is supported on T(N), R6, and R12

RFC 6428, Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile is supported on T(N), R6, and R12

RFC 6478, Pseudowire Status for Static Pseudowires is supported on T(N), R6, and R12

RFC 7213, MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing is supported on T(N), R6, and R12

MPLS — OAM

RFC 6424, Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6425, Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping is supported on T(N), Mxp, R6, and R12

MPLS — RSVP-TE

RFC 2702, Requirements for Traffic Engineering over MPLS is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2747, RSVP Cryptographic Authentication is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2961, RSVP Refresh Overhead Reduction Extensions is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3097, RSVP Cryptographic Authentication -- Updated Message Type Value is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3209, RSVP-TE: Extensions to RSVP for LSP Tunnels is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3477, Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE) is supported on M(N), T(N), X, Mxp, R6, and R12

RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4561, Definition of a Record Route Object (RRO) Node-Id Sub-Object is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4875, Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs) is supported on T(N), Mxp, R6, and R12

RFC 4950, ICMP Extensions for Multiprotocol Label Switching is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5712, MPLS Traffic Engineering Soft Preemption is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5817, Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

OSPF

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1765, OSPF Database Overflow is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2328, OSPF Version 2 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3101, The OSPF Not-So-Stubby Area (NSSA) Option is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3509, Alternative Implementations of OSPF Area Border Routers is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3623, Graceful OSPF Restart Graceful OSPF Restart (Helper Mode) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3630, Traffic Engineering (TE) Extensions to OSPF Version 2 is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4222, Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4552, Authentication/Confidentiality for OSPFv3 is supported on M(N), T(N), X, Mxp, R6, and R12

RFC 4576, Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4577, OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4970, Extensions to OSPF for Advertising Optional Router Capabilities is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5185, OSPF Multi-Area Adjacency is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5187, OSPFv3 Graceful Restart (Helper Mode) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5243, OSPF Database Exchange Summary List Optimization is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5250, The OSPF Opaque LSA Option is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5340, OSPF for IPv6 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5838, Support of Address Families in OSPFv3 is supported on M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6987, OSPF Stub Router Advertisement is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, OAM Procedures for VPWS Interworking is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3916, Requirements for Pseudo- Wire Emulation Edge-to-Edge (PWE3) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3985, Pseudo Wire Emulation Edge-to-Edge (PWE3) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4385, Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4446, IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4447, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5659, An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6073, Segmented Pseudowire is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6310, Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6391, Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network is supported on Mxp, R6, and R12

RFC 6718, Pseudowire Redundancy is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6870, Pseudowire Preferential Forwarding Status bit is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 7023, MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 7267, Dynamic Placement of Multi-Segment Pseudowires is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

Quality of Service

RFC 2430, A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE) is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 2598, An Expedited Forwarding PHB is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3140, Per Hop Behavior Identification Codes is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 3260, New Terminology and Clarifications for Diffserv is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RIP

RFC 1058, Routing Information Protocol is supported on Mxp

RFC 2082, RIP-2 MD5 Authentication is supported on Mxp

RFC 2453, RIP Version 2 is supported on Mxp

Timing

GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

IEEE 1588-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx-1/10GE, R6, and R12

ITU-T G.781, Synchronization layer functions, issued 09/2008 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T G.8261, Timing and synchronization aspects in packet networks, issued 04/2008 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T G.8262, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T G.8264, Distribution of timing information through packet networks, issued 10/2008 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

ITU-T G.8265.1, Precision time protocol telecom profile for frequency synchronization, issued 10/2010 is supported on D-ETR, K5, K12, M(A,N), T(A,N), X, Mxp, Sx-1/10GE, R6, and R12

ITU-T G.8275.1, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014 is supported on X, Mxp, R6, and R12

RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification is supported on D, E, K5, K12, M(A,N), T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

VPLS

RFC 4761, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4762, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5501, Requirements for Multicast Support in Virtual Private LAN Services is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 6074, Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs) is supported on K12, M(N), T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

INDEX

C

customers
33, 79

D

default SAP 39

DHCP
CLI 571

E

encapsulation types
Ethernet 37
SAPs 37

Epipe
overview 147
configuring 203, 204
creating a service 203, 204
SAP 206
local 207

ETH-CFM Support Matrix 233

I

IES
overview 494, 540
filter policies 500
IP interfaces 496, 541
SAP encapsulation 498
configuring
creating a service 504
IES interface 504
management tasks 507
SAPs on IES interface 506

Ipipe
creating
management tasks 218

P

pseudowire
redundancy 160

S

SAPs
overview 34
configuration considerations 43
encapsulation types
Ethernet 37

service access points (SAP) 34

Services
Epipe 147
IES 494, 540
VPLS 279

Services command reference
Epipe 221
Internet Enhances Service (IES) 509, 551
Virtual Leased Line (VLL) 221
Virtual Private LAN Service (VPLS) 395
Virtual Private Routed Network 609

split horizon 931, 937
configuration 934
overview 932

Subscriber services command reference 103

V

VLL
MC-LAG and pseudowire redundancy 176

VPLS
overview 279
MAC learning 288
packet walkthrough 280
STP 299
VPLS over QinQ spokes 288
configuring
basic 344
creating a service 347
management tasks 390
SAP 356
local 356

TSTP bridge parameters 351

Customer Document and Product Support

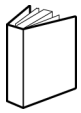

Customer documentation

Customer Documentation Welcome Page

Technical Support

Product Support Portal

Documentation feedback

Customer Documentation Feedback

© 2016, 2017 Nokia.
3HE11478AAAJTQZZA