

Alcatel-Lucent 7705

SERVICE AGGREGATION ROUTER OS | RELEASE 7.0.R4

SERVICES GUIDE

Alcatel-Lucent – Proprietary & Confidential
Contains proprietary/trade secret information which is the property of Alcatel-Lucent. Not to be made available to, or copied or used by anyone who is not an employee of Alcatel-Lucent except when there is a valid non-disclosure agreement in place which covers such information and contains appropriate non-disclosure and limited use obligations.
Copyright © 2015-2016 Alcatel-Lucent. All rights reserved.

Alcatel·Lucent

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

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

Copyright 2015-2016 Alcatel-Lucent.

All rights reserved.

Disclaimers

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

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

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

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

Table of Contents

Preface 25
About This Guide 25
Audience 26
List of Technical Publications 26
Technical Support 27

Getting Started 29
In This Chapter 29
Alcatel-Lucent 7705 SAR Services Configuration Process 29

Services Overview 31
In This Chapter 31
Introduction to Services on the 7705 SAR 32
Overview 32
Service Types 34
Service Policies 36
Alcatel-Lucent Service Model 37
Service Entities 38
Customers 39
Service Types 39
Service Names 39
Service Access Points (SAPs) 40
SAP Encapsulation Types and Identifiers 41
SAP Configuration Considerations 45
Service Destination Points (SDPs) 46
SDP Binding 47
Spoke and Mesh SDPs 48
SDPs and BGP Route Tunnels 49
SDP Encapsulation Types 49
Spoke SDP Terminations 55
SDP Ping 56
SDP Keepalives 56
Mobile Solutions 58
HSDPA Offload 58
Failure Detection 61
Pure DSL Backhaul 62
DSL Channel Bonding 64
ETH-CFM (802.1ag and Y.1731) 65
802.1ag and Y.1731 Terminology 66
MDs, MD Levels, MAs, and MEPs (802.1ag) 67
MEG Levels, MEGs, and MEPs (Y.1731) 69
ETH-CFM Frame Format 69
ETH-CFM OAMPDU 71
CFM Frame Processing 72
MEG-ID and ICC-Based Format 73
ETH-CFM Functions and Tests 74

ETH-CFM Ethernet OAM Tests 75
MEP Support (802.1ag and Y.1731) 81
802.1ag MEP Support on Ethernet SAPs 81
802.1ag MEP Support on Ethernet Spoke SDPs 83
Y.1731 MEP Support on Ethernet SAPs 85
Priority Mapping (802.1ag and Y.1731) 86
Priority Mapping for SAP Up MEPs 86
Priority Mapping for SAP Down MEPs 88
QinQ Support 89
Overview of QinQ 89
QinQ Support with Forced C-Tag Forwarding (VPLS only) 91
Example 1: General QinQ Implementation 92
Example 2: QinQ using VPLS with Ethernet SAPs 92
Example 3: QinQ Using VPLS with ATM SAPs 93
QinQ Support on Ethernet Ports 94
Special QinQ SAP Identifiers 94
QinQ Dot1p Match Behavior 95
QinQ Top-Only Mark Option 96
Maximum Number of VLAN Tags 96
QinQ Configuration Overview 96
Service Creation Overview 97
Port and SAP CLI Identifiers 99
Reference Sources 99
Configuring Global Service Entities with CLI 101
Service Model Entities 102
Basic Configuration 103
Common Configuration Tasks 105
Configuring Customer Accounts 105
Configuring SDPs 106
SDP Configuration Considerations 106
Configuring an SDP 107
Configuring Service Names 109
ETH-CFM (802.1ag and Y.1731) Tasks 110
Configuring ETH-CFM Parameters (802.1ag and Y.1731) 110
Applying ETH-CFM Parameters 112
Service Management Tasks 114
Modifying Customer Accounts 114
Deleting Customers 115
Modifying SDPs 115
Deleting SDPs 116
Deleting LSP Associations 116
Global Service Command Reference 117
Command Hierarchies 117
Global Service Configuration Commands 118
Show Commands 120
Command Descriptions 121
Global Service Configuration Commands 122
Show Commands 137

VLL Services 147
In This Chapter 147
ATM VLL (Apipe) Services 148
ATM VLL for End-to-End ATM Service 148
ATM Virtual Trunk Over an IP/MPLS Packet-Switched Network 149
ATM SAP-to-SAP Service 151
ATM Traffic Management Support 151
Network Ingress Classification 151
ATM Access Egress Queuing and Shaping 152
Control Word 152
Circuit Emulation VLL (Cpipe) Services 153
Cpipe Service Overview 153
TDM SAP-to-SAP Service 154
Cpipe Service Modes 154
Unstructured Mode (SAToP) 154
Structured Mode (CESoPSN) 155
TDM PW Encapsulation 158
Circuit Emulation Parameters and Options 161
Unstructured 163
Structured DS1/E1 CES without CAS 164
Structured T1/E1 CES with CAS 167
Packet Payload Size 170
Jitter Buffer 170
Asymmetric Delay Control 172
RTP Header 174
Control Word 174
Error Situations 175
Ethernet VLL (Epipe) Services 177
Epipe Service Overview 177
Ethernet Access Egress Queuing and Scheduling 179
Ethernet SAP-to-SAP 179
MEF 8 179
Epipe Service Modes 183
Ethernet OAM 185
Control Word 185
MTU 186
Raw and Tagged Modes 186
Raw Mode 187
Tagged Mode 187
VLAN Translation 187
Tagging Rules for Epipe 188
IP Filters (Epipe) 194
Frame Relay VLL (Fpipe) Services 195
Fpipe Service Overview 195
FR SAP-to-SAP Service 196
FR Traffic Management 197
Ingress SAP Classification and Marking 197
FR Encapsulation 197
FR PW Control Word 198

Status Signaling and OAM Propagation 199
HDLC VLL (Hpipe) Services 200
Hpipe Service Overview 200
HDLC VLL for End-to-End HDLC Service 201
HDLC SAP-to-SAP Service 202
HDLC Encapsulation 202
HDLC PW Control Word and Payload Size 203
Status Signaling 205
OAM Propagation 205
IP Interworking VLL (Ipipe) Services 206
Ipipe Service Overview 207
IP Interworking VLL Datapath 208
CE IP Address Discovery and Distribution 209
Manual IP Configuration 209
CE IP Address Discovery Using ARP 209
CE IP Address Discovery for Frame Relay 210
CE IP Address Discovery for cHDLC 210
CE IP Address Discovery for PPP/MLPPP 210
CE IP Address Distribution to Remote PE Nodes 211
IP SAP-to-SAP Service 211
Hardware Support for Interworking IP PWs 211
Control Word 213
Termination at Access 214
PPP and MLPPP Termination 214
FR Termination 214
cHDLC Termination 215
PPP and MLPPP Termination on 2-port OC3/STM1 Channelized Adapter Cards 216
Traffic Management 217
Traffic Management for FR IW SAPs 217
Status Signaling 217
FR LMI 217
IP Filters (Ipipe) 218
Pseudowire Switching 220
Overview 220
Pseudowire Switching with Pseudowire Redundancy 223
Pseudowire Switching Behavior 223
Pseudowire Switching With IP Tunnels 225
Pseudowire Switching TLV 226
VLL Service Considerations 227
Service Support 227
SDPs 231
SDP Statistics for VLL Services 231
SAP Encapsulations and Pseudowire Types 231
ATM PWE3 N-to-1 Cell Mode Encapsulation 234
N-to-1 Cell Mode Encapsulation (N = 1) 235
N-to-1 Cell Mode Encapsulation (N > 1) 237
SAP Aggregation Groups 241
Configuration 241
QoS and Traffic Descriptor Profiles (N > 1) 243

Statistics and Counters 243
Fault Management 244
QoS Policies 245
IP Filter Policies 245
MTU Settings 246
Targeted LDP and MTU 251
QinQ (VLL Service) 252
Pseudowire Control Word 252
Pseudowire Redundancy 252
PW Redundancy Operation 253
Selecting the Active Spoke SDP for PW Redundancy Configuration 255
PW Redundancy and Inter-Chassis Backup 257
AIS Fault Propagation 258
Active/Standby Mode for Pseudowire Redundancy (Standby Signaling) 259
PW Status Signaling Label Withdrawal Option 262
Configuring a VLL Service with CLI 265
Common Configuration Tasks 266
Configuring VLL Components 267
Creating an Apipe Service 268
Configuring Apipe SAP Parameters 269
Configuring Apipe SDP Bindings 271
Configuring Apipe SAP Aggregation Groups 272
Configuring Apipe SAPs as Aggregation Group Members 273
Creating a Cpipe Service 275
Configuring Cpipe SAP Parameters 276
Configuring Cpipe SDP Bindings 279
Creating an Epipe Service 280
Configuring Epipe SAP Parameters 280
Configuring Epipe SAP MEF 8 Parameters 283
Configuring Epipe SAP Microwave Link Parameters for Interworking with TDM2Ethernet 286
Configuring Epipe Spoke SDP Bindings 287
Creating an Fpipe Service 290
Configuring Fpipe SAP Parameters 291
Configuring Fpipe SDP Bindings 293
Creating an Hpipe Service 294
Configuring Hpipe SAP Parameters 295
Configuring Hpipe SDP Bindings 297
Creating an Ipipe Service 298
Configuring Ipipe SAP Parameters 298
Configuring Ipipe SDP Bindings 300
Configuring PW Switching 300
Configuring Ingress and Egress SAP Parameters 303
Using the Control Word 304
Configuring PW Redundancy 305
Configuring PW Redundancy—Standby Signaling 306
Configuring PW Redundancy—ICB 307
Service Management Tasks 309
Modifying Service Parameters 309
Disabling a Service 311

Re-enabling a Service 313
Deleting a Service 313
VLL Services Command Reference 315
Command Hierarchies 315
VLL Services Configuration Commands 316
Show Commands 326
Clear Commands 326
Command Descriptions 327
VLL Service Configuration Commands 328
Show Commands 382
Clear Commands 471

VPLS 475
In This Chapter 475
VPLS Overview 476
VPLS Packet Walkthrough 479
Bridged Mobile Backhaul 483
Multi-Tenant Unit (MTU) Termination 485
VPLS Features 487
VPLS Enhancements 487
Fabric Mode 489
Subscriber VLAN 489
ATM Encapsulated Residential SAP 490
VPLS over MPLS 490
VPLS MAC Learning and Packet Forwarding 491
Pseudowire Control Word 491
Agent Circuit ID Insertion 491
MAC Filters 493
FDB Table Management 493
FDB Size 493
FDB Size Alarms 494
Local and Remote Aging Timers 494
Unknown MAC Discard 495
VPLS and Rate Limiting Via QoS Policy 495
MAC Move 496
Split Horizon Groups (SAP and Spoke SDP) 497
Residential Split Horizon Groups 498
Multicast for VPLS and Routed VPLS (IGMP and MLD Snooping) 498
Sample Scenarios 499
Group and Addressing Support 500
IP Multicast in r-VPLS 501
Multicast Router Ports 504
Tagged Access Traffic 504
Hardware Support 505
Routed VPLS 506
IES or VPRN IP Interface Binding 506
Assigning a Service Name to a VPLS Service 507
Binding a Service Name to an IP Interface 507
Removing a Bound VPLS Service or Service Name 508

IP Interface and VPLS Operational State Coordination 508
IP Interface MTU and Fragmentation 508
ARP/ND and VPLS FIB Interactions 509
Routed VPLS Specific ARP/ND Cache Behavior 510
The allow-ip-int-binding VPLS Flag 511
VPLS Feature Restrictions (due to allow-ip-int-binding) 511
DSCP Marking 511
VPLS SAP Ingress IP Filter Override 512
Routed VPLS Supported Routing-Related Protocols 512
VPLS and Spanning Tree Protocol 513
VPLS Redundancy 513
Spoke SDP Redundancy for Metro Interconnection 514
Spoke SDP-Based Redundant Access 515
VPLS Access Redundancy 515
STP-Based Redundant Access to VPLS 515
MAC Flush Message Processing 516
Dual Homing to a VPLS Service 518
ATM PVC Access and Termination on a VPLS Service 520
VPLS Service Considerations 522
SAP Encapsulations 522
VLAN Processing 522
Tagging Rules for VPLS 523
QinQ (VPLS) 523
Configuration Notes 524
Reference Sources 524
Configuring a VPLS Service with CLI 525
Basic Configuration 526
Common Configuration Tasks 528
Configuring VPLS Components 529
Creating a VPLS Service 529
Creating a Split Horizon Group 530
Enabling MAC Move 530
Configuring STP Bridge Parameters in a VPLS 532
Bridge STP Admin State 532
Mode 533
Bridge Priority 533
Hello Time 533
Hold Count 534
Configuring a VPLS SAP 535
Local VPLS SAPs 535
Distributed VPLS SAPs 536
Configuring SAP-Specific STP Parameters 537
SAP STP Administrative State 537
SAP Virtual Port Number 538
SAP Priority 538
SAP Path Cost 539
SAP Edge Port 539
SAP Auto Edge 540
SAP Link Type 540

STP SAP Operational States 541
Operationally Disabled 541
Operationally Discarding 541
Operationally Learning 542
Operationally Forwarding 542
Configuring VPLS SAPs with Split Horizon 542
Configuring SDP Bindings 543
Configuring Mesh SDP Bindings 544
Configuring Spoke SDP 545
Configuring Spoke SDP Binding 545
Configuring Spoke SDP-Specific STP Parameters 546
Configuring VPLS Spoke SDPs with Split Horizon 550
Configuring Selective MAC Flush 551
Configuring Routed VPLS 552
Configuring IP Multicast in VPLS 553
Configuring IP Multicast in r-VPLS 554
Configuring Multicast Parameters for VPLS and r-VPLS 556
Configuring a Static Multicast Group 558
Service Management Tasks 559
Modifying VPLS Service Parameters 559
Modifying Management VPLS Parameters 560
Deleting a Management VPLS 560
Disabling a Management VPLS 560
Deleting a VPLS Service 561
Disabling a VPLS Service 561
Re-enabling a VPLS Service 561
VPLS Command Reference 563
Command Hierarchies 563
VPLS Service Configuration Commands 564
Show Commands 572
Clear Commands 573
Debug Commands 575
Command Descriptions 576
VPLS Service Configuration Commands 577
Routed VPLS Configuration Commands 643
VPLS Show Commands 645
VPLS Clear Commands 732
VPLS Debug Commands 739

Internet Enhanced Service 745
In This Chapter 745
IES for In-band Management 746
Setting Up Connections Between the 5620 SAM and the 7705 SAR 747
Encapsulation 748
Layer 2 and Layer 3 Traffic Management 748
Troubleshooting and Fault Detection Services 748
IP ECMP Load Balancing 749
IES for Customer Traffic 750
DHCP Relay and DHCPv6 Relay 753

DHCP Relay 753
DHCPv6 Relay 754
IPCP 755
IPSec Support 756
Security Zones and IES 756
Proxy ARP 757
Configurable ARP Retry Timer 758
Unnumbered Interfaces 758
Troubleshooting and Fault Detection Services 759
SAPs 760
Encapsulations 760
Routing Protocols 760
QoS Policies 760
QinQ (IES) 761
IP Filter Policies on an IES SAP 761
Spoke SDP Termination to IES 762
Configuring IES with CLI 765
Common Configuration Tasks 766
Configuring IES Components 767
Creating an IES Service 767
Configuring Interface Parameters 768
IES Management Service 768
IES Service 769
IES IPv6 Service 771
Configuring IES SAP Parameters 773
IES Management SAP 773
IES Service SAP 774
Configuring IES Spoke SDP Parameters 775
Configuring VRRP 776
Configuring a Security Zone within IES 777
Service Management Tasks 779
Modifying IES Service Parameters 779
Disabling an IES Service 779
Re-enabling an IES Service 780
Deleting an IES Service 780
IES Command Reference 781
Command Hierarchies 781
Configuration Commands 782
Show Commands 788
Clear Commands 789
Debug Commands 789
Command Descriptions 790
IES Management Configuration Commands 791
IES Service Configuration Commands 803
Show Commands 857
Clear Commands 880
Debug Commands 882

VPRN Services 883
In This Chapter 883
VPRN Service Overview 884
Routing Prerequisites 885
BGP Support 886
BGP Fast Reroute with Prefix-Independent Convergence in a VPRN 886
IPSec Support 887
Security Zones and VPRN 887
Route Distinguishers 890
PE-to-CE Route Exchange 891
In-Band Management using a VPRN 893
VPRN Features 898
IP Interfaces 898
Unnumbered Interfaces 899
DHCP and DHCPv6 899
IPCP 905
Troubleshooting and Fault Detection Services 906
IP ECMP Load Balancing 906
Proxy ARP 907
Configurable ARP Retry Timer 907
SAPs 908
Encapsulations 908
QoS Policies 908
CoS Marking for Self-generated Traffic 909
QinQ (VPRN) 909
Filter Policies on a VPRN SAP 909
PE-to-CE Routing Protocols 910
Using OSPF in IP VPNs 910
TTL Security 911
PE-to-PE Tunneling Mechanisms 911
Per-VRF Route Limiting 911
RIP Metric Propagation in VPRNs 912
Spoke SDPs 912
Spoke SDP Termination to VPRN 912
IPv6 on Virtual Private Edge Router 916
IPv6 over IPv4 LAN-to-LAN IPSec tunnels 918
Configuring a VPRN Service with CLI 919
Basic Configuration 920
Common Configuration Tasks 921
Configuring VPRN Components 922
Creating a VPRN Service 923
Configuring Global VPRN Parameters 923
Configuring Router Interfaces 924
Configuring BGP for VPRN 925
Configuring VPRN BGP Group and Neighbor Parameters 926
VPRN BGP CLI Syntax 926
Configuring IPv6 Parameters for VPRN BGP 927
Configuring VPRN IPv6 In-Band Management Parameters 928
Configuring VPRN IPv6 Neighbor Discovery Parameters 928

Configuring OSPF for VPRN 929
Configuring RIP for VPRN 930
Configuring a VPRN Interface 932
Configuring a VPRN IPv6 Interface 935
Configuring VPRN Interface Routed VPLS IPv6 Parameters 937
Configuring VPRN Interface SAP Parameters 938
Configuring VPRN Interface SAP IPv6 Parameters 941
Configuring VPRN Interface Spoke SDP Parameters 941
Configuring VPRN Interface Spoke SDP IPv6 Parameters 942
Configuring VRRP 943
Configuring a Security Zone within a VPRN 944
Configuring VPRN Router Advertisement 945
Service Management Tasks 947
Modifying VPRN Service Parameters 947
Deleting a VPRN Service 948
Disabling a VPRN Service 948
Re-enabling a VPRN Service 949
VPRN Services Command Reference 951
Command Hierarchies 951
Configuration Commands 952
Show Commands 973
Clear Commands 974
Debug Commands 975
Command Descriptions 976
VPRN Service Configuration Commands 977
Show Service Commands 1128
Show Router Commands 1170
Clear Service Commands 1221
Debug Service Commands 1227

IPSec 1231
In This Chapter 1231
IPSec Overview 1232
IPSec Implementation 1233
IPSec Overview 1233
Hardware Support 1234
IPSec Encryption Features 1235
SHA2 Support 1236
IPSec Security Policy, IKE Policy, and IPSec Transform 1236
Tunnel Group 1236
Tunnel Interfaces and SAPs 1237
IPSec Tunnel Configuration 1238
X.509v3 Certificate Overview 1239
X.509v3 Certificate Support on the 7705 SAR 1240
Local Storage 1240
CA Profile 1241
CA Chain Computation 1242
Certificate Enrollment 1242
Certificate Revocation Check 1243

Certificate, CRL, and Key Cache 1244
Using Certificates For IPSec Tunnel Authentication 1244
Trust Anchor Profile 1245
Certificate Profile 1245
Certificate Management Protocol Version 2 (CMPv2) 1247
OCSP 1248
Applications 1249
Metrocell Deployment 1249
Small Business Deployment 1250
NAT-Traversal for IKEv2 and IPSec 1250
BFD over IPSec Tunnel 1253
QoS for IPSec 1253
Network and Access Ingress QoS (Decryption QoS) 1253
Network and Access Egress QoS (Encryption QoS) 1254
Fragmentation and IP MTU 1255
Fragmentation Configuration 1256
Reassembly 1256
Support for Private VPRN Service Features 1256
Routing in Private Services 1256
IPSec on the 10-port 1GigE/1-port 10GigE X-Adapter Card 1257
IPSec Sequence Number 1257
PBR and MFC 1257
PBR 1257
MFC 1258
Statistics 1259
Security Support 1259
Public Key Infrastructure (PKI) 1260
CA Role in PKI 1260
Digital Signature and Certificates 1261
Certificates 1261
Certificate Storage 1263
CMPv2 Certificate Management 1264
CMPv2 Initial Registration 1265
Key Update 1265
CRL 1266
OCSP 1266
Best Practices Recommendations 1267
IPSec Best Practices 1267
Configuration Notes 1268
Reference Sources 1268
Configuring IPSec with CLI 1269
Basic Configuration Overview 1270
Common Configuration Tasks 1271
Configuring an IPSec Tunnel Group 1271
Configuring Router Interfaces for IPSec 1272
Configuring IPSec Parameters 1273
Configuring IPSec and IPSec Tunnels in Services 1274
Configuring IPSec IPv6 Parameters for a VPRN Service 1276
Configuring X.509v3 Certificate Parameters 1277

Configuring CMPv2 1279
Configuring OCSP 1280
Service Management Tasks 1281
Deleting an IPSec IKE Policy or an IPSec Transform 1281
Deleting a Public IPSec Tunnel SAP and Interface 1282
Deleting a Private IPSec Tunnel SAP and Interface 1282
Deleting an IPSec Security Policy 1282
Deleting an IPSec Tunnel 1283
IPSec Command Reference 1285
Command Hierarchies 1285
IPSec Configuration Commands 1286
PKI Configuration Commands 1289
Show Commands 1292
Clear Commands 1292
Debug Commands 1293
Command Descriptions 1294
IPSec Configuration Commands 1295
PKI Configuration Commands 1320
Show Commands 1345
Clear Commands 1360
Debug Commands 1361

Network Group Encryption 1363
In This Chapter 1363
NGE Overview 1364
NGE Security Domains and Encryption Partitions 1365
5620 SAM Management 1367
NGE Global Encryption Label 1368
Key Groups 1369
Key Group Algorithms 1370
Encapsulating Security Payload (ESP) 1371
Security Associations (SAs) 1372
Active Outbound SA 1372
Services Encryption 1373
Services Encryption Overview 1373
Assigning Key Groups to Services 1375
Pseudowire Switching for NGE Traffic 1376
Pseudowire Control Word for NGE Traffic 1376
VPRN Layer 3 Spoke-SDP Encryption and MP-BGP-based VPRN Encryption Interaction 1377
NGE and RFC 3107 1377
NGE Packet Overhead and MTU Considerations 1378
QoS for NGE Traffic 1380
Network Ingress 1380
Network Egress 1381
Statistics 1382
Remote Network Monitoring (RMON) Support 1383
Configuration Notes 1384
Reference Sources 1384
Configuring NGE with CLI 1385

Basic NGE Configuration Overview 1386
Configuring NGE Components 1387
Configuring the Global Encryption Label 1387
Configuring a Key Group 1388
Assigning a Key Group to an SDP or Service 1389
Service Management Tasks 1391
Modifying a Key Group 1391
Removing a Key Group from an SDP or VPRN Service 1393
Changing the Key Group for an SDP or VPRN Service 1394
Deleting a Key Group from a 7705 SAR 1395
NGE Command Reference 1397
Command Hierarchies 1397
Configuration Commands 1398
Show Commands 1399
Command Descriptions 1400
Configuration Commands 1401
Show Commands 1407

Standards and Protocol Support 1439

List of Tables

Getting Started 29
Table 1 7705 SAR Configuration Process 29

Services Overview 31
Table 2 Pseudowire Service Types 35
Table 3 Service Types and SAP Encapsulations 44
Table 4 GRE Header Descriptions 51
Table 5 GRE Pseudowire Payload Packet Descriptions 53
Table 6 IP Pseudowire Payload Packet Descriptions 54
Table 7 Spoke SDP Termination Support 56
Table 8 802.1ag Terminology 66
Table 9 Y.1731 Terminology 67
Table 10 CFM Frame Processing 72
Table 11 FC and VLAN Priority Mappings for Up and Down MEPs as per Frame Type 87
Table 12 Special QinQ SAP Identifiers 94
Table 13 SDP Echo Reply Response Conditions 133
Table 14 Show Customer Command Output Fields 138
Table 15 Show Service SDP Output Fields 140
Table 16 Show Service SDP-Using Output Fields 143
Table 17 Show Service service-using Output Fields 146

VLL Services 147
Table 18 T1 Framing for CAS (RBS) Support in a T1 ESF Multiframe 157
Table 19 SAToP and CESoPSN Support on the 7705 SAR 161
Table 20 Unstructured Payload Defaults 164
Table 21 Default and Minimum Payload Size for CESoPSN without CAS 166
Table 22 Default Values for the Payload Size for T1 and E1 CESoPSN with CAS 168
Table 23 Control Word Bit Descriptions 174
Table 24 MEF 8 Support on the 7705 SAR 182
Table 25 Ingress SAP Tagging Rules 188
Table 26 Egress SAP Tagging Rules 189
Table 27 VLAN Tagging Examples (Epipe) 190
Table 28 Control Word Bit Descriptions 199
Table 29 Control Word Bit Descriptions 204
Table 30 Hardware Support for Interworking PWs 211
Table 31 Supported Pseudowire Switching Options 222
Table 32 Pseudowire Switching TLV Field Descriptions 226
Table 33 MTU Points and Descriptions 247
Table 34 Service MTU Default and Maximum Values 248
Table 35 MTU Calculator—Service Creation (Worst Case) Access Ports and SAPs 249
Table 36 MTU Calculator—Service Creation (Worst Case) Network Ports 250
Table 37 Matching MTU or Payload Values for Signaled VLL Services 251
Table 38 Service MTU Values 340
Table 39 SAP ID Configurations 342
Table 40 Port and Encapsulation Values 345
Table 41 SAP ID Preconfiguration for SAP Aggregation Groups 350

Table 42 Match-QinQ-Dot1p Matching Behavior 359
Table 43 Dot1P Re-marking Behavior for the QinQ-mark-top-only Command 361
Table 44 Show Service-ID All Command Output Fields 385
Table 45 Show Service-ID Base Output Fields 426
Table 46 Show Service Egress Label Output Fields 428
Table 47 Show Service Ingress Label Output Fields 430
Table 48 Service-ID Endpoint Output Fields 431
Table 49 Service-ID Labels Output Fields 432
Table 50 Service-ID SAP Output Fields 435
Table 51 Service-ID SAP Aggregation Group Output Fields 454
Table 52 Show Service SAP-Using Output Fields 460
Table 53 Show Service SCADA Bridge SAP Fields 461
Table 54 SDP Output Fields 463
Table 55 Sdp-using Output Fields 468
Table 56 Services-using Output Fields 469

VPLS 475
Table 57 Ingress Behavior for VPLS Next-Hop Routing 510
Table 58 Egress Behavior for VPLS Next-Hop Routing 510
Table 59 Show Service Service-MTU Output Fields 593
Table 60 Match-QinQ-Dot1p Matching Behavior 624
Table 61 Dot1P Re-marking Behavior for the QinQ-mark-top-only Command 625
Table 62 Show Service Egress Labels Output Fields 646
Table 63 Show Service FDB-Info Output Fields 647
Table 64 Show Service FDB-MAC Output Fields 649
Table 65 Show Service Service-ID (All) Output Fields 658
Table 66 Show Service Service-ID (Base) Output Fields 660
Table 67 Show Service Service-ID (DHCP Statistics) Output Fields 663
Table 68 Show Service Service-ID (DHCP Summary) Output Fields 664
Table 69 Show Service Service-ID (Endpoint) Output Fields 665
Table 70 Show Service Service-ID (FDB) Output Fields 668
Table 71 Show Service Service-ID (IGMP and MLD Snooping) Output Fields 676
Table 72 Show Service Service-ID IGMP and MLD Snooping (Base) Output Fields 684
Table 73 Show Service Service-ID IGMP and MLD Snooping (Port-DB) Output Fields 688
Table 74 Show Service Service-ID IGMP and MLD Snooping (Proxy-DB) Output Fields 692
Table 75 Show Service Service-ID IGMP and MLD Snooping (Querier) Output Fields 693
Table 76 Show Service Service-ID IGMP and MLD Snooping (Static) Output Fields 695
Table 77 Show Service Service-ID IGMP and MLD Snooping (Statistics) Output Fields 698
Table 78 Show Service Service-ID (Labels) Output Fields 700
Table 79 Show Service Service-ID (MAC Move) Output Fields 702
Table 80 Show Service ID (SAP) Output Fields 708
Table 81 Show Service ID (SDP) Output Fields 717
Table 82 Show Service ID (Split Horizon Group) Output Fields 722
Table 83 Show Service ID (STP) Output Fields 724
Table 84 Show Service Ingress-Label Output Fields 729
Table 85 Show Service SAP-Using Output Fields 731

Internet Enhanced Service 745
Table 86 Security Zone Interfaces per Context 757

Table 87 Match-QinQ-Dot1p Matching Behavior 840
Table 88 Dot1P Re-marking Behavior for the QinQ-mark-top-only Command 842
Table 89 Show Service Customer Output Fields 857
Table 90 Show Service Egress Output Fields 859
Table 91 Show Service ID All Output Fields 862
Table 92 Show Service ID ARP Output Fields 867
Table 93 Show Service ID Base Output Fields 868
Table 94 Show Service ID DHCP Statistics Output Fields 869
Table 95 Show Service ID DHCP Summary Output Fields 871
Table 96 Show Service ID Interface Output Fields 872
Table 97 Show Service ingress Output FIelds 875
Table 98 Show Service SAP-Using Output Fields 878
Table 99 Show Service Service-Using Output Fields 879

VPRN Services 883
Table 100 BGP FRR Scenarios 887
Table 101 Security Zone Interfaces per Context 888
Table 102 Route Distinguisher Type-Value Fields 890
Table 103 IPv4 and IPv6 GRT-Supported Management Protocols 894
Table 104 6VPE Access Control List, SAP 917
Table 105 6VPE Access Control List, SDP 917
Table 106 6VPE Access Control List, r-VPLS Override 917
Table 107 Applications and Default Values for DSCP or dot1p Markings 990
Table 108 Route Preference Defaults by Route Type 1038
Table 109 VPRN Interface State and IP Address 1073
Table 110 Match-QinQ-Dot1p Matching Behavior 1108
Table 111 Dot1P Re-marking Behavior for the QinQ-mark-top-only Command 1110
Table 112 Show Service Egress Output Fields 1129
Table 113 Show Service ID All Output Fields 1134
Table 114 Show Service ID ARP Output Fields 1144
Table 115 Show Service ID Base Output Fields 1146
Table 116 Show Service ID DHCP Statistics Output Fields 1148
Table 117 Show Service ID DHCP Summary Output Fields 1149
Table 118 Show Service ID Interface Detailed Output Fields 1152
Table 119 Show Service ID SAP Detailed Output Fields 1158
Table 120 Show Service ingress Output FIelds 1162
Table 121 Show Service ID SDP Detailed Output Fields 1164
Table 122 Show TWAMP Light Output Fields 1167
Table 123 Show Service Service-Using Output Fields 1168
Table 124 Show Aggregate Table Output Fields 1171
Table 125 Show ARP Table Output Fields 1172
Table 126 Show BGP Damping Output Fields 1177
Table 127 Show BGP Group Output Fields 1180
Table 128 Show BGP Neighbor (Standard and Detailed) Output Fields 1186
Table 129 Show BGP Neighbor (Advertised-Routes and Received-Routes) Output Fields 1193
Table 130 BGP Neighbor (Graceful Restart) Output Fields 1194
Table 131 Show BGP Next-Hop Output Fields 1197
Table 132 Show BGP Path Output Fields 1199
Table 133 Show BGP Route Output Fields 1204

Table 134 Show BGP Summary Output Fields 1208
Table 135 Show DHCP Statistics Output Fields 1210
Table 136 Show DHCP Summary Output Fields 1211
Table 137 Show IP Interface Output Fields 1212
Table 138 Show Route Table Output Fields 1213
Table 139 Application QoS Output Fields 1215
Table 140 DSCP-to-FC Mapping Output Fields 1217
Table 141 Show Static ARP Table Output Fields 1218
Table 142 Show Static Route Output Fields 1219
Table 143 Show Tunnel Table Output Fields 1220

IPSec 1231
Table 144 Show IPSec IKE-Policy Command Output Fields 1349
Table 145 Show IPSec Security Policy Command Output Fields 1352
Table 146 Show IPSec Transform Command Output Fields 1353
Table 147 Show IPSec Tunnel Command Output Fields 1356

Network Group Encryption 1363
Table 148 NGE Overhead 1378
Table 149 Accounting for NGE Overhead SDP and Service MTU — Sample Calculations 1379
Table 150 Show Encryption Key Group Output Fields 1409
Table 151 Show Group Encryption Summary Output Fields 1412

List of Acronyms 1413
Table 152 Acronyms 1413

Standards and Protocol Support 1439
Table 153 EMC Industrial Standards Compliance 1440
Table 154 EMC Regulatory and Customer Standards Compliance 1441
Table 155 Environmental Standards Compliance 1443
Table 156 Safety Standards Compliance 1444
Table 157 Telecom Interface Compliance 1445
Table 158 Directives, Regional Approvals and Certifications Compliance 1447

List of Figures

Services Overview 31

Figure 1 Service Entities and the Service Model 38
Figure 2 Service Access Point (SAP) 41
Figure 3 Multiple SAPs on a Single Port/Channel 42
Figure 4 SDP Tunnel Pointing from ALU-A to ALU-B 48
Figure 5 GRE Header 51
Figure 6 GRE Pseudowire Payload Packet over Ethernet 52
Figure 7 IP Example of Pseudowire Payload Packet over Ethernet 54
Figure 8 7705 SAR-F HSDPA Offload Example 59
Figure 9 7705 SAR-M With 8-port xDSL Module for HSDPA Offload 60
Figure 10 7705 SAR-M With 6-port DSL Combination Module for HSDPA Offload 60
Figure 11 Pure DSL Backhaul Network Example 62
Figure 12 "Sample DSL Network Traffic 63
Figure 13 MEPs and MAs 68
Figure 14 ETH-CFM Frame Format 69
Figure 15 ETH-CFM OAMPDU Message 71
Figure 16 ICC-based MEG-ID Format 74
Figure 17 Dot1ag Loopback Test 76
Figure 18 MEP on Ethernet Access 82
Figure 19 Down MEP at Ethernet SAP 83
Figure 20 Dot1ag Down MEPs on Spoke SDPs 84
Figure 21 Y.1731 MEP Support on the 7705 SAR 85
Figure 22 QinQ Frame 89
Figure 23 QinQ Tagging Example 91
Figure 24 QinQ Using VPLS Ethernet SAPs 92
Figure 25 QinQ Using VPLS ATM SAPs 93
Figure 26 Service Creation and Implementation Flowchart 98

VLL Services 147

Figure 27 ATM VLL for End-to-End ATM Service 149
Figure 28 ATM Virtual Trunk Over IP/ MPLS Packet-Switched Network 150
Figure 29 E1 Framing for CAS Support in an E1 Multiframe 156
Figure 30 SAToP MPLS Encapsulation 159
Figure 31 CESoPSN MPLS Encapsulation 159
Figure 32 CESoPSN Packet Payload Format for Trunk-Specific n x 64 kb/s (With and Without CAS Transport) 160
Figure 33 Control Word Bit Structure 174
Figure 34 Ethernet VLL Frame with MPLS Encapsulation 178
Figure 35 Epipe Service 178
Figure 36 CESoETH Encapsulation 180
Figure 37 TDM SAP to Ethernet SAP 181
Figure 38 TDM SAP to Ethernet SAP via 7705 SAR-M GPON Module 181
Figure 39 TDM SAP to Spoke SDP 182
Figure 40 FR VLL for End-to-End FR Service 196
Figure 41 FR Frame 197
Figure 42 FR PW 1-to-1 MPLS PSN Encapsulation 198

Figure 43 FR PW Control Word 198
Figure 44 HDLC VLL for End-to-End HDLC Service....201
Figure 45 HDLC VLL Frame 202
Figure 46 HDLC VLL Frame with MPLS Encapsulation....203
Figure 47 HDLC PW Control Word 204
Figure 48 IP Pseudowires Between SAR Nodes 207
Figure 49 FR Header with NLPID Support for IP Interworking PW....215
Figure 50 cHDLC Header Frame 215
Figure 51 SLARP Keepalive Frame....216
Figure 52 PPP and MLPPP to IP PWs on 2-port OC3/STM1 Channelized Adapter Cards 216
Figure 53 Simplex to Simplex Pseudowire Switching 221
Figure 54 Simplex to Redundant Pseudowire Switching 223
Figure 55 Pseudowire Switching Network....224
Figure 56 Pseudowire Switching TLV 226
Figure 57 N-to-1 Cell Mode Encapsulation 234
Figure 58 VCC Cell Mode Encapsulation with N > 1 238
Figure 59 MTU Points on the 7705 SAR....246
Figure 60 Pseudowire Redundancy....253
Figure 61 Implicit and Explicit Endpoint Objects....254
Figure 62 Pseudowire Redundancy with Four Spoke SDPs....256
Figure 63 Access Side Failure with ICB Protection....257
Figure 64 Network Side Failure with ICB Protection....258
Figure 65 Active/Standby Mode for Redundant Pseudowires....259
Figure 66 Standby-signaling-master Enabled....260
Figure 67 Standby-signaling-slave Enabled....261
Figure 68 Mixed Microwave Link Scenario 286
Figure 69 SDPs—Unidirectional Tunnels....287

VPLS 475
Figure 70 VPLS Service Architecture....479
Figure 71 Access Port Ingress Packet Format and Lookup....480
Figure 72 Network Port Egress Packet Format and Flooding....481
Figure 73 Access Port Egress Packet Format and Lookup 482
Figure 74 Typical Pseudowire-based Mobile Backhaul 483
Figure 75 Local VPLS on 7705 SAR in Mobile Backhaul....484
Figure 76 Spoke-SDP Termination to VPLS using 7705 SAR-18 Routers 486
Figure 77 ATM and IP DSLAM Backhaul....488
Figure 78 PPPoE Initialization and Agent-Id Push Function 492
Figure 79 Agent Circuit ID Information....492
Figure 80 Video over Mobile Backhaul 499
Figure 81 Metro Cell Multicast 500
Figure 82 Multiple Hosts in an r-VPLS Receiving the Same Channel 502
Figure 83 Multiple Hosts in an r-VPLS Receiving Different Channels 503
Figure 84 H-VPLS with Spoke Redundancy 514
Figure 85 Dual Homed MTUs in Two-Tier Hierarchical VPLS 516
Figure 86 Dual-Homed CE Connection to VPLS 519
Figure 87 Example of ATM PVC Access and Termination on a VPLS 520
Figure 88 SDPs—Unidirectional Tunnels....544

Internet Enhanced Service 745
Figure 89 IES for Customer Access to the Internet 750
Figure 90 SDP ID and VC Label Service Identifiers (Conceptual View of the Service) 762
Figure 91 IES Spoke SDP Termination 763
Figure 92 Pseudowire-Based Backhaul (Spoke SDP Termination at 7750 SR) 764

VPRN Services 883
Figure 93 Virtual Private Routed Network 885
Figure 94 Firewall Protection for the Network Core 889
Figure 95 Route Distinguisher Structure 890
Figure 96 Directly Connected IP Target 892
Figure 97 Multiple Hops to IP Target 892
Figure 98 IPv4 In-Band Management Using a VPRN Configured with GRT Lookup 895
Figure 99 IPv6 In-Band Management Using a VPRN Configured with GRT Lookup 896
Figure 100 RIP Metric Propagation in VPRNs 912
Figure 101 SDP ID and VC Label Service Identifiers (Conceptual View of the Service) 913
Figure 102 VPRN Spoke SDP Termination 913
Figure 103 Pseudowire-Based Backhaul (Spoke SDP Termination at 7750 SR) 915
Figure 104 VPRN in Mobile Backhaul Application 915
Figure 105 Spoke-SDP Termination to VPRN 916
Figure 106 Access IPv6 Traffic Aggregation and Encryption 918

IPSec 1231
Figure 107 IPSec Implementation Architecture 1234
Figure 108 Typical Metrocell Deployment 1249
Figure 109 Typical Small Business Deployment 1250
Figure 110 UDP Header Injected by a NAT-T-enabled IPSec Tunnel 1252

Network Group Encryption 1363
Figure 111 Any-to-Any Encryption of MPLS-based Services 1364
Figure 112 Key-Group Partitioning 1366
Figure 113 Key Groups and a Typical NGE Packet 1370
Figure 114 NGE MPLS/GRE Label Stack 1374
Figure 115 NGE Encryption and Packet Formats 1374
Figure 116 Inbound and Outbound Key-Group Assignments 1375
Figure 117 QoS for NGE Traffic (Network Ingress) 1381

List of Figures

Preface

About This Guide

This guide describes subscriber services support provided by the 7705 Service Aggregation Router (7705 SAR) and presents examples to configure and implement various protocols and services.

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



Note: This manual generically covers Release 7.0 content and may contain some content that will be released in later maintenance loads. Please refer to the 7705 SAR OS 7.0.Rx Software Release Notes, part number 3HE10099000xTQZZA, for information on features supported in each load of the Release 7.0 software.



Note:

In Release 7.0, support for the following hardware has been deprecated:

- CSMv1
- 7705 SAR-F
- 8-port Ethernet Adapter card, version 1
- 16-port T1/E1 ASAP Adapter card, version 1

These components are no longer recognized in the release.

Audience

This guide is intended for network administrators who are responsible for configuring the 7705 SAR routers. It is assumed that the network administrators have an understanding of networking principles and configurations. Concepts described in this guide include the following:

- CLI concepts
- subscriber services

List of Technical Publications

The 7705 SAR OS documentation set is composed of the following guides:

- 7705 SAR OS Basic System Configuration Guide

  This guide describes basic system configurations and operations.
- 7705 SAR OS System Management Guide

  This guide describes system security and access configurations as well as event logging and accounting logs.
- 7705 SAR OS Interface Configuration Guide

  This guide describes card and port provisioning.
- 7705 SAR OS Router Configuration Guide

  This guide describes logical IP routing interfaces, filtering, and routing policies.
- 7705 SAR OS MPLS Guide

  This guide describes how to configure Multiprotocol Label Switching (MPLS), Resource Reservation Protocol for Traffic Engineering (RSVP-TE), and Label Distribution Protocol (LDP).
- 7705 SAR OS Services Guide

  This guide describes how to configure service parameters such as service access points (SAPs), service destination points (SDPs), customer information, and user services.
- 7705 SAR OS Quality of Service Guide

  This guide describes how to configure Quality of Service (QoS) policy management.
- 7705 SAR OS Routing Protocols Guide

  This guide provides an overview of dynamic routing concepts and describes how to configure them.

- 7705 SAR OS OAM and Diagnostics Guide

  This guide provides information on Operations, Administration and Maintenance (OAM) tools.

Technical Support

If you purchased a service agreement for your 7705 SAR router and related products from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance. If you purchased an Alcatel-Lucent service agreement, follow this link to contact an Alcatel-Lucent support representative and to access product manuals and documentation updates:

http://support.alcatel-lucent.com

Getting Started

In This Chapter

This chapter provides the process flow information required to configure services.

Alcatel-Lucent 7705 SAR Services Configuration Process

Table 1 lists the tasks necessary to configure subscriber services. 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: 7705 SAR Configuration Process

Area	Task	Reference
Global entities	Configure global entities	Configuring Global Service Entities with CLI
VLL services	Apipe service	ATM VLL (Apipe) Services
	Cpipe service	Circuit Emulation VLL (Cpipe) Services
	Epipe service	Ethernet VLL (Epipe) Services
	Fpipe service	Frame Relay VLL (Fpipe) Services
	Hpipe service	HDLC VLL (Hpipe) Services
	Ipipe service	IP Interworking VLL (Ipipe) Services
VPLS service	Configure VPLS service	VPLS
Internet Enhanced Service	Configure a routed-IP connectivity or Internet access service and configure in-band management of the 7705 SAR over ATM links	Internet Enhanced Service

Table 1: 7705 SAR Configuration Process (Continued)

Area	Task	Reference
VPRN	Configure a Layer 3 multipoint-to-multipoint VPN service as defined in RFC 2547bis	VPRN Services
IPSec	Configure IPSec service	IPSec
Network group encryption	Configure network group encryption	Network Group Encryption
Reference	List of IEEE, IETF, and other proprietary entities	Standards and Protocol Support

Services Overview

In This Chapter

This chapter provides an overview of the 7705 SAR subscriber services, service model, and service entities. Additional details on the individual subscriber services are found in subsequent chapters.

Topics in this chapter include:

- Introduction to Services on the 7705 SAR
- Alcatel-Lucent Service Model
- Service Entities
- Mobile Solutions
- ETH-CFM (802.1ag and Y.1731)
- QinQ Support
- Service Creation Overview
- Port and SAP CLI Identifiers
- Configuring Global Service Entities with CLI
- Global Service Command Reference

Introduction to Services on the 7705 SAR

Topics in this section include:

- Overview
- Service Types
- Service Policies

Overview

A service is a type of telecommunications connection from one place to another. These telecommunications connections have the particular attributes and characteristics that are needed to provide a specific communications link through which an information flow or exchange can occur. The 7705 Service Access Router (7705 SAR) offers VLL services, Layer 2 multipoint VPN services (VPLS), Layer 3 MPLS VPN services (VPRN), and Layer 3 routed/IP services.

The 7705 SAR service model uses (logical) service entities to construct a service. These logical entities provide a uniform, service-centric configuration, management, and billing model for service provisioning (see Alcatel-Lucent Service Model for more information). Many services can be created on the same 7705 SAR at the same time, and each service is uniquely identified by a service ID.

The 7705 SAR offers Virtual Leased Line (VLL) services (also referred to as pseudowire (PW) services or pipes), which emulate a Layer 1/2 entity, such as a wire or a leased line. These emulated services provide connectivity between a service access point (SAP) on one 7705 SAR and on another SAP on the same router, or on a remote 7705 SAR, 7710 SR, or 7750 SR. VLL services offer SAP logical entities—such as a VLAN or a virtual connection—Layer 2 visibility or processing (IMA termination). A SAP is the point where customer traffic enters and exits the service.

A connection between two SAPs on the same router is known as a local service. A connection between SAPs on a local and a remote router is known as a distributed service. SAP-to-SAP connections are supported for ATM, Ethernet, FR, HDLC, and TDM VLLs.

The 7705 SAR also supports the aggregation of multiple ATM VCC SAPs into a single aggregation group within a service. Aggregation groups enable service providers to make more efficient use of the ATM VCC VLL services that are deployed in a network.

Distributed services use service destination points (SDPs) to direct traffic from a local router to a remote router through a service tunnel. An SDP is created on the local router and identifies the endpoint of a logical unidirectional service tunnel. Traffic enters the tunnel at the SDP on the local router and exits the tunnel at the remote router. Hence, a service tunnel provides a path from a 7705 SAR to another service router, such as another 7705 SAR, a 7710 SR, or a 7750 SR. Because an SDP is unidirectional, two service tunnels are needed for bidirectional communication between two service routers (one SDP on each router).

SDPs are configured on each participating 7705 SAR or service router, specifying the address of the source router (the 7705 SAR participating in the service communication) and the address of the destination router, such as another 7705 SAR or service router. After SDPs are created, they are bound to a specific service. The binding process is needed to associate the far-end devices to the service; otherwise, far-end devices are not able to participate in the service.

The 7705 SAR also offers IES, VPLS, and VPRN services.

IES provides IP connectivity between customer access points. From the customer's perspective, IES provides direct IP connectivity. The customer is assigned an IP interface and a SAP designates the customer access point where the customer IP device is connected—one SAP binding per IP interface. Supported SAP encapsulations are MC-MLPPP, PPP/MLPPP and null/dot1q/qinq Ethernet. SDP binding is not required because traffic is routed rather than being encapsulated in a tunnel.

A Virtual Private Routed Network (VPRN) consists of a set of customer sites connected to one or more PE routers. VPRNs are based on RFC 2547bis, which details a method of distributing routing information and forwarding data to provide a Layer 3 Virtual Private Network (VPN) service to end customers. VPRN traffic is transported over LDP and RSVP-TE tunnels, as well as static LSPs.

A Virtual Private LAN Service (VPLS) enables Layer 2 multipoint connections within an enterprise infrastructure. Supported SAP encapsulations are null/dot1q/qinq Ethernet (on the 8-port Ethernet Adapter card, 6-port Ethernet 10Gbps Adapter card, 7705 SAR-X, 8-port Gigabit Ethernet Adapter card, and the 10-port 1GigE/1-port 10GigE X-Adapter card) and ATM (on the 4-port OC3/STM1 Clear Channel Adapter card).

Note:

- The 10-port 1GigE/1-port 10GigE X-Adapter card supports qinq only on version 2 when it is in 10-port 1GigE mode.
- The 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X support qinq only when the card is in access mode.

VPLS traffic can also be transported over existing tunnel types like GRE tunnels, LDP tunnels, RSVP-TE tunnels, and static LSPs using SDPs. For the ATM SAPs, the Layer 2 Ethernet frames are encapsulated in llc-snap bridged PDUs, as per RFC 2684, widely referred to with the obsoleted RFC 1483.

Service Types

Services are commonly called customer or subscriber services. The 7705 SAR offers the following types of services, which are described in more detail in the referenced chapters:

- Virtual Leased Line (VLL) services
  - → ATM VLL (Apipe)—a pseudowire emulation edge-to-edge (PWE3) ATM service over MPLS, GRE, or IP tunnels on 7705 SAR nodes. See ATM VLL (Apipe) Services.
  - → Circuit emulation VLL (Cpipe)—a PWE3 circuit emulation service over MPLS or GRE tunnels on 7705 SAR nodes. See Circuit Emulation VLL (Cpipe) Services.
  - → Ethernet VLL (Epipe)—a PWE3 Ethernet service over MPLS or GRE tunnels for Ethernet frames on 7705 SAR nodes. See Ethernet VLL (Epipe) Services.
  - → FR VLL (Fpipe)—a PWE3 FR service over MPLS or GRE tunnels for FR frames on 7705 SAR nodes. See Frame Relay VLL (Fpipe) Services.
  - → HDLC VLL (Hpipe)—a PWE3 HDLC service over MPLS or GRE tunnels for HDLC frames on 7705 SAR nodes. See HDLC VLL (Hpipe) Services.
  - → IP interworking VLL (Ipipe)—a PWE3 IP service between two hosts connected by any combination of point-to-point access circuits. IP interworking VLLs can operate over MPLS or GRE networks. Some typical examples are Ethernet SAP to Ethernet SAP, Ethernet SAP to MLPPP SAP, Ethernet SAP to LAG SAP, FR SAP to Ethernet SAP, or cHDLC SAP to Ethernet SAP. See IP Interworking VLL (Ipipe) Services.
- Internet Enhanced Service (IES)
  - → IES is used both as a direct Internet access service where the customer is assigned an IP interface for routed connectivity and for in-band management of the 7705 SAR. See Internet Enhanced Service.
- Virtual Private LAN Service (VPLS)
  - → VPLS provides a Layer 2 multipoint VPN service to end customers. VPLS includes Hierarchical VPLS (H-VPLS), which is an enhancement of VPLS that extends pseudowire-style signaled or static virtual circuit labeling outside the fully meshed VPLS core. The 7705 SAR can participate in hierarchical VPLS. In addition, the 7705 SAR supports management VPLS (mVPLS) via Rapid Spanning Tree Protocol (RSTP). See VPLS.

- Virtual Private Routed Network Service (VPRN)
  → VPRN provides a Layer 3 VPN service to end customers. VPRN services provide MP-BGP peering with other PEs, configurable QoS policy and filtering, VRF import and export policies, and SGT-QoS marking. See VPRN Services.

Table 2 lists the supported pseudowire (PW) service types. The values are as defined in RFC 4446.

Table 2: Pseudowire Service Types

PW Service Type (Ethertype)	Value
IP Layer 2 transport	0x000B
Ethernet tagged mode	0x0004
Ethernet raw	0x0005
HDLC	0x0006
ATM N-to-one VCC cell mode [1]	0x0009
ATM N-to-one VPC cell mode	0x000A
ATM transparent cell transport mode	0x0003
SAToP E1	0x0011
SAToP T1	0x0012
CESoPSN basic mode	0x0015
CESoPSN TDM with CAS	0x0017
FR DLCI	0x0019

Note:

1. "N-to-one" is expressed as "N-to-1" throughout this guide.

Service Policies

Common to all 7705 SAR connectivity services are policies that are assigned to the service. Policies are defined at the global level, then applied to a service on the router. Policies are used to define 7705 SAR service enhancements.

The types of policies that are common to all 7705 SAR connectivity services are SAP Quality of Service (QoS) policies and accounting policies. Filter policies are supported on network interfaces, Epipes, Ipipes, VPLS SAPs and SDPs (mesh and spoke), VPRN SAPs and spoke SDPs, IES SAPs and spoke SDPs, and IES in-band management SAPs.

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

  QoS ingress and egress policies determine the QoS characteristics for a SAP. A QoS policy applied to a SAP specifies the number of queues, queue characteristics (such as forwarding class, committed information rates, and peak information rates) 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 a single egress QoS policy can be associated with a SAP.

  QoS ingress and egress policies also apply to SAP aggregation groups.
- Accounting policies define how to count the traffic usage for a service for billing purposes.

  The 7705 SAR routers provide a comprehensive set of service-related counters. Accounting data can be collected on a per-service, per-forwarding class basis, which enables network operators to accurately measure network usage and bill each customer for each individual service using any of a number of different billing models.
- Filter policies, also referred to as access control lists (ACLs), allow selective blocking or forwarding of traffic that matches criteria that is set in the policy. The resulting action (block or forward) is applied to that traffic.

  Filter policies control the traffic allowed into a SAP or SDP, and are based on IP or MAC match criteria. The ability to configure and apply a filter depends on the combination of service, traffic type and direction, and entity type (SAP or SDP). Assigning a filter policy to a SAP or SDP is optional. Filter policies are identified by a unique filter policy ID. A filter policy must be created before it can be applied. A single ingress and a single egress filter policy can be assigned to a SAP (if supported), and a single ingress filter policy can be assigned to an SDP (if supported).

For more information on provisioning QoS policies, including queuing behaviors, refer to the 7705 SAR OS Quality of Service Guide. For information on configuring IP and MAC filter policies, refer to the 7705 SAR OS Router Configuration Guide.

Alcatel-Lucent Service Model

The 7705 SAR routers are deployed at the provider edge (PE). Services are provisioned on the 7705 SAR and other network equipment in order to facilitate the transport of telecommunications data across an IP/MPLS provider's core network. The data is formatted so that it can be transported in encapsulation tunnels created using generic routing encapsulation (GRE), IP encapsulation, or MPLS label switched paths (LSPs).

The service model has four main logical components, referred to as (logical) service entities. The entities are: customers, service types, service access points (SAPs), and service destination points (SDPs) (see Service Entities). In accordance with the service model, the operator uses the (logical) service entities to construct an end-to-end service. The service entities are designed to provide a uniform, service-centric model for service provisioning. This service-centric design implies the following characteristics.

- 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 service entity rather than to multiple ports on multiple devices. It is easier to change one tunnel rather than several services.
- The operational integrity of a service entity (such as a service tunnel or service endpoint) can be verified by one operation rather than through the verification of dozens of parameters, thereby simplifying management operations, network scalability, and performance.
- A failure in the network core can be correlated to specific subscribers and services.
- The following policies are applied to various services:
  - → QoS policies
  - → accounting policies
  - → filter policies (IP and MAC)

Additional properties can be configured for bandwidth assignments, class of service, and accounting and billing on the appropriate entity.

Service Entities

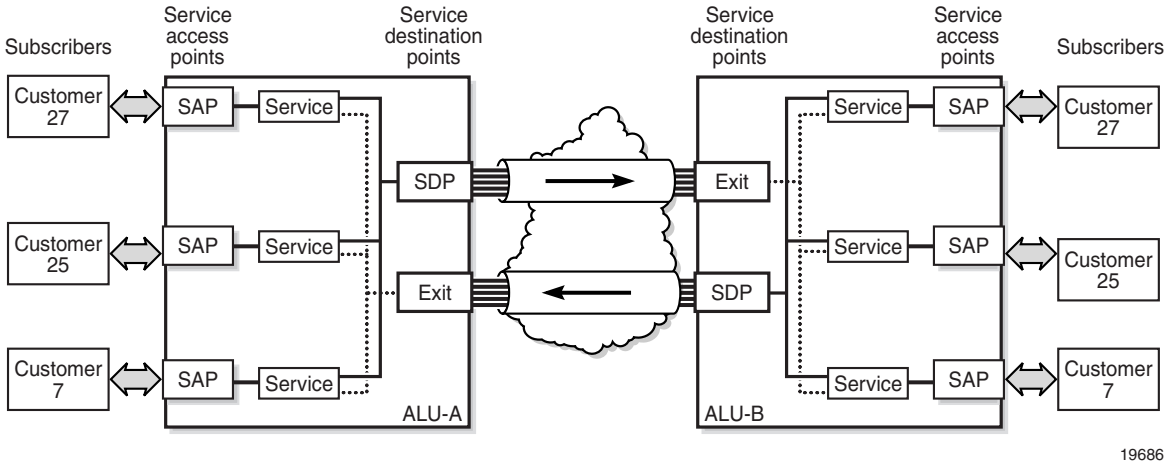
The basic (logical) service entities in the service model used to construct an end-to-end service are:

- Customers
- Service Types
- Service Access Points (SAPs)
- Service Destination Points (SDPs)

Figure 1 shows an example of how the service entities relate to the service model. A subscriber (or customer) attachment circuit connects to a SAP. SDPs define the entrance and exit points of unidirectional service tunnels, which carry one-way traffic between the two routers (ALU-A and ALU-B). After SDPs have been configured, they are bound to a service, which is the final step in making the end-to-end service connection. In Figure 1, the entrance point is labeled SDP and the exit point is labeled Exit.

Traffic encapsulation occurs at the SAP and SDP. The SAP encapsulation types are SONET/SDH, Ethernet, and TDM. The SDP encapsulation types are MPLS, GRE, and IP. For information on SAP encapsulation types, see SAP Encapsulation Types and Identifiers. For information on SDP encapsulation types, see SDP Encapsulation Types.

Figure 1: Service Entities and the Service Model



Customers

The terms customers and subscribers are used synonymously. Every customer account must have a customer ID, which is 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 Types

Service types provide the traffic adaptation needed by customer attachment circuits (ACs). This (logical) service entity adapts customer traffic to service tunnel requirements. The 7705 SAR provides six types of VLL service (that is, point-to-point MPLS-based emulation service, also called Virtual Private Wire Service (VPWS)):

- ATM VLL (Apipe)
- circuit emulation VLL (Cpipe)
- Ethernet VLL (Epipe)
- Frame relay VLL (Fpipe)
- HDLC VLL (Hpipe)
- IP interworking VLL (Ipipe)

The 7705 SAR also provides Ethernet layer (MAC-based) VPLS service (including management VPLS), as well as IP layer VPRN and IES services, that offer any-to-any connectivity within a Virtual Routing Domain or Generic Routing Domain, respectively.

Service Names

A service ID number must be associated with a service at the time of service creation. Once the service is created, an optional service name can be assigned to the service for use by commands that reference the service.

Service Access Points (SAPs)

A service access point (SAP) is the point at which a service begins (ingress) or ends (egress) and represents the access point associated with a service. A SAP may be a physical port or a logical entity within a physical port. For example, a SAP may be a channel group within a DS1 or E1 frame, an ATM endpoint, an Ethernet port, or a VLAN that is identified by an Ethernet port and a VLAN tag. Each subscriber service connection on the 7705 SAR is configured to use only one SAP.

A SAP identifies the customer interface point for a service on a 7705 SAR router. Figure 2 shows one customer connected to two services via two SAPs. The SAP identifiers are 1/1/5 and 1/1/6, which represent the physical ports associated with these SAPs. The physical port information should be configured prior to provisioning a service. Refer to the 7705 SAR OS Interface Configuration Guide for more information about configuring a port. See Port and SAP CLI Identifiers for more information about identifiers.

The 7705 SAR supports the following services types: ATM pseudowires (Apipe), TDM pseudowires (Cpipe), IP pseudowires (Ipipe), Ethernet pseudowires (Epipe), FR pseudowires (Fpipe), HDLC pseudowires (Hpipe), IES, VPLS, and VPRN services. Customer access to these services is provided via SAPs. For each service type, the SAP has slightly different parameters.

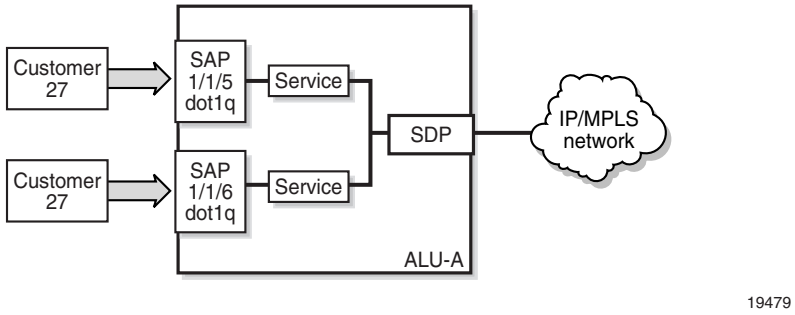
In general, SAPs are logical endpoints that are local to the 7705 SAR and are uniquely identified by:

- the physical Ethernet port, SONET/SDH port, or TDM channel group
- the encapsulation type for the service (for example, ATM)
- the encapsulation identifier (ID), which is, for example, the optional VLAN ID for Epipes, or the channel group ID for Cpipes

Depending on the encapsulation, a physical port or channel can have more than one SAP associated with it (for example, a port may have several circuit groups, where each group has an associated SAP). SAPs can only be created on ports or channels designated as “access” in the physical port configuration.

SAPs cannot be created on ports designated as core-facing “network” ports because these ports have a different set of features enabled in software.

Figure 2: Service Access Point (SAP)



SAP Encapsulation Types and Identifiers

The SAP encapsulation type is an access property of the Ethernet port, SONET/SDH port, or TDM channel group used for the service. It identifies the protocol that is used to provide the service.

The 7705 SAR supports three SAP encapsulation types:

- Ethernet Encapsulations
- SONET/SDH Encapsulations
- TDM and Serial (TDM) Encapsulations

Encapsulation types may have more than one option to choose from. For example, the options for TDM encapsulation type include “cem” (for circuit emulation service) and “atm” (for ATM service), among others.

In the case of SAPs configured on the 16-port T1/E1 ASAP Adapter card version 2, the 32-port T1/E1 ASAP Adapter card, and the 4-port DS3/E3 Adapter card, the cards must be configured to support the appropriate encapsulation methods before the encapsulation type can be configured. This is done using the `mda-mode` command. See the 7705 SAR OS Interface Configuration Guide for more information.

The encapsulation ID is an optional suffix that is appended to a *port-id* to specify a logical sub-element for a SAP. For example, *`port-id:qtag1`* represents a port that can be tagged to use IEEE 802.1Q encapsulation (referred to as dot1q), where each individual tag can identify with an individual service. Similarly, *`port-id.channel-group:vpi/vci`* represents the encapsulation ID for an ATM SAP, which is a special case because it requires that a channel group identifier (which always uses the value 1) precede the VPI/VCI value.

Note:

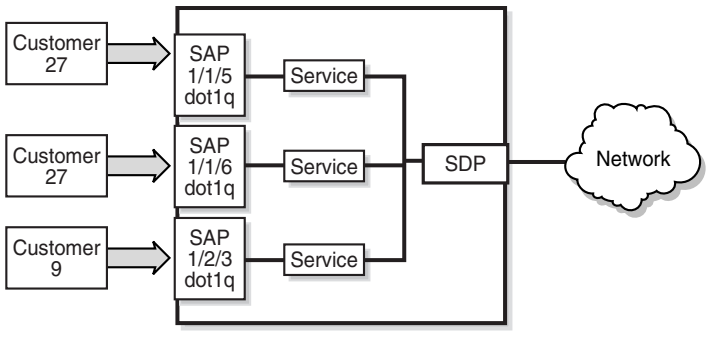
- Throughout this guide, the term “channel group” is often simplified to “channel”.
- Do not confuse the term “encapsulation ID” (described here) with the term “Encapsulation ID”, which is used with the SNMP and MIBs for the 7705 SAR.

Ethernet Encapsulations

The following encapsulation service options are available 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.
- dot1q—supports multiple services for one customer or services for multiple customers (see Figure 3). An example of dot1q use might be the case where the Ethernet port is connected to a multi-tenant unit device with multiple downstream customers. The encapsulation ID used to distinguish an individual service is the VLAN ID in the IEEE 802.1Q header.
- qinq— supports multiple services for one customer or services for multiple customers. The encapsulation IDs used to distinguish an individual service are the QinQ VLAN IDs in the IEEE 802.1ad header, producing a double-tagged frame. For more information on QinQ, see QinQ Support.

Figure 3: Multiple SAPs on a Single Port/Channel



Default SAP on a Dot1q and QinQ Port

The 7705 SAR supports default SAP functionality on dot1q- and qinq-encapsulated ports. On dot1q- and qinq-encapsulated ports where a default SAP is configured, all packets with Q-tags not matching any other explicitly defined SAPs are assigned to the default SAP for transport.

A default SAP is identified in the CLI by the use of the character "*" as a Q-tag, where the "*" means "all". For example, *port-id*:vlan-x.* and *port-id*:*.* are default SAPs. The former case (vlan-x.*) is a specific example of the latter case (*.*), where the outer tag (vlan-x) matches an existing SAP VLAN ID. See Special QinQ SAP Identifiers for more information.

Note: On the 7705 SAR, the *.* notation for QinQ behaves in the same way as the * notation for Dot1q. This behavior is different from the 7750 SR implementation. On the 7750 SR, the "*.*" notation is used only for VPLS service as a capture SAP.

One of the applications where the default SAP feature can be used is for 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 (the use of a whole port) can be provided by a null-encapsulated port. A dedicated VLAN (not used by the user) can be used to provide CPE management.

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

There are a few constraints related to the use of a default SAP on a dot1q- or a qinq-encapsulated port:

- The default SAP is supported only on VPLS and Epipe services and cannot be created in IES and VPRN services because IES and VPRN services 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. By not allowing IGMP snooping of a default SAP, all IGMP packets will be transparently forwarded.
- The default SAP and the SAP defined by explicit null encapsulation are mutually exclusive (for example, 1/1/1:* and 1/1/1:0 are mutually exclusive, and 1/1/2:1.* and 1/1/2:1.0 are mutually exclusive). This avoids conflict as to which SAP untagged frames should be associated with.

SONET/SDH Encapsulations

The following service encapsulation option is available on SONET/SDH ports:

- atm—supports multiple service instances for one customer, as well as bridged llc-snap encapsulated ATM SAP termination to VPLS

TDM and Serial (TDM) Encapsulations

The following service encapsulation options are available on TDM and SDI ports:

- atm—supports multiple services for one customer (TDM ports only)
- cem—supports multiple services for one customer. Structured cem service (circuit emulation service over packet switched network (CESoPSN (n × DS0)) and unstructured cem service (structure-agnostic TDM over packet (SAToP)) are supported. (TDM and SDI ports)
- ipcp—supports a single IP service per TDM channel group on channelized interfaces. This is typically used for router interconnection using the point-to-point protocol (PPP). (TDM ports, and SDI V.35 and X.21 ports at super-rate speeds)
- frame-relay (TDM ports, and SDI V.35 and X.21 ports at super-rate speeds)
- cisco-hdlc (TDM on DS1/E1 ports, and SDI V.35 and X.21 ports at super-rate speeds)
- hdlc (TDM on DS1/E1 ports, and SDI V.35 and X.21 ports at super-rate speeds)

Service Types and SAP Encapsulations—Summary

Table 3 lists the SAP encapsulations available to 7705 SAR service types. These encapsulations apply to access-facing ports. The service (port) type and encapsulations are configured at the port level. See the 7705 SAR OS Interface Configuration Guide for more information about the cards and ports that support each of the service types.

Table 3: Service Types and SAP Encapsulations

Service (Port) Type	Encapsulation Option
Ethernet	null
Ethernet	dot1q
Ethernet	qinq
SONET/SDH	atm

Table 3: Service Types and SAP Encapsulations (Continued)

Service (Port) Type	Encapsulation Option
TDM	cem
TDM	atm
TDM	ipcp
TDM	frame-relay
TDM	cisco-hdlc
TDM	hdlc

SAP Configuration Considerations

In addition to being an entry or exit point for service traffic, a SAP has to be configured for a service and, therefore, has properties. When configuring a SAP, consider the following.

- A SAP is a local entity and is only locally unique to a given device. The same SAP ID value can be used on another 7705 SAR.
- There are no default SAPs. All subscriber service SAPs 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 are also deleted.
- A SAP is owned by and associated with the service in which it is created.
- An Ethernet port or channel with a dot1q encapsulation type means that 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 is 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. QinQ encapsulation means that the SAP is identified based on specific IEEE 802.1ad VLAN ID values.
- A TDM circuit emulation service (for example, CESoPSN) requires a channel group. The channel group must be created before it can be assigned to a SAP.
- An ATM service (for example, ATM N-to-1 VCC cell transport) on a 16-port T1/E1 ASAP Adapter card, 2-port OC3/STM1 Channelized Adapter card, or a 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card requires a channel group. For this case, the channel group requires the assignment of all 24 timeslots (T1) or 30 timeslots (E1). The timeslot assignments are made automatically after a channel group is configured for ATM encapsulation.

- If a port or channel is administratively shut down, all SAPs on that port or channel will be operationally out of service.
- A SAP cannot be deleted until it has been administratively disabled (shut down).
- Each SAP can have one of the following policies assigned to it:
  - → Ingress QoS policy
  - → Egress QoS policy
  - → Accounting policy
  - → Ingress filter policy (for Epipe SAPs, Ipipe SAPs, VPLS SAPs, VPRN SAPs, IES SAPs, and IES in-band management SAPs)
  - → Egress filter policy (for VPRN and IES SAPs, and for VPLS SAPs (Ethernet SAPs only))

Service Destination Points (SDPs)

An SDP identifies the endpoint of a logical unidirectional service tunnel. The service tunnel provides a path from one 7705 SAR to another network device, such as another 7705 SAR, a 7710 SR, or a 7750 SR.

In more general terms, SDP refers to the service tunnel itself. The SDP terminates at the far-end router, which is responsible for directing the flow of packets to the correct service egress SAPs on that device.



Note: In this document and in command line interface (CLI) usage, SDP is defined as Service Destination Point. However, it is not uncommon to find the term SDP defined in several different ways, as in the following list. All variations of SDP have the same meaning:

- Service Destination Point
- Service Distribution Point
- Service Destination Path
- Service Distribution Path
- Service Delivery Path

When an SDP is bound to a service, the service is referred to as a distributed service. 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 that binds the service to the service tunnel.

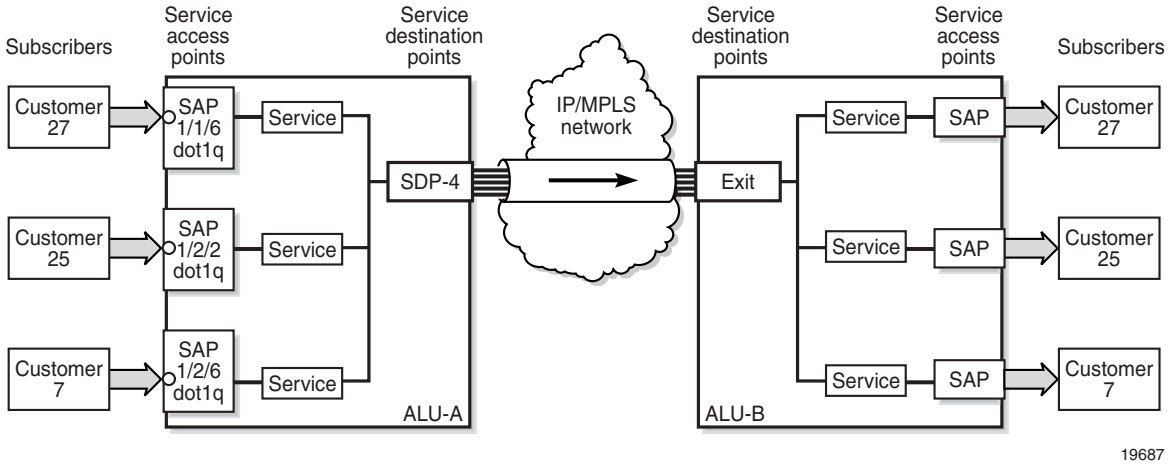
An SDP has the following characteristics.

- An SDP is locally unique to a participating 7705 SAR. The same SDP ID can appear on other 7705 SAR routers.
- An SDP uses the system IP address of the far-end edge router to locate its destination.
- An SDP is not specific to any one service or to 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 bound to an SDP use the same SDP (transport) encapsulation type defined for the SDP (GRE, IP, or MPLS).
- An SDP is a service entity used for service management. Even though the SDP configuration and the services carried within it 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 tunnel from the local device (typically, a 7705 SAR) to the far-end device (router) requires a return SDP tunnel from the far end back to the local device. Each device must have an SDP defined for every remote router to which it wants to provide service. The SDP must be created before a distributed service can be configured.
- An SDP can be used to provide PW redundancy, where up to four spoke SDPs can be assigned to a service endpoint that acts as the managing entity to ensure service connection. See Pseudowire Redundancy.

SDP Binding

To configure a distributed service pointing from ALU-A to ALU-B, the SDP ID on the ALU-A side (see Figure 4) must be specified during service creation in order to bind the service to the tunnel (the SDP). Otherwise, service traffic is not directed to a far-end point and the far-end 7705 SAR device(s) cannot participate in the service (there is no service). To configure a distributed service pointing from ALU-B to ALU-A, the SDP ID on the ALU-B side must be specified.

Figure 4: SDP Tunnel Pointing from ALU-A to ALU-B



Spoke and Mesh SDPs

There are two types of SDPs: spoke and mesh. The type of SDP defines how flooded traffic (or broadcast traffic, such as an ARP request) is propagated. For point-to-point PW/VLL services, spoke SDPs are the only way to bind services to the far-end router. For VPLS, mesh and spoke SDP bindings are allowed.

A spoke SDP that is bound to a service operates like a traditional bridge port. Flooded traffic that is received on the spoke SDP is transmitted to all the spoke SDPs, mesh SDPs, and SAPs to which it is connected. Flooded traffic is not transmitted back toward the port from which it was received.

In contrast, a mesh SDP that is bound to a service operates like a single bridge port. Flooded traffic received on a mesh SDP is transmitted to all spoke SDPs and SAPs to which it is connected. Flooded traffic is not transmitted to any other mesh SDPs or back toward the port from which it was received. This property of mesh SDPs is important for multi-node networks; mesh SDPs are used to prevent the creation of routing loops.

SDPs and BGP Route Tunnels

SDP can use BGP route tunnels to extend inter-AS support for Layer 2 and Layer 3 VPN services as defined in RFC 3107. An SDP can be configured based on service transport method (for example, GRE or MPLS tunnel). 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 LSP, RSVP-TE LSP or BGP route tunnel).

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

Only BGP route labels can be used to transition from an ASBR to the next-hop ASBR. The global BGP route tunnel transport configuration option must be configured to select an LSP to reach the PE node from the ASBR.

For more information on BGP route tunnels, refer to the 7705 SAR OS Routing Protocols Guide, “BGP Route Tunnel”.

SDP Encapsulation Types

The Alcatel-Lucent service model uses encapsulation tunnels (also referred to as service tunnels) through the core to interconnect 7705 SAR and SR routers. An SDP is a logical way of referencing the entrance to an encapsulation tunnel.

The following encapsulation types are supported:

- Layer 2 within multiprotocol label switching (MPLS Encapsulation)
- Layer 2 or Layer 3 within generic routing encapsulation (GRE Encapsulation)
- Layer 2 within IP (IP Encapsulation)

Each SDP service tunnel has an entrance and an exit point for the pseudowires contained within it.

MPLS Encapsulation

Multiprotocol label switching (MPLS) encapsulation has the following characteristics.

- An MPLS 7705 SAR router supports both signaled and non-signaled LSPs through the network.
- Non-signaled paths are defined at each hop through the network.

An SDP has an implicit Maximum Transmission Unit (MTU) value because services are carried in encapsulation tunnels and an SDP is an entrance to the tunnel. The MTU is configurable (in octets), where the transmitted frame can be no larger than the MTU.

With MPLS, the MTU for the network port permits the addition of labels for transmission across the MPLS network. Ethernet frames that are sent out of a network port toward the MPLS core network (or a P router) are allowed to be oversized in order to include the MPLS labels without the need to fragment large frames. See MTU Settings for more information.

The following ways of configuring an MPLS tunnel are supported:

- LDP signaled
- RSVP-TE signaled
- user-configured (static LSP)

GRE Encapsulation

Generic routing encapsulation (GRE) is one of the most common tunneling techniques in the industry. GRE tunnels are used to transport various network layer packets and are especially useful for facilitating pseudowires over IP networks. Since MPLS is a Layer 2.5 protocol, MPLS packets cannot be natively transported over a Layer 3 (IP) network. Therefore, GRE is the ideal alternative for applications where traffic must travel over a Layer 3 network; for example, in DSL applications.

For the HSDPA offload application (see HSDPA Offload), ATM pseudowires are transported over IP using GRE tunneling. For other applications, Ethernet and TDM pseudowires over GRE are also supported.

GRE SDPs are supported on all network interfaces.

GRE Format

In accordance with RFC 2784, a GRE encapsulated packet has the following format:

- delivery header
- GRE header
- payload packet

Delivery Header

The delivery header is always an IP header.

GRE Header

The GRE header format is shown in Figure 5 and described in Table 4.

Figure 5: GRE Header

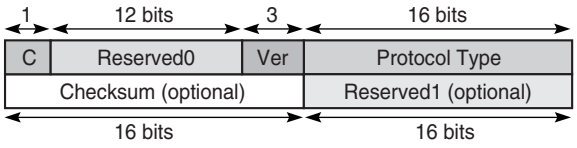


Table 4: GRE Header Descriptions

Field	Description
C	Specifies whether there is a checksum in the header If set to 1, both the checksum and reserved1 fields must be present On the 7705 SAR, in the network egress (transmit) direction, the C bit is always set to 0; therefore, the checksum and reserved1 fields are omitted from the header. The GRE header is therefore always 4 bytes (32 bits) in the network egress direction. In the network ingress direction, the C bit validity is checked. If it is set to a non-zero value, the GRE packet is discarded and the IP discards counter is increased.

Table 4: GRE Header Descriptions (Continued)

Field	Description
Reserved0	Indicates whether the header contains optional fields Not applicable to the 7705 SAR—first 5 bits of the field are always set to 0 and bits 6 to 12 are reserved for future use and also set to 0 by the 7705 SAR
Ver	Always set to 000 for GRE At network ingress, if a GRE packet is received with the version field set to any value other than 000, the packet is discarded and the IP discards counter is increased
Protocol Type	Specifies the protocol type of the original payload packet—identical to Ethertype with the only supported option being MPLS unicast (0x8847)
Checksum (optional)	Not applicable
Reserved1 (optional)	Not applicable

Payload Packet

The payload encapsulation format for pseudowires over GRE is shown in Figure 6 and described in Table 5.

Figure 6: GRE Pseudowire Payload Packet over Ethernet

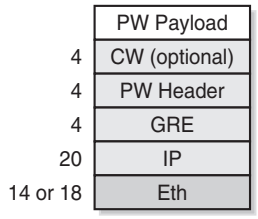


Table 5: GRE Pseudowire Payload Packet Descriptions

Field	Description
Eth	The Layer 2 transport header The only Layer 2 protocol supported is Ethernet MTU size depends on the encapsulation type (14 bytes for null encapsulation, 18 bytes for dot1q encapsulation, and 22 bytes for qinq encapsulation)
IP	Indicates the transport protocol The Ethertype is always set to IP (0x800), and in case of a mismatch, the unexpected or illegal Ethertype counters are increased [1]
GRE	Indicates the encapsulation protocol
PW header	The pseudowire header identifies a service within the GRE tunnel
CW (optional)	The pseudowire control word (CW) is a 32-bit (4-byte) field that is inserted between the VC label and the Layer 2 frame For more information on the control word, see Pseudowire Control Word
PW payload	The PW payload is the payload of the service being encapsulated (Ethernet, ATM, or TDM)

Note:

1. The only exception to the Ethertype is if the packets are address resolution protocol (ARP) packets. For information on ARP, refer to the 7705 SAR OS Router Configuration Guide.

When using GRE, the service MTU might have to be set to a value smaller than 2102 octets. For more information on MTU, see MTU Settings.

At the network egress of the 7705 SAR, the source address of the IP header is always set to the system IP address. The destination IP address is set to the system IP address of the service router on which the GRE SDP is configured. Using the system IP addresses to bring up the GRE session ensures that any IP link between the two routers can be used to transport GRE/IP packets. It might therefore be necessary to use static IP address configuration over DSL networks to ensure connectivity between the routers (especially if the DSL modem is in bridge mode).

IP Encapsulation

IP encapsulation is added to the 7705 SAR in response to a growing demand for more pseudowire-based solutions in mobile backhaul. IP encapsulation is similar to GRE encapsulation but allows pseudowires to be transported natively over IP packets. Only static pseudowires are supported for IP SDPs, because there is no label path to define except for the endpoints. The path is an IP routed path.

Since Release 3.0, the 7705 SAR has supported the transport of pseudowires over IP tunnels. All pseudowire types supported since Release 3.0 can be transported over IP tunnels. Figure 7 shows an example of an application using Apipes over IP over Ethernet.

Payload Packet

A typical payload encapsulation format for pseudowires over IP is shown in Figure 7 and described in Table 6.

Figure 7: IP Example of Pseudowire Payload Packet over Ethernet

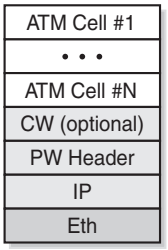


Table 6: IP Pseudowire Payload Packet Descriptions

Field	Description
Eth	The Layer 2 transport header The only Layer 2 protocol supported is Ethernet MTU size depends on the encapsulation type (14 bytes for null encapsulation, 18 bytes for dot1q encapsulation, and 22 bytes for qinq encapsulation)
IP	Indicates the transport protocol The only supported option is MPLS in IP (0x89) The Ethertype is always set to IP (0x0800), and in case of a mismatch, the unexpected or illegal Ethertype counters are increased [1]

Table 6: IP Pseudowire Payload Packet Descriptions (Continued)

Field	Description
PW header	The pseudowire header identifies a service within the IP tunnel. The pseudowire header is like an MPLS header that has context only to the encapsulating and decapsulating LERs. This means that the IP transport network has no knowledge about the pseudowires that it carries. Only the edge LERs are aware of the pseudowire because the IPv4 Protocol Number field is set to 137 (0x89), indicating an MPLS unicast packet.
CW (optional)	The pseudowire control word (CW) is a 32-bit (4-byte) field that is inserted between the VC label and the Layer 2 frame For more information on control word, see Pseudowire Control Word
ATM Cell #1 to ATM Cell #N	Indicates the payload of the service being encapsulated (ATM)

Note:

1. The only exception to the Ethertype is if the packets are address resolution protocol (ARP) packets. For information on ARP, refer to the 7705 SAR OS Router Configuration Guide.

Spoke SDP Terminations

The 7705 SAR supports spoke SDP as termination points for IES and VPRN services. Table 7 shows which service interfaces and spoke SDPs can be connected to each other. For example, an Epipe spoke SDP can connect to an IES or VPRN interface. Refer to Spoke SDP Termination to IES and to Spoke SDP Termination to VPRN for more information.

Table 7: Spoke SDP Termination Support

	Epipe Spoke SDP	Epipe Spoke SDP Redundancy (standby-signal-master enabled)	IES Interface	VPRN Interface	VPLS Spoke SDP	VPLS Spoke SDP Redundancy (suppress-standby-signaling disabled)
Epipe Spoke SDP	✓	✓	✓	✓	✓	✓
Epipe Spoke SDP Redundancy (standby-signal-master enabled)	✓	✓	✓	✓	✓	✓
IES Interface	✓	✓	✓	✓	✓	✓
VPRN Interface	✓	✓	✓	✓	✓	✓
VPLS Spoke SDP	✓	✓	✓	✓	✓	✓
VPLS Spoke SDP Redundancy (suppress-standby-signaling disabled)	✓	✓	✓	✓	✓	✓

SDP Ping

Ping is an application that allows a user to test whether a particular host is reachable. SDP Ping is an application that allows a user to test whether a particular SDP endpoint is reachable.

SDP ping uses the SDP identifier that is stored in the 7705 SAR that originates the ping request. SDP ping responses can be configured to return through the corresponding return tunnel as a round-trip ping, or out-of band when unidirectional pings are requested. Refer to the 7705 SAR OS OAM and Diagnostics Guide, “SDP Ping”, for more information.

SDP Keepalives

The SDP keepalive application allows a system operator to actively monitor the SDP operational state using periodic Alcatel-Lucent SDP Echo Request and Echo Reply messages. Automatic SDP keepalives work in a manner that is similar to a manual SDP ping command. The SDP Echo Request and Echo Reply messages provide a mechanism for exchanging far-end SDP statuses.

SDP keepalive Echo Request messages are only sent after the SDP has been completely configured and is administratively up and the SDP keepalives are 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.

The SDP is immediately brought operationally down when:

- the Max Drop Count Echo Request messages do not receive an Echo Reply
- a keepalive response is received that indicates an error condition

After a response is received that indicates the error has cleared and the Hold Down Time interval has expired, the SDP is eligible to be put into the operationally up state. If no other condition prevents the operational change, the SDP enters the operational state.

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

- Hello Time
- Message Length
- Max Drop Count
- Hold Down Time
- Timeout

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

Mobile Solutions

Topics in this section include:

- HSDPA Offload
- Pure DSL Backhaul

The Mobile Radio Access Network (RAN) is rapidly growing to meet the increased demand in mobile services. This in turn increases demands on carriers to provide high-bandwidth, mobile broadband services. Today, at a typical cell site, 2G and 3G base stations are connected to high-cost, T1/E1 leased lines that are used to backhaul both voice and data traffic to the MTSO. For mission-critical, delay-sensitive, and low-bandwidth traffic such as voice, signaling, and synchronization traffic, it is vital that the high availability of these leased lines is ensured. SLA agreements also promise a high level of availability for customers.

Currently, however, best-effort traffic such as high-speed downlink packet access (HSDPA) is also switched over these SLA-enabled leased lines. HSDPA is a 3G mobile telephony communications service that allows UMTS networks to have higher data transfer speeds and capacity, allowing the mobile customer (end user) to browse the Internet or to use the mobile device. The increasing use of HSDPA is having a dramatic impact on the ability of the T1/E1 leased lines to scale with the traffic growth as well as on the operating costs of these lines.

Similar issues confront CDMA EVDO networks today.

Alcatel-Lucent provides a solution that enables mobile operators to keep their existing infrastructure (circuit-based leased lines), while gradually migrating to a packet-based infrastructure that will allow scalability, decrease costs, and ease the transition to the next-generation, all-IP network solutions.

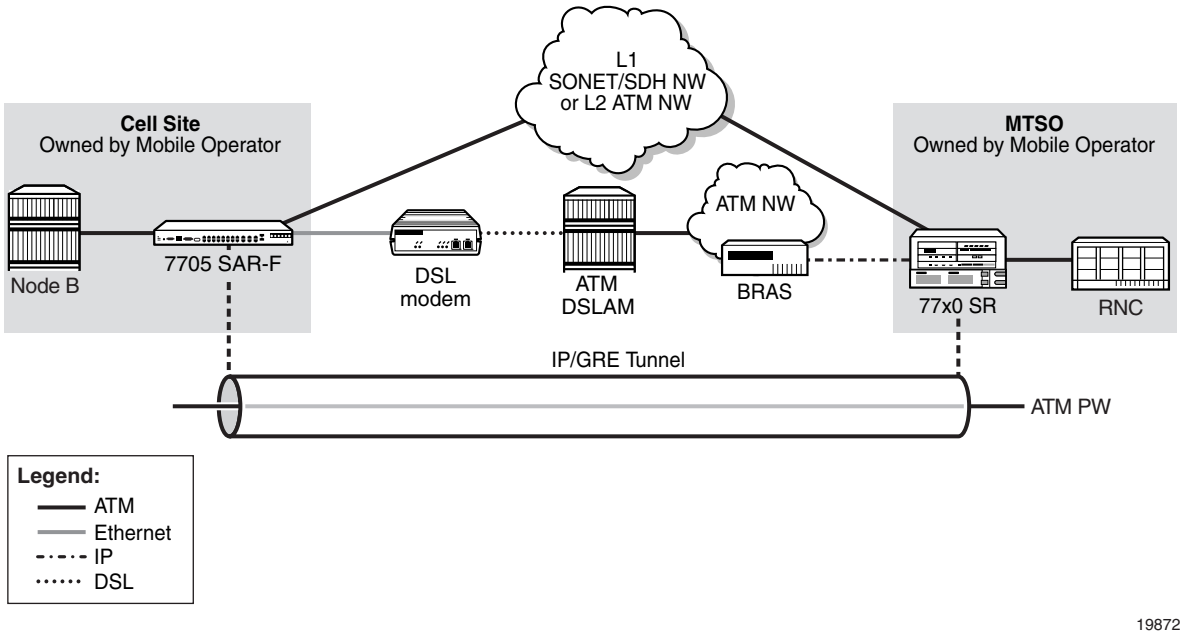
HSDPA Offload

The Alcatel-Lucent solution is to make use of widely available DSL networks and split the traffic being backhauled. Mission-critical traffic (voice, signaling, synchronization) remains on the T1/E1 leased line circuits, while the best-effort, bandwidth-hungry HSDPA traffic is off-loaded to DSL networks.

The 7705 SAR-F and 7705 SAR-M with DSL modules are ideal candidates for this scenario. They are small-scale versions of the 7705 SAR product family, optimized for use in standalone small or midsized sites where traffic aggregation from multiple cell sites is not needed. For more information on the 7705 SAR-F and 7705 SAR-M, refer to the 7705 SAR-F Chassis Installation Guide and 7705 SAR-M Chassis Installation Guide.

Figure 8 shows an example of HSDPA offload with the 7705 SAR-F.

Figure 8: 7705 SAR-F HSDPA Offload Example



A 3G Node B is connected to a 7705 SAR over an ATM/IMA access port (SAP endpoint). An ATM SAP-to-SAP connection is set up in the 7705 SAR and a pseudowire is configured between the two endpoints to emulate local ATM switching. Traffic from the Node B enters an ATM/IMA port, the VCs transporting mission-critical traffic are locally switched (SAP-to-SAP) to another ATM/IMA port (SAP endpoint), and then switched over the leased lines to the MTSO.

Note: ATM SAP-to-SAP connections are supported between any T1/E1 ASAP port that is in access mode with ATM/IMA encapsulation and another port with the same encapsulation configuration. One endpoint of a SAP connection can be an IMA group, while the other endpoint can be on a single ATM port. ATM SAP-to-SAP connections are also supported between any two OC3/STM1 ports and between any T1/E1 ASAP port and OC3/STM1 port, as long as both SAPs support ATM.

For non-mission-critical traffic, for example, HSDPA traffic, an Ethernet interface on the 7705 SAR is connected to an external DSL modem. HSDPA traffic is interworked to ATM pseudowires and transported over the DSL network to the BRAS, then forwarded to the service router at the MTSO.

Figure 9 shows an example of HSDPA offload with the 7705 SAR-M using an 8-port xDSL module. Figure 10 shows an example of HSDPA offload with the 7705 SAR-M using a 6-port DSL Combination module.

Figure 9: 7705 SAR-M With 8-port xDSL Module for HSDPA Offload

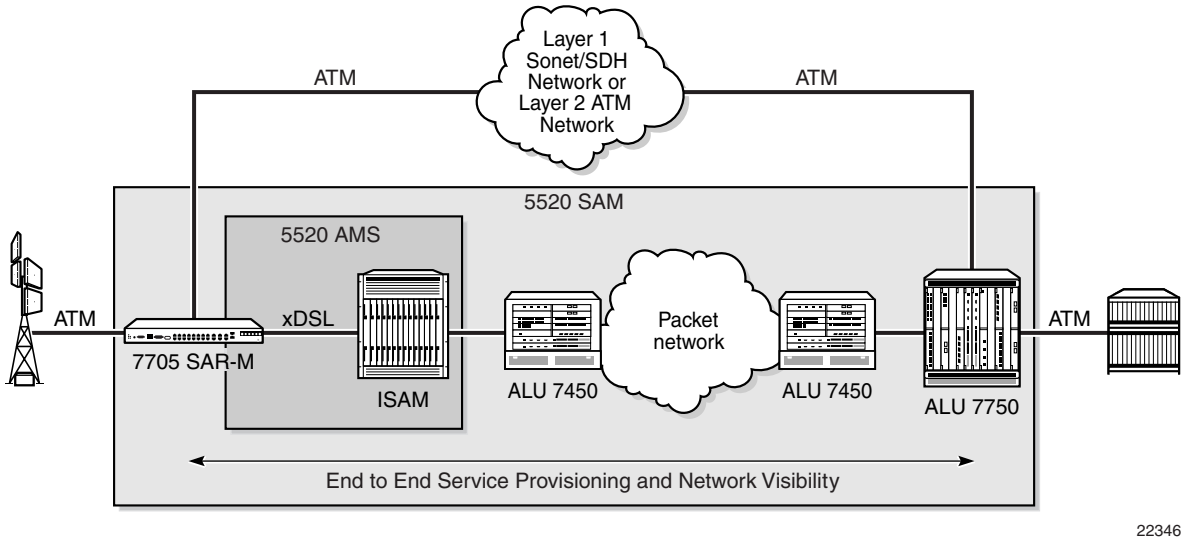
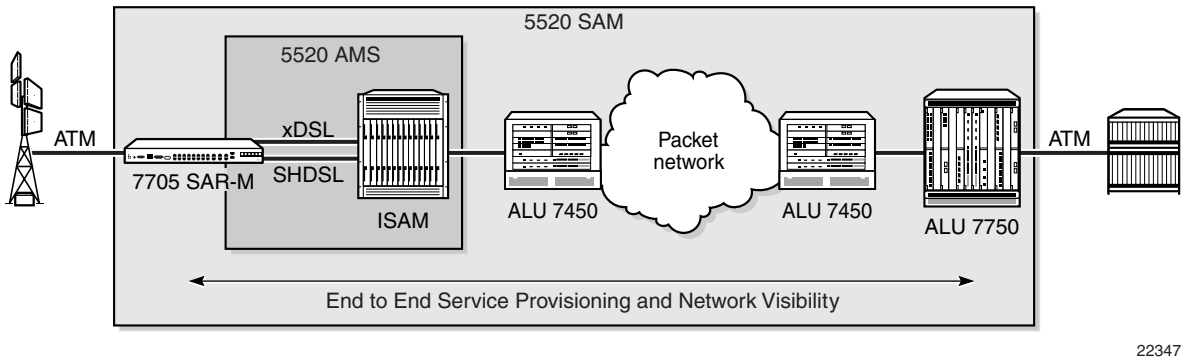


Figure 10: 7705 SAR-M With 6-port DSL Combination Module for HSDPA Offload



On a 7705 SAR-M with an 8-port xDSL module, non-real-time UMTS traffic is backhauled over the xDSL connection and all voice traffic remains on the existing attached TDM network. For this type of network, timing recovery via DSL is typically not required as synchronization can be derived from the attached TDM network.

On a 7705 SAR-M with a 6-port DSL Combination module, non-real-time UMTS traffic is backhauled over the xDSL connection and all voice traffic is backhauled over the SHDSL interface. Using SHDSL rather than existing ATM or TDM networks as the primary means to backhaul voice and signaling is an additional step toward improving the economies of backhaul. With this functionality, all voice, signaling, and HSDPA offload can be backhauled over a low-cost DSL network. The network characteristics of SHDSL are based on the design considerations for T1/E1 lines and are subject to lower latency and delay compared with xDSL. Clock recovery with NTR is supported over both SHDSL and xDSL interfaces on the 6-port DSL Combination module and over any xDSL interface on the 8-port xDSL module.

Failure Detection

Failure of the GRE SDP or the IP network it rides over can be detected by OAM tools as well as by BFD. With SAA, OAM tools can be configured to run periodically in order to facilitate faster failure detection. If a failure occurs, the ATM SAPs must be rerouted by the 5620 SAM to the ATM ports used for backhauling the traffic. The mission-critical traffic is still serviced before the best-effort HSDPA traffic.

For information on OAM and SAA tools, refer to the 7705 SAR OS OAM and Diagnostics Guide. For information on BFD, refer to the 7705 SAR OS Router Configuration Guide.

Pure DSL Backhaul

There are two applications for pure backhaul over DSL:

- wholesale mobile backhaul aggregation or business services provided via DSL
- cell site aggregation from one or more base stations

Timing and/or synchronization are typically not requirements for deployments of the first application since mobile operators have their own cell site aggregator for timing recovery. Figure 11 shows a typical network using pure DSL backhaul with a 7705 SAR-M. Pure DSL backhaul is also supported on networks with distributed MPLS services.

Figure 11: Pure DSL Backhaul Network Example

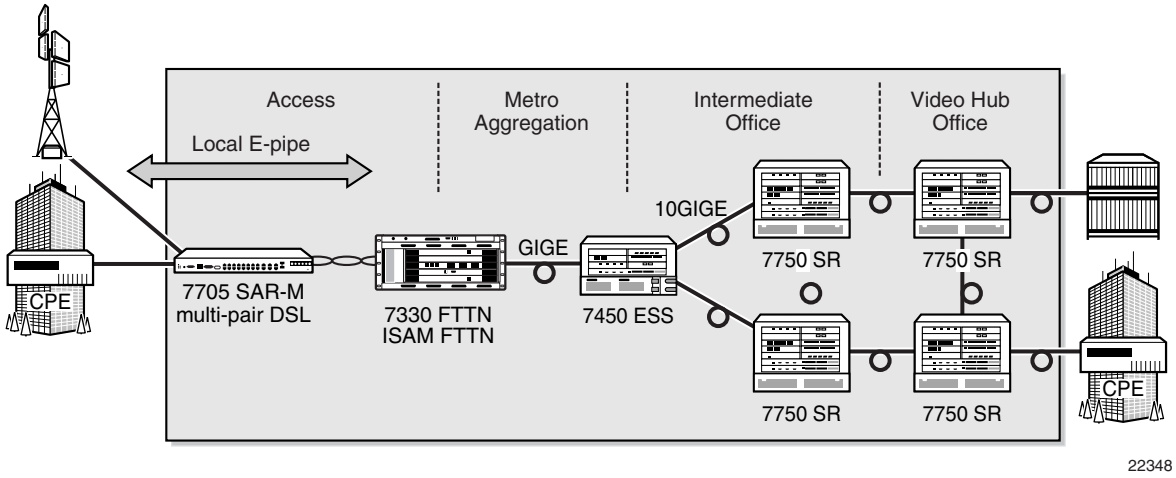
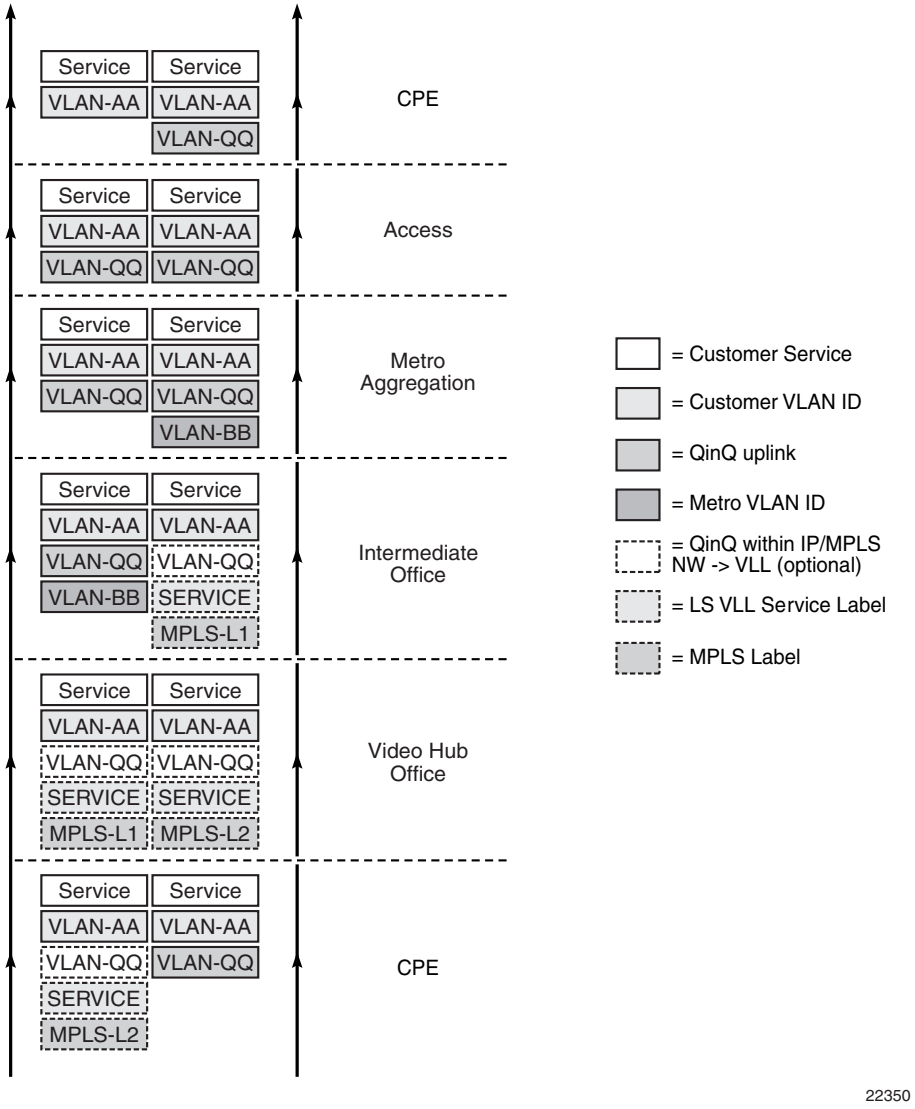


Figure 12 shows a breakdown of the traffic on this network.

Figure 12: ‘‘Sample DSL Network Traffic



In the case of pure backhaul via DSL for a mobile operator, legacy 2G and 3G traffic is aggregated at the 7705 SAR-M then fed upstream to a connected ISAM. Services remain terminated on a 7750 SR as shown in Figure 11, but without the north ATM path. Timing recovery is provided by a local GPS at the cell site, by NTR, or by 1588 for frequency synchronization.

Note: Using 1588 for timing recovery over DSL requires careful consideration of the end-to-end network characteristics for 1588v2 PTP delivery, and for the underlying PDV characteristics of the DSL used to transport the PTP packets.

Pure DSL backhaul maximizes savings for mobile operators since leased lines are not used at all. SHDSL can also be more expensive to deploy and operate because it is generally not deployed for residential broadband services. Additionally, when mobile backhaul services are combined with triple-play services reusing the same broadband platform (in this case the ISAM), the pure DSL backhaul architecture offers a compelling, cost-effective solution across a wide range of scenarios.

The pure DSL backhaul deployment model may be Layer 2-based with the use of a local Epipe on the 7705 SAR-M as shown in Figure 11. The SAP facing the ISAM must support an Epipe SAP and an IES SAP for in-band management. Alternatively, distributed services via MPLS or GRE tunnels can be used with pure DSL backhaul to offer distributed VPLS, VPRN, IES, or VLL services.

DSL Channel Bonding

DSL channel bonding provides an ideal mix of features, including higher bandwidth for all users and the ability to extend the distance that can be reached at certain bandwidths. DSL bonding distributes traffic over a bundle of copper pairs instead of just a single copper pair. For example, to achieve an effective bandwidth of 12 Mb/s, three DSL lines of 4 Mb/s are bundled with a channel bonding processor at each of the Central Office and CPE devices.

In order to meet the growing needs of mobile backhaul and business services in areas where fiber connectivity is not available, DSL access solutions can be used as a viable alternative through the use of channel bonding for increased rate and reach. For more information about DSL bonding and DSL modules, see the 7705 SAR OS Interface Configuration Guide and the 7705 SAR-M Chassis Installation Guide.

ETH-CFM (802.1ag and Y.1731)

Topics in this section include:

- 802.1ag and Y.1731 Terminology
- ETH-CFM Frame Format
- ETH-CFM Functions and Tests
- MEP Support (802.1ag and Y.1731)
- Priority Mapping (802.1ag and Y.1731)

Ethernet Connectivity Fault Management (ETH-CFM) is defined in two complementary standards: IEEE 802.1ag (dot1ag) and ITU-T Y.1731. Both standards specify protocols, procedures, and managed objects in support of transport fault management, including discovery and verification of the path, and detection and isolation of a connectivity fault for each Ethernet service instance.

Dot1ag and Y.1731 provide fault management (FM) functions for loopback, linktrace, and connectivity checks, as well as Up and Down MEP support for Ethernet SAPs, Down MEP support for spoke SDPs (dot1ag only), and facility MEP support for network interfaces.

Y.1731 fault management (Y.1731 FM) extends dot1ag CFM by providing functions for alarm indication signal (AIS) and ETH-Test testing. Furthermore, Y.1731 provides performance monitoring (Y.1731 PM) functions for delay and loss measurements. For more information on Y.1731 PM, refer to the “ITU-T Y.1731 Performance Monitoring (PM)” section in the 7705 SAR OS OAM and Diagnostics Guide.

For information on running Ethernet OAM tests, refer to the “ETH-CFM (802.1ag and Y.1731)” section in the 7705 SAR OS OAM and Diagnostics Guide.

The information in this section is specific to Ethernet SAPs and spoke SDPs, although most of it also applies to Ethernet network interfaces. For information on ETH-CFM support specific to network interfaces, refer to the 7705 SAR OS Router Configuration Guide, “ETH-CFM Support”.

CFM uses Ethernet frames that are distinguished by their Ethertype value and special Ethernet multicast address. For more information on the Ethernet frame, and the Ethertype and Ethernet multicast address values, see ETH-CFM Frame Format.

Using CFM, interoperability can be achieved between different vendor equipment in the service provider network, up to and including customer premises bridges.

Note: In the 7705 SAR CLI command hierarchy, commands for 802.1ag and Y.1731 are found under the `eth-cfm` context that appears at the following levels:

- global (`config>eth-cfm`)
- Epipe service (`config>service>epipe>sap>eth-cfm`)
- network interface (`config>router>if>eth-cfm`)
- show (`show>eth-cfm`)
- oam (`oam eth-cfm`)

802.1ag and Y.1731 Terminology

Table 8 defines 802.1ag terms. Table 9 illustrates the similarities and differences between Y.1731 and 802.1ag terms.

Table 8: 802.1ag Terminology

Term	Expansion	Definition
MA	Maintenance Association	A grouping of maintenance entities (MEs) that need to be managed as part of a given service Typically, on the 7705 SAR, the MEs are a pair of MEPs (one local and one remote)
MA-ID	Maintenance Association Identifier	A unique combination of *md-index*, MD *level* and *ma-index*, where *md-index, level,* and *ma-index* are user-configured values An MA is identified by its MA-ID
MD	Maintenance Domain	A set of Ethernet network elements or ports that are controlled by an operator, where boundaries are set by MEPs
MD level	Maintenance Domain level	A user-configured value of 0 to 7 representing a level of hierarchy within a CFM architecture. The value 7 is the highest MD level and 0 is the lowest. The MD level is transmitted as part of the Ethernet CFM frame. A CFM message is said to have a higher MD level when its MD level value is higher than the MD value configured on the receiving MEP 7705 SAR. Higher-level CFM messages are relayed as data frames by MEPs and ignored by the MEP entity
ME	Maintenance Entity	An Ethernet port or endpoint that is managed as part of dot1ag OAM An endpoint can be a SAP or a spoke SDP

Table 8: 802.1ag Terminology (Continued)

Term	Expansion	Definition
MEP	Maintenance Association End Point	An (edge) endpoint that can terminate, respond to, or initiate the OAM messages for a configured MD-MA combination
MEP-ID	Maintenance Association End Point Identifier	A MEP is identified by its MEP-ID, which is a unique combination of *md-index* and *ma-index*, where *md-index* and *ma-index* are user-configured values
MIP	Maintenance Association Intermediate Point	An intermediate point that can respond to OAM messages initiated by MEPs in the same MD. Connectivity fault management (CFM) messages destined for other MIPs or the destination MEP are transparent to MIPs. MIPs are not supported on the 7705 SAR

Table 9: Y.1731 Terminology

Term	Expansion	Definition
MEG	Maintenance Entity Group	Same as MA but applies to Y.1731
MEG-ID	Maintenance Entity Group Identifier	Same as MA-ID but applies to Y.1731
MEG level	Maintenance Entity Group Level	Same as MD level but applies to Y.1731 Note that although MEG level and MD level are equivalent terms, there is no Y.1731 equivalent to an MD
MEP	Maintenance Association End Point	Same as MEP for 802.1ag

MDs, MD Levels, MAs, and MEPs (802.1ag)

Maintenance domains (MDs) and maintenance associations (MAs) are configured at the global level. Maintenance association endpoints (MEPs) are configured at the service level.

An MD is a set of network elements that have a common CFM OAM purpose. MDs are identified by their MD index and can be given an MD name. An MD is assigned a maintenance domain level (MD level). There are eight MD levels. Use MD levels to set up a messaging hierarchy for the CFM architecture.

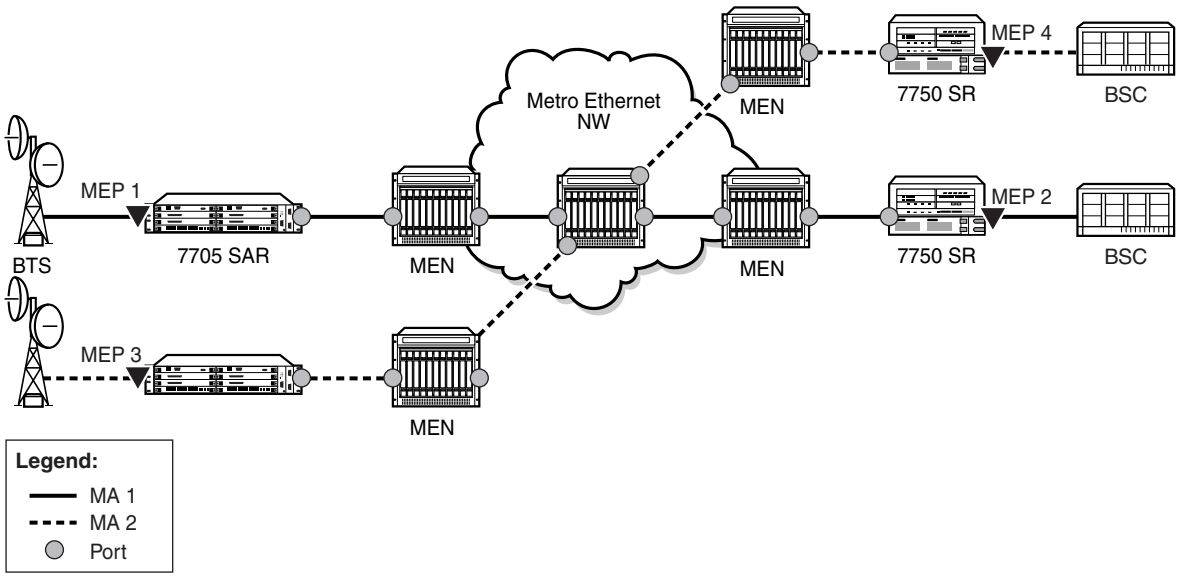
An MA consists of a pair of maintenance endpoints (one local and one remote MEP) that are managed as part of an Ethernet VLL service. The MA and the VLL service are associated by configuring the MA bridge identifier parameter to have the same value as the VLL service ID parameter of the VLL service that supports the MEPs. MAs are identified by their MA index and can be given an MA name. The MA is used to verify the integrity of a single service instance.

A MEP is configured as part of an Ethernet SAP or spoke SDP. MEPs can generate or terminate CFM OAM messages. MEPs only communicate within the same MD level, where the value of the MD level (0 to 7) is carried in a CFM OAMPDU. MEPs are identified by their MEP identifier and MA-ID. The MA-ID is configured at the global level.

Figure 13 depicts a high-level view of MEPs in a CFM-enabled network. Two MAs are shown. The endpoints of MA 1 are MEPs 1 and 2, while MEPs 3 and 4 are the endpoints for MA 2.

For more information on MEP support, see Ethernet OAM.

Figure 13: MEPs and MAs



MEG Levels, MEGs, and MEPs (Y.1731)

On the 7705 SAR, the implementation of Y.1731 Fault Management (FM) is similar to that of dot1ag CFM, except that Y.1731 does not have a maintenance domain (MD). For Y.1731 and 802.1ag, the following terms are equivalent:

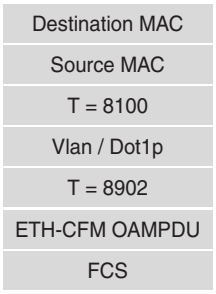
- MEG level is equivalent to MD level
- MEG is equivalent to MA
- a Y.1731 MEP is equivalent to a dot1ag MEP

To access Y.1731 functions, including Y.1731 Performance Monitoring (PM) functions, configure a MEP to have the domain format set to **none** and the association format set to **icc-based** or **string** (the **string** keyword enables the Y.1731 MEP to interoperate with a dot1ag MEP).

ETH-CFM Frame Format

ETH-CFM OAMPDU messages for 802.1ag and Y.1731 use a standard Ethernet frame (see Figure 14). The parts of the frame are described below.

Figure 14: ETH-CFM Frame Format



Destination and Source Addresses

The destination and source MAC addresses of the CFM message must match at the send and the receive routers. For example, a 7705 SAR-initiated ETH-CFM message would use the spoke SDP MAC address of the 7705 SAR as the source MAC address and the spoke SDP MAC address of the far-end router as the destination MAC address. At the far end, the source and destination MAC addresses would be the reverse of the near end.

An exception to the matching source-destination MAC address requirement occurs for linktrace and continuity messages, where the destination MAC address is set to a multicast group address. The designated multicast group address for linktrace and CCM is 01-80-C2-00-00-3*x*; where *x* represents the maintenance domain (MD) level (for 802.1ag) or the MEG level (for Y.1731). For example, a dot1ag CCM message destined for 01-80-C2-00-00-31 corresponds to MD level 1.

CCM packets using source-destination multicast MAC addresses are for user-initiated messages only (that is, loopbacks).

Ethertype (T)

If dot1q or qinq encapsulation is not configured, then the Ethertype value is 0x8902 and there are no VLAN tags. If dot1q or qinq encapsulation is configured, the VLAN tag (Ethertype value 0x8100) is present and is followed by the Ethertype value of 0x8902, which indicates ETH-CFM messages. The Ethertype is not hard-coded to 0x8100 and can be changed via the port configuration command.

VLAN/dot1p

This is the VLAN/dot1p identifier. If null encapsulation is configured (for Ethernet SAPs or spoke SDP bindings to a VC-type, `ether` or `vlan`), the frame is tagged with NULL.

ETH-CFM OAMPDU

The contents of the Ethernet OAMPDU depend on whether dot1ag or Y.1731 standards are being used. For details on the dot1ag or Y.1731 OAMPDU, see ETH-CFM OAMPDU.

FCS

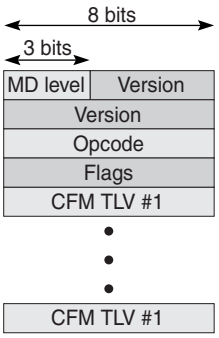
This is the frame check sequence field.

ETH-CFM OAMPDU

As shown in Figure 15, each ETH-CFM OAMPDU message contains the following fields:

- MD level or MEG level: user-configured value, 0 to 7
- version: current version is 0
- opcodes: as defined in IEEE 802.1ag and Y.1731 standards, for messages such as:
  - → Continuity Check Message (CCM)
  - → Loopback Message (LBM)
  - → Loopback Reply (LBR)
  - → Linktrace Message (LTM)
  - → Linktrace Reply (LTR)
- flags: as defined in IEEE 802.1ag and Y.1731 standards
- one or more TLVs, which include:
  - → End TLV
  - → Data TLV
  - → Reply Ingress TLV
  - → Reply Egress TLV
  - → LTM egress identifier TLV
  - → LTR egress identifier TLV
  - → Test TLV

Figure 15: ETH-CFM OAMPDU Message



CFM Frame Processing

Table 10 shows whether a CFM frame received by various MEP types is processed or not. Frames that are processed are extracted from the datapath for CFM processing; unprocessed frames are treated as user traffic and follow the user traffic rules.

Table 10: CFM Frame Processing

MEP Details				Received CFM Frame Treatment	
MEP Type	**Direction**	**VC-Type**	**Port**	**Untagged**	**Tagged** [1]
Spoke SDP	Down	Raw	Any	Processed	Not processed
		VLAN	Any	Not processed	Processed
SAP	Down	Any	Dot1q or QinQ	Not processed [2]	Processed
		Any	Null	Processed	Not processed
	Up	Any	Dot1q or QinQ	Not processed	Processed
		Any	Null	Processed	Not processed

Notes:

1. Tagged frames are single-tagged frames (dot1q) or double-tagged frames (qinq).
2. Untagged frames received on a dot1q-encapsulated port are processed by the Epipe configured to handle untagged frames. The SAP identifier on this Epipe uses VLAN ID 0, also referred to as SAP 0 (for example, 1/1/2:0). Untagged frames are also processed on the *.* QinQ SAP, and on a vlan-x.0 QinQ SAP only if the outer tag matches.

Processing SAP 0 and SAP 0.* OAM Packets

Note the following points with respect to processing OAM packets on SAP 0 (dot1q) and SAP 0.* (qinq). Also note that SAP 0.0 is not supported on OAM packets.

- the 7705 SAR transmits untagged OAM frames on SAP 0 and SAP 0.*
- the 7705 SAR does not process OAM frames tagged with VLAN ID 0 or VLAN ID 0.* on a port configured with null encapsulation
- (for consistency with the 7710 SR) the 7705 SAR processes double-tagged or triple-tagged OAM frames under the following configuration scenario: there is a SAP Up MEP on a dot1q- or qinq-encapsulated port, using SAP 0 or 0.*, and having VC-type VLAN

In this case, the top VLAN tag is removed and the bottom VLAN tag is assumed to be SAP 0 or 0.*. If the bottom VLAN tag is not SAP 0 or 0.*, the VLAN ID is changed to 0. If the frame is an OAM frame, double-tagged (or triple-tagged), the frame is extracted from the SAP Up MEP and the reply is with a single-tagged frame with its VLAN ID set to 0.

- on a SAP Down MEP (dot1q or qinq), untagged frames are processed on SAP 0 or 0.*

MEG-ID and ICC-Based Format

Similar to an 802.1ag MA-ID, a Y.1731 MEG-ID uniquely identifies a group of MEs that are associated at the same MEG level in one administrative domain. The features of MEG-IDs are:

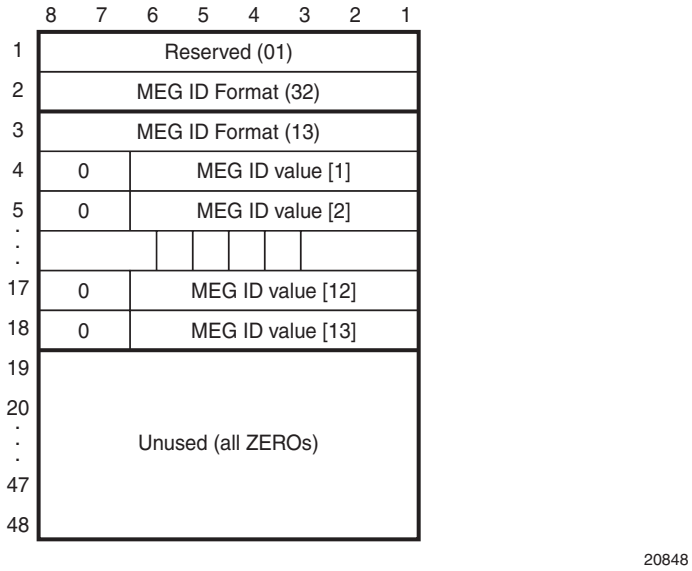
- each MEG-ID must be globally unique
- if the MEG may be required for path setup across interoperator boundaries, then the MEG-ID must be available to other network operators
- the MEG-ID should not change while the MEG remains in existence
- the MEG-ID should be able to identify the network operator that is responsible for the MEG

The 7705 SAR supports the ITU Carrier Code (ICC-based) MEG-ID format (TLV value 32). The generic and ICC-based MEG-ID formats are defined in the ITU-T Y.1731 standard. Figure 16 shows the ICC-based MEG-ID format.

The MEG-ID value has exactly 13 characters and consists of two subfields, the ITU Carrier Code (ICC) followed by a Unique MEG-ID Code (UMC). The ITU Carrier Code consists of between 1 and 6 left-justified characters (alphabetic or leading alphabetic with trailing numeric). The UMC code immediately follows the ICC and consists of between 7 and 12 characters, with trailing NULLs (if necessary to complete the 13 characters).

The UMC is the responsibility of the organization to which the ICC has been assigned, provided that uniqueness is guaranteed.

Figure 16: ICC-based MEG-ID Format



ETH-CFM Functions and Tests

The following list of ETH-CFM functions applies to both dot1ag and Y.1731 Ethernet OAM:

- ETH-CFM—ETH-CFM can be enabled or disabled on a SAP or spoke SDP
- MD levels—eight MD levels can be assigned
- MD name—the following MD name formats are supported:
  - → none (no MD name; used for specifying a Y.1731 functionality)
  - → DNS name
  - → MAC address and 2-octet integer
  - → character string
- MAs—MAs for each MD level can be configured, modified, or deleted
  - → each MA is defined by a unique combination of MD index, MD level, and MA index. This unique combination of values is called the MA identifier (MA-ID).
  - → the following MA name formats are supported:
    - primary VLAN ID (VID)
    - character string (when used with MD name format none, specifies Y.1731 interoperability with 802.1ag)
    - 2-octet integer

  - RFC 2685, Virtual Private Networks Identifier
  - icc-based (used for specifying a Y.1731 functionality)

  → when a VID is used as the MA name, CFM will not support VLAN translation because the unique MA-ID must match all the MEPs

  → the default format for an MA name is a 2-octet integer; integer value 0 means that the MA is not attached to a VID.

- MEPs—Up and Down MEPs on a SAP, and Down MEPs on a spoke SDP

  → MEPs can be configured, modified, or deleted for each MD level (both associations for the Up or Down MEP are with the same Bridge Port; as described in Section 19.2.1 of IEEE Standard 802.1ag-2007)

  → each MEP is uniquely identified by its MEP identifier and MA-ID combination

- MEP creation—MEP creation on a SAP is allowed only for Ethernet ports (with null, dot1q, or qinq encapsulations)

ETH-CFM Ethernet OAM Tests

This section describes Ethernet OAM tests for ETH-CFM on the 7705 SAR, including:

- loopbacks
- linktrace
- throughput measurement
- continuity check
- remote defect indication
- alarm indication signal
- Ethernet (signal) test

Loopback (LB)

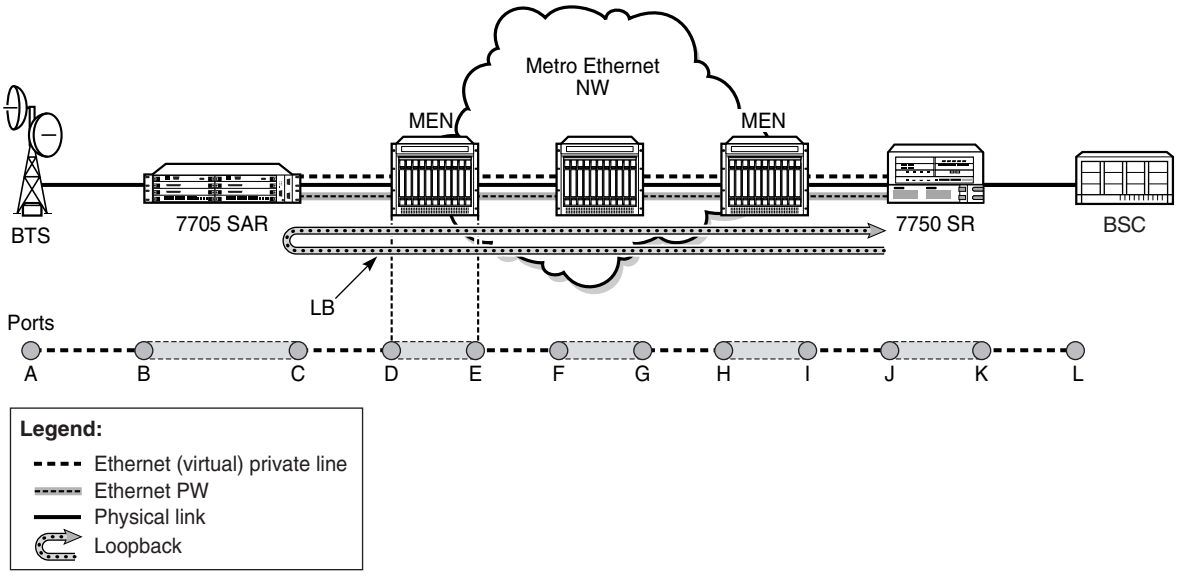
The loopback function is supported by 802.1ag and Y.1731 on the 7705 SAR. A Loopback Message (LBM) is generated by a MEP to its peer MEP. Both dot1ag and dot3ah loopbacks are supported. The loopback function is similar to IP or MPLS ping in that it verifies Ethernet connectivity between the nodes on a per-request basis. That is, it is non-periodic and is only initiated by a user request.

In Figure 17, the line labeled LB represents the dot1ag loopback message between the 7750 SR (source) and 7705 SAR (target). The 7750 SR-generated LBM is switched to the 7705 SAR, where the LBM message is processed. Once the 7705 SAR generates the Loopback Reply message (LBR), the LBR is switched over the PW to the 7750 SR.

The following loopback-related functions are supported:

- loopback message functionality on a MEP can be enabled or disabled
- MEP—supports generating loopback messages and responding to loopback messages with loopback reply messages
- displays the loopback test results on the originating MEP

Figure 17: Dot1ag Loopback Test



Linktrace (LT)

The linktrace function is supported by 802.1ag and Y.1731 on the 7705 SAR. A Linktrace Message (LTM) is originated by a MEP and targeted to a peer MEP in the same MA and within the same MD level. Its function is similar to IP traceroute. The peer MEP responds with a Linktrace Reply (LTR) message after successful inspection of the LTM.

The following linktrace related functions are supported:

- enables/disables LT functions on an MEP
- MEP—supports generating LTMs and responding with LTR messages
- displays linktrace test results on the originating MEP

Throughput Measurement

Throughput measurement is performed by sending frames to the far end at an increasing rate (up to wire speed) and measuring the percentage of frames received back. In general, the rate is dependent on frame size; the larger the frame size, the lower the rate.

The Y.1731 specification recommends the use of unicast ETH-LB and ETH-Test frames to measure throughput.

On the 7705 SAR, LBM processing and LBR generation are enhanced and occur on the datapath, allowing the 7705 SAR to respond to loopback messages at wire speed and making in-service throughput tests possible. Thus, if the 7705 SAR receives LBMs at up to wire speed, it can generate up to an equal number of LBRs.

In order to process LBMs at wire speed, there must be either no TLVs or a single TLV (which is a Data TLV) in the LBM frame. The End TLV field (0) must be present and the frame can be padded with data after the End TLV field in order to increase the size of the frame. Note, however, that the MAC address cannot be a multicast MAC address; it must be the MEP MAC destination address (DA).

Datapath processing of LBMs is supported for the following MEPs:

- dot1ag
  - → SAP Up MEP
  - → SAP Down MEP
  - → spoke SDP Down MEP
- Y.1731
  - → SAP Up MEP
  - → SAP Down MEP

For spoke SDP Down MEPs, fastpath (datapath) LBM processing requires that both interfaces—the LBM receiver and the LBR transmitter—reside on the same adapter card. For example, if the 7705 SAR must perform a reroute operation and needs to move the next-hop interface to another adapter card (that is, LBMs are received on one card and LBRs are transmitted on another), then the fastpath processing of LBMs is terminated and LBM processing continues via the CSM.

Continuity Check (CC)

The continuity check function is supported by 802.1ag and Y.1731 on the 7705 SAR. A Continuity Check Message (CCM) is a multicast frame that is generated by a MEP and sent to its remote MEPs in the same MA. The CCM does not require a reply message. To identify faults, the receiving MEP maintains a MEP database with the MAC addresses of the remote MEPs with which it expects to maintain connectivity checking. The MEP database can be provisioned manually. If there is no CCM from a monitored remote MEP in a preconfigured period, the local MEP raises an alarm.

The following CC capabilities are supported:

- enable and disable CC for a MEP
- automatically put local MEPs into the database when they are created
- manually configure and delete the MEP entries in the CC MEP monitoring database. Note that the only local provisioning required to identify a remote MEP is the remote MEP identifier (using the `remote-mepid` *mep-id* command).
- CCM transmit interval: 10ms, 100ms, 1s, 10s, 1m, 10m (default: 10s)
- transmit interval: 10ms, 100ms, 1s, 10s, 1m, 10m (default: 10s)
- CCM declares a fault when it:
  - → stops hearing from one of the remote MEPs for a period of 3.5 times the CC interval
  - → hears from a MEP with a lower MD level
  - → hears from a MEP that is not in the same MA
  - → hears from a MEP that is in the same MA but is not in the configured MEP list
  - → hears from a MEP that is in the same MA with the same MEP ID as the receiving MEP
  - → recognizes that the CC interval of the remote MEP does not match the local configured CC interval
  - → recognizes that the remote MEP declares a fault

    An alarm is raised and a trap is sent if the defect is greater than or equal to the configured low-priority-defect value.
- CC must be enabled in order for RDI information to be carried in the CCM OAMPDU

ETH-RDI

The Ethernet Remote Defect Indication function (ETH-RDI) is used by a MEP to communicate to its peer MEPs that a defect condition has been encountered. Defect conditions such as signal fail and AIS may result in the transmission of frames with ETH-RDI information. ETH-RDI is used only when ETH-CC transmission is enabled.

ETH-RDI has the following two applications:

- single-ended fault management—the receiving MEP detects an RDI defect condition, which gets correlated with other defect conditions in this MEP and may become a fault cause. The absence of received ETH-RDI information in a single MEP indicates the absence of defects in the entire MEG.
- contribution to far-end performance monitoring—the transmitting MEP reflects that there was a defect at the far end, which is used as an input to the performance monitoring process

A MEP that is in a defect condition transmits frames with ETH-RDI information. A MEP, upon receiving frames with ETH-RDI information, determines that its peer MEP has encountered a defect condition.

The specific configuration information required by a MEP to support the ETH-RDI function is as follows:

- MEG level—the MEG level at which the MEP exists
- ETH-RDI transmission period—application-dependent and is the same value as the ETH-CC transmission period
- priority—the priority of frames containing ETH-RDI information and is the same value as the ETH-CC priority

The PDU used to carry ETH-RDI information is the CCM.

ETH-AIS

The Ethernet Alarm Indication Signal function (ETH-AIS) is a Y.1731 CFM enhancement used to suppress alarms at the client (sub) layer following detection of defect conditions at the server (sub) layer.

Transmission of frames with ETH-AIS information can be enabled or disabled on a Y.1731 MEP.

Frames with ETH-AIS information can be issued at the client MEG level by a MEP, including a server MEP, upon detecting the following conditions:

- signal failure conditions in the case where ETH-CC is enabled
- AIS condition in the case where ETH-CC is disabled

For a point-to-point Ethernet connection at the client (sub) layer, a client layer MEP can determine that the server (sub) layer entity providing connectivity to its peer MEP has encountered a defect condition upon receiving a frame with ETH-AIS information. Alarm suppression is simplified by the fact that a MEP is expected to suppress only those defect conditions associated with its peer MEP.

Only a MEP, including a server MEP, is configured to issue frames with ETH-AIS information. Upon detecting a defect condition, the MEP can immediately start transmitting periodic frames with ETH-AIS information at a configured client MEG level. A MEP continues to transmit periodic frames with ETH-AIS information until the defect condition is removed. Upon receiving a frame with ETH-AIS information from its server (sub) layer, a client (sub) layer MEP detects the AIS condition and suppresses alarms associated with all its peer MEPs. Once the AIS condition is cleared, a MEP resumes alarm generation upon detecting defect conditions.

The following specific configuration information is required by a MEP to support ETH-AIS:

- client MEG level—the MEG level at which the most immediate client layer MEPs exist
- ETH-AIS transmission period—the transmission period of frames with ETH-AIS information
- priority—the priority of frames with ETH-AIS information

ETH-Test

The Ethernet Test (signal) function (ETH-Test) is a Y.1731 CFM enhancement used to perform one-way, on-demand, in-service diagnostics tests, which include verifying frame loss, bit errors, and so on.

Note: The out-of-service diagnostics test is not supported on the 7705 SAR.

When configured to perform such tests, a MEP inserts frames with ETH-Test information such as frame size and transmission patterns.

When an in-service ETH-Test function is performed, data traffic is not disrupted and the frames with ETH-Test information are transmitted.

To support ETH-Test, a Y.1731 MEP requires the following configuration information:

- MEG level—the MEG level at which the MEP exists
- unicast MAC address—the unicast MAC address of the peer MEP for which ETH-Test is intended
- data—an optional element with which to configure data length and contents for the MEP. The contents can be a test pattern and an optional checksum.

  Examples of test patterns include all 0s or all 1s. At the transmitting MEP, this configuration information is required for a test signal generator that is associated with the MEP. At the receiving MEP, this configuration is required for a test signal detector that is associated with the MEP.
- priority—the priority of frames with ETH-Test information

A MEP inserts frames with ETH-Test information towards a targeted peer MEP. The receiving MEP detects the frames with ETH-Test information and performs the requested measurements.

MEP Support (802.1ag and Y.1731)

The 7705 SAR supports Up and Down maintenance association endpoints (MEPs) on Ethernet SAPs for both 802.1ag and Y.1731. It also supports Down MEPs on Ethernet spoke SDPs for 802.1ag only.

802.1ag MEP Support on Ethernet SAPs

The 7705 SAR supports Up and Down MEPs on Ethernet SAPs. Figure 18 shows that the 7705 SAR can terminate and respond to CFM messages received from connected devices, such as base stations, when port B is a Down MEP on a SAP. A CFM message coming from port A would be terminated on port B of the 7705 SAR. As well, port B on the 7705 SAR can generate and send a CFM message towards port A.

Figure 18: MEP on Ethernet Access

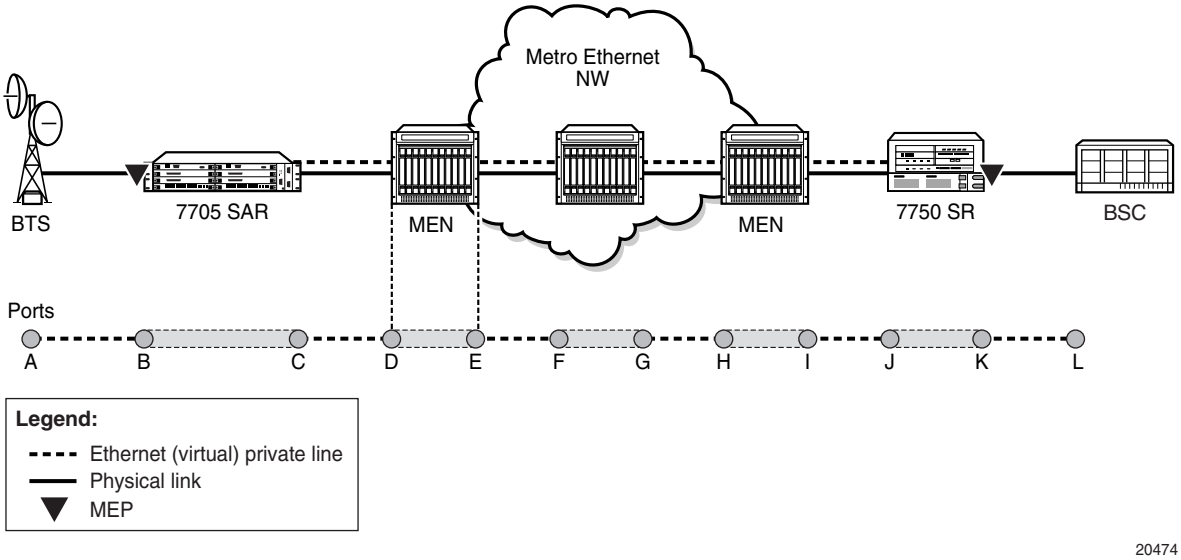
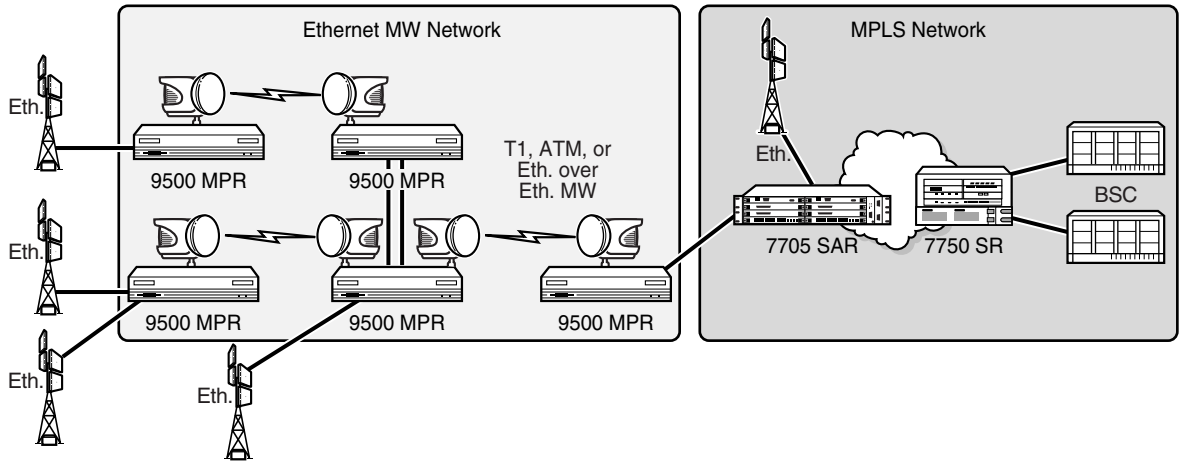


Figure 19 shows how a Down MEP on an Ethernet SAP might be used. In this example, an Ethernet network connects to an access Ethernet port on the 7705 SAR and there are multiple SAPs on that port (that is, multiple endpoints). Since CFM offers OAM capabilities on a per-service basis, which in this case means per SAP (or endpoint), each service can run CFM. Note that if BSC end devices were directly connected to the 7705 SAR (and a VLAN was not used to separate services from each other), EFM would offer capabilities similar to CFM for Ethernet OAM.

In the example shown in Figure 19, separate dot1ag instances initiated on the 9500 MPR nodes can be used to ensure Ethernet layer connectivity on a per-base-station basis. All the traffic from these base stations is aggregated and switched to a single port on the 7705 SAR. Each base station is recognized through a different VLAN, where the VLANs are bound to different services. CFM with traffic in the Down MEP OAMPDU direction at the Ethernet SAP offers the flexibility to run OAM tests on a per-base-station basis.

Figure 19: Down MEP at Ethernet SAP



802.1ag MEP Support on Ethernet Spoke SDPs

The 7705 SAR supports down maintenance association endpoints (MEPs) on Ethernet spoke SDP endpoints. Figure 20 illustrates a Down MEP on an Ethernet spoke SDP.

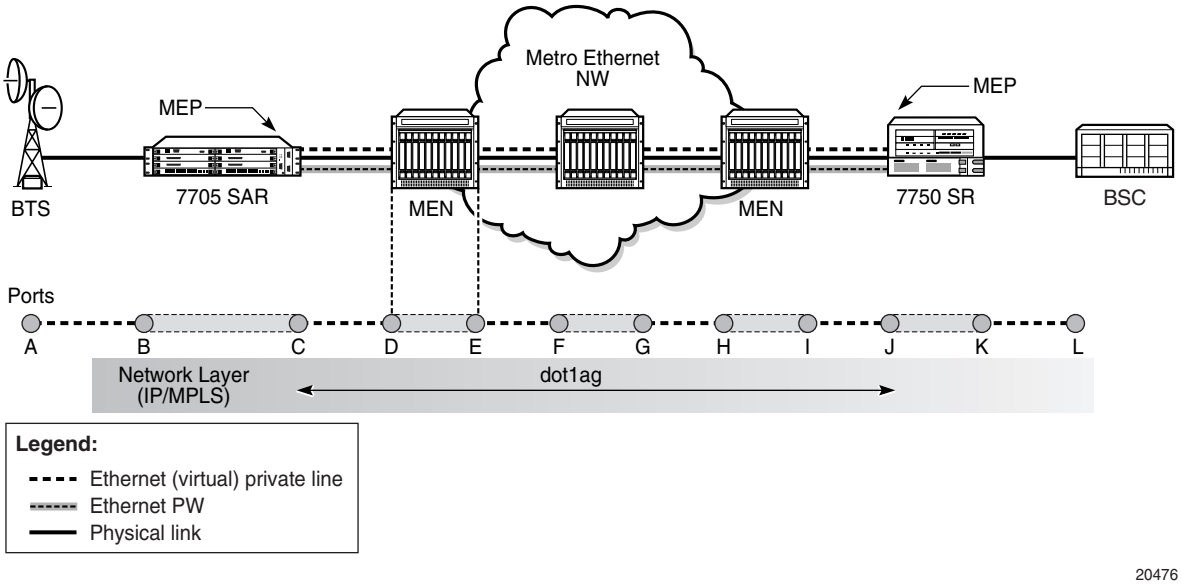
CFM messages can be generated and switched across an Ethernet PW. CFM messages that are received and have an MD that matches the value configured on the 7705 SAR are extracted and processed. Any received CFM messages with an MD level that does not match the configured value are not terminated and are switched transparently to the Ethernet SAP.

Down MEPs on Ethernet spoke SDPs on the 7705 SAR support the following:

- termination of the CFM messages destined for the MEP-ID of the 7705 SAR
- termination of CFM messages at the user-configured domain only
- discarding of OAMPDUs at a lower MD level than the configured one (an alarm message is raised)
- transparent pass-through of upper layer CFM messages
  → MD of the CFM messages that are higher than the one configured on the 7705 SAR

MIP functionality (that is, forwarding of CFM messages with the same MD level) is not supported. Only Down MEP functionality is supported on Ethernet spoke SDPs (that is, termination of CFM messages that are ingress from the Ethernet PW, or generation of CFM packets that are destined for the 7750 SR spoke SDP MEP-ID).

Figure 20: Dot1ag Down MEPs on Spoke SDPs



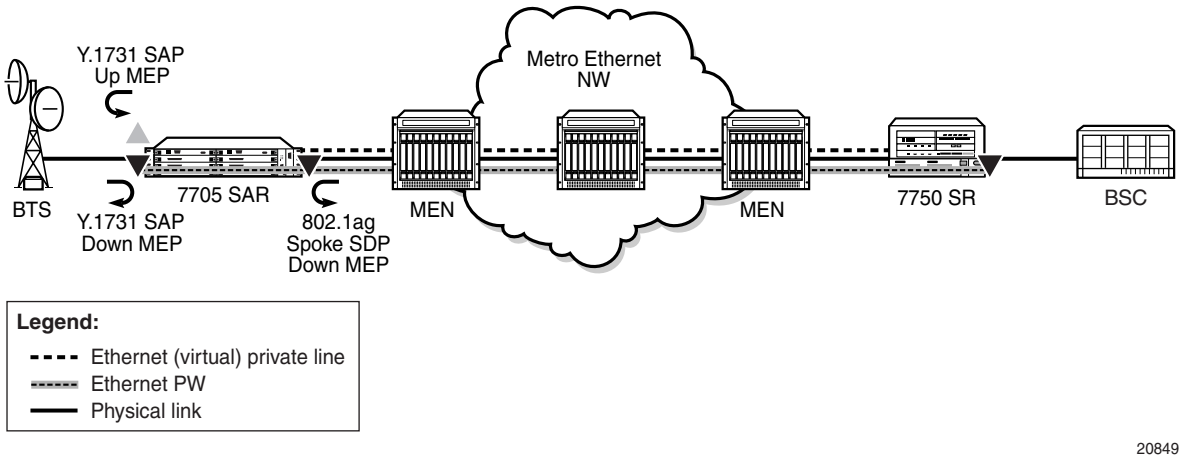
In Figure 20, assuming that the MEP is enabled both on the SR and the 7705 SAR spoke SDP endpoints, the 7705 SAR can generate CFM messages and can terminate any received CFM messages that are destined for the 7705 SAR MEP-ID and have a matching configured domain. Any 7705 SAR-generated CFM packets would traverse the Ethernet PW and would be processed first by the SR node. The Ethernet PW running between the 7705 SAR and the SR generates a pipe-like connectivity; thus, no intermediate Ethernet node can process the CFM messages. All the CFM messages are transported over Ethernet PWs, and PW termination only takes place on SR and 7705 SAR endpoints.

Y.1731 MEP Support on Ethernet SAPs

As shown in Figure 21, the 7705 SAR supports Y.1731 Up and Down MEPs on Ethernet SAPs that are bound to an Ethernet PW service.

Figure 21 also shows an 802.1ag Down MEP on an Ethernet spoke SDP in order to illustrate that when performing CFM tests on the 7705 SAR, a Y.1731 Up MEP on an Ethernet SAP should be used instead of an 802.1ag Down MEP on an Ethernet spoke SDP. Using a Y.1731 SAP Up MEP means that CFM packets verify the switching fabric and SAP status before the packet is processed, because the SAP is on the access side of the 7705 SAR whereas a spoke SDP is on the network side. If a spoke SDP Down MEP is used, packets are terminated and extracted on the network side without being switched through the switching fabric.

Figure 21: Y.1731 MEP Support on the 7705 SAR



Priority Mapping (802.1ag and Y.1731)

Operators often run OAM tests over a single, specific forwarding class (FC). For example, an operator might be mapping OAM traffic to FC2 (AF – Assured Forwarding) and, in order to examine the delay, jitter, or loss qualities of the OAM traffic, would need to run OAM tests using FC2. To provide operators with the ability to control which FC the OAM packets will follow, the `priority` *priority* command is included in several OAM test commands.

When the 7705 SAR generates an Ethernet OAM frame, it uses the priority as per the user's configuration of the `priority` keyword and then sends the frame through the datapath. Thus, the OAM frame follows the entire datapath and receives the same treatment as any other user frame before it is switched over the port.

For example, a CCM frame generated by a SAP Up MEP with a priority value of 7 will receive the following treatment.

- First, the CCM frame is classified as per the access ingress and QoS policy settings. For example, the CCM frame can be mapped to the BE forwarding class if the assigned QoS policy has its priority 7 mapped to BE.
- Then, the OAM packet is mapped to the associated queue (the queue hosting the BE forwarding class) and follows ingress scheduling like any other datapath frame.
- Next, the CCM frame is switched through the fabric and reclassified to the network egress queues, as per the assigned QoS policy classifiers.
- Finally, the CCM frame is scheduled again, as per the queue type and profile state of the queue.

This implementation replicates the user experience since the OAM packet follows the same path as the data packets.

Priority Mapping for SAP Up MEPs

For Up MEPs on a SAP, priority mapping operates as described in the following list, which indicates how the messages or replies generated on ingress have their FC and VLAN tag priority set.

The resulting frames (CCM, LMM, DMM, 1DM, LBM, LMR, DMR, or LBR) are inserted in the access ingress datapath and are processed in the same way as any other frame. That is, they are classified based on the sap-ingress policy.

- Continuity Check messages (CCMs) generated on ingress are based on the setting of the `ccm-ltm-prio` command for the MEP (that is, the VLAN tag priority is set according to the `ccm-ltm-prio` command for the MEP).

- Loss Measurement messages (LMMs), two-way Delay Measurement messages (DMMs), one-way Delay Measurement messages (1DMs), and Loopback messages (LBMs) generated on ingress are based on the priority specified during the LMM, DMM, 1DM, or LBM test (that is, the VLAN tag priority is set according to the priority specified during the test).
- Loss Measurement replies (LMRs), two-way Delay Measurement replies (DMRs), and Loopback replies (LBRs) generated on ingress keep the VLAN tag priority of their corresponding LMM or DMM frame.

Table 11 summarizes the 7705 SAR FC and VLAN priority mappings for SAP Up and Down MEPs based on the frame type.

Table 11: FC and VLAN Priority Mappings for Up and Down MEPs as per Frame Type

MEP Type	Frame Type (Tx)	FC	VLAN priority
SAP Up MEP	CCM	derived from vlan-prio after classification	ccm-ltm-prio
	LMM, DMM, 1DM, LBM	derived from vlan-prio after classification	user-specified
	LMR, DMR, LBR	derived from vlan-prio after classification	preserve incoming query priority
SAP Down MEP	CCM	ccm-ltm-prio	ccm-ltm-prio
	LMM, DMM, 1DM, LBM	user-specified	user-specified
	LMR, DMR, LBR	ccm-ltm-prio	incoming tag priority

Priority Mapping for SAP Down MEPs

For SAP Down MEPs, priority mapping operates as described in the following list, which indicates how the messages or replies generated on egress have their FC and VLAN tag priority set.

- CCMs generated on egress are based on `ccm-ltm-prio` of the MEP (that is, FC and VLAN tag priority are set according to the `ccm-ltm-prio` of the MEP).
- LMM, DMM and 1DM generated on egress are based on the priority specified during the test (that is, FC and VLAN tag priority are set according to the priority specified during the test).
- LMR, DMR, and LBR generated on egress use the `ccm-ltm-prio` of the MEP as FC. The VLAN tag priority is not replaced (that is, the VLAN tag priority of LMM and DMM are kept).

QinQ Support

This section contains the following topics:

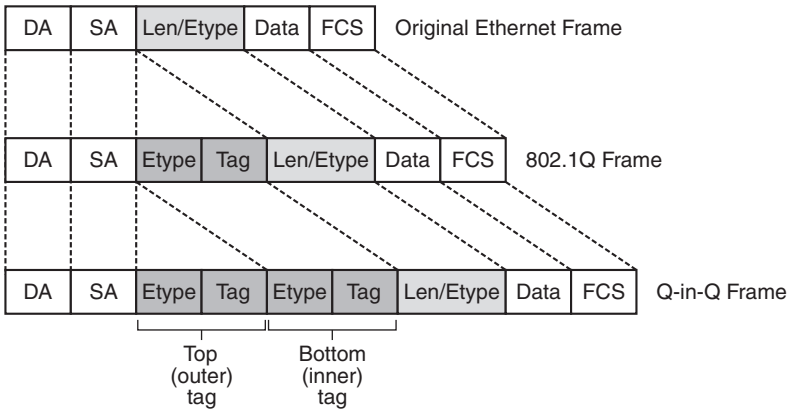
- Overview of QinQ
- QinQ Support with Forced C-Tag Forwarding (VPLS only)
- QinQ Support on Ethernet Ports
- QinQ Configuration Overview

Overview of QinQ

On the 7705 SAR, QinQ (also referred to as VLAN stacking) is based on the IEEE 802.1ad specification. QinQ allows a service provider to use a single VLAN for customers needing multiple VLANs by adding a second VLAN tag to an Ethernet frame, as shown in Figure 22. QinQ encapsulation can be thought of as a dot1q within a dot1q encapsulation.

QinQ operates from end-to-end by receiving customer frames on an ingress SAP, transporting the frames over a service tunnel, and receiving them at the far-end router, where they are unpacked and sent through the egress SAP to the customer.

Figure 22: QinQ Frame



On the 7705 SAR, QinQ ports and QinQ SAPs offer the same feature set as dot1q ports and dot1q SAPs for the following features:

- OAM
- redundancy
- synchronization
- QoS
- card and node limits

QinQ is supported on the following service SAPs (including SAP-to-SAP configurations):

- VLL services (Epipe and Ipipe)
- VPLS
- VPRN
- IES

QinQ is supported on the following adapter cards, modules, and nodes:

- 8-port Ethernet Adapter card, version 2
- 6-port Ethernet 10Gbps Adapter card
- 8-port Gigabit Ethernet Adapter card (all versions)
- 10-port 1GigE/1-port 10GigE X-Adapter Card, version 2 (10-port 1GigE mode)
- Packet Microwave Adapter card
- 4-port SAR-H Fast Ethernet module
- 6-port SAR-M Ethernet module
- 7705 SAR-M
- 7705 SAR-A
- 7705 SAR-W
- 7705 SAR-Wx (Ethernet ports and DSL ports)
- 7705 SAR-F
- 7705 SAR-H
- 7705 SAR-Hc
- 7705 SAR-X

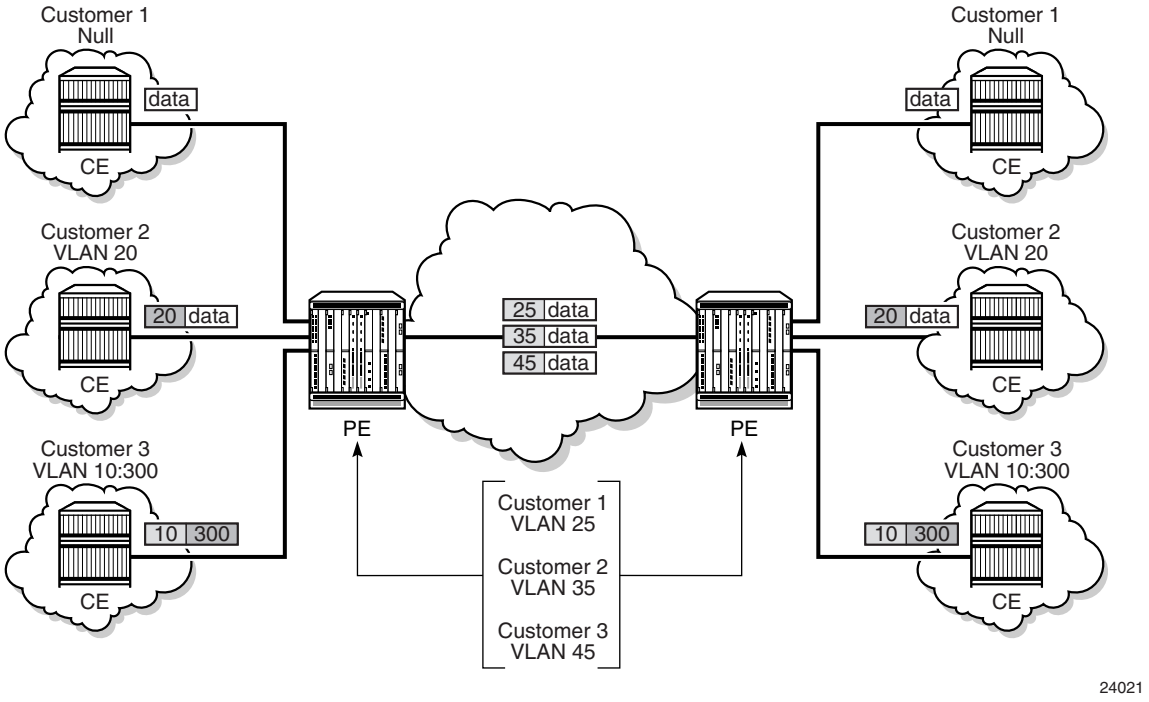


Note: QinQ is not supported by DSL and GPON modules on the 7705 SAR-M.

Figure 23 shows an example of QinQ tagging, where three customer sites each use the same service provider PE router to transport multiple VLANs across an MPLS network. Customer 1 sends an Ethernet frame without a VLAN tag (null encapsulation) and the provider uses MPLS label 25. Customer 2 sends dot1q frames (VLAN ID 20) and the provider uses MPLS label 35. Lastly, customer 3 sends a qinq frame (VLAN ID 10:300) and the provider uses MPLS label 45.

For details on VLAN translation in an Ethernet frame from ingress SAP to egress SAP, see Raw and Tagged Modes.

Figure 23: QinQ Tagging Example



QinQ Support with Forced C-Tag Forwarding (VPLS only)

For VPLS, `force-c-vlan-forwarding` can be user-configured as enabled or disabled. When `force-c-vlan-forwarding` is enabled at the ingress and egress SAPs, the VLAN tag transmitted at the far-end egress SAP is the same as the VLAN tag received at the near-end ingress SAP.

The following examples illustrate the QinQ implementation. Example 1 is a general example. Examples 2 and 3 are examples where the 7705 SAR parses a single frame that has three VLAN tags. For examples 2 and 3, the innermost VLAN Ethertype is always set to 0x8100.

Example 1: General QinQ Implementation

When a service has multiple SAPs configured that can match an incoming frame, the SAP with the most explicit match transmits the frame. For this example, assume that the following SAPs are configured on a port and that the port is in QinQ mode:

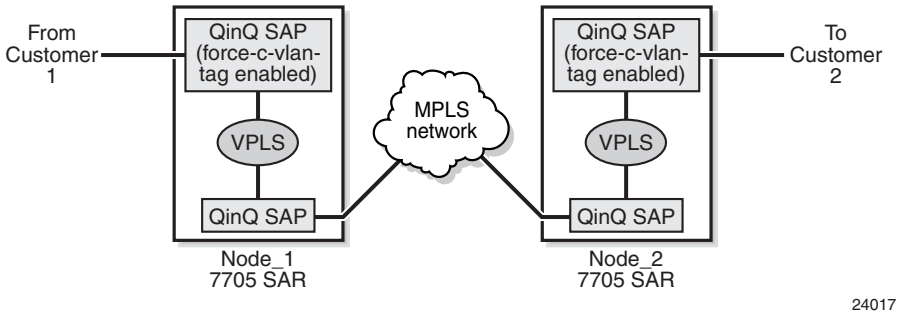
- SAP_1 identifier is *port-id*:vlan-x.vlan-y
- SAP_2 identifier is *port-id*:vlan-x.*
- SAP_3 identifier is *port-id*:*.*

For an incoming frame with VLAN tags vlan-x.vlan-y, SAP_1 processes the frame because it is the most explicit match. Although SAP_2 and SAP_3 also match the frame, the most explicit SAP prevails.

Example 2: QinQ using VPLS with Ethernet SAPs

In this example, assume that the 7705 SAR parses a single frame that has three VLAN tags (innermost VLAN Ethertype is always set to 0x8100). This might occur in the scenario shown in Figure 24, where VPLS QinQ SAPs with `force-c-vlan-forwarding` enabled connect Customer 1 and Customer 2 and preserve the ingress VLAN tag.

Figure 24: QinQ Using VPLS Ethernet SAPs



In Figure 24, the following events occur.

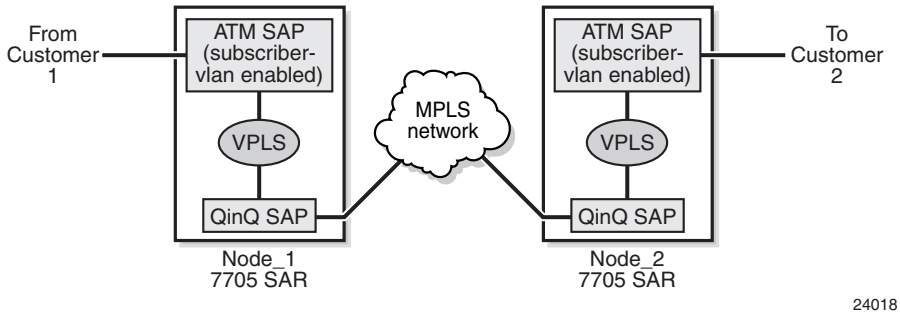
- At Node_1, an ingress QinQ SAP with `force-c-vlan-forwarding` enabled receives a frame and ensures that the bottom (inner) tag is preserved. The frame is then sent to an egress qinq-encapsulated SAP, where the top (outer) tag is swapped and an additional (third) tag is pushed on top.

- At Node_2, the frame with three tags is received on a qinq-encapsulated SAP. The frame is then sent to an egress QinQ SAP with `force-c-vlan-forwarding` enabled, where the egress qinq-encapsulation SAP:
  - → removes the first (top) tag
  - → replaces the second (middle) tag
  - → replaces only the Ethertype of the third (bottom/innermost) tag (that is, the VLAN ID is not replaced)

Example 3: QinQ Using VPLS with ATM SAPs

A second example of triple-tag behavior is shown in Figure 25, where customers are connected using ATM SAPs with `subscriber-vlan` enabled.

Figure 25: QinQ Using VPLS ATM SAPs



In this example, the ATM SAP with `subscriber-vlan` enabled pushes a tag on the frame at ingress. The frame is then sent toward a qinq-encapsulated SAP on egress, which pushes two more tags onto the frame. At the far end, the frame (with three tags) is received on a qinq-encapsulated SAP and sent toward the egress ATM SAP with `subscriber-vlan` enabled. In this case, the egress ATM-encapsulated SAP must be able to remove the three tags.

QinQ Support on Ethernet Ports

This section contains the following topics:

- Special QinQ SAP Identifiers
- QinQ Dot1p Match Behavior
- QinQ Top-Only Mark Option
- Maximum Number of VLAN Tags

QinQ requires that the `encap-type` for the associated port be set for `qinq`, and that the *`sap-id`* include two Q-tags (VLAN IDs). An ingress QinQ SAP can be configured so that dot1p bits for packet QoS classification can come from the top or the bottom Q-tag. At egress, if dot1p re-marking is configured, both Q-tags are re-marked. However, you can use `qinq-mark-top-only` to re-mark only the top Q-tag.

Special QinQ SAP Identifiers

Table 12 describes the special SAP identifiers used on the 7705 SAR. In Table 12, a *`qtag`* value represents a VLAN ID. The * (asterisk) represents a reserved VLAN that is used to carry traffic from any unused VLAN on the port. An unused VLAN is a VLAN that is not explicitly defined on the port.

Table 12: Special QinQ SAP Identifiers

QinQ SAP Type	SAP Identifier	Example	Notes
Explicit (specified)	*`port-id`*`:`*`qtag1`*`.`*`qtag2`* [1]	1/1/1:20.2000	*`qtag1`* and *`qtag2`* range: 0 to 4094, and *
Null	*`port-id`* *`port-id`*:0 *`port-id`*:0.*	1/1/1 1/1/1:0 1/1/1:0.*	No suffix means no VLANs 0 suffix means SAP 0 0.* suffix means SAP 0 and any unused VLAN on the port
Default	*`port-id`*:* *`port-id`*:*.* *`port-id`*`:`*`qtag1`*.*	1/1/1:* 1/1/1:*.* 1/1/1:40.*	*.* means any unused VLAN on the port [2]

Notes:

1. *qtag1* is the top (outer) tag. *qtag2* is the bottom (inner) tag.
2. The behavior of the *.* SAP on the 7705 SAR is different from the *.* SAP behavior on the 7750 SR. On the 7750 SR, *.* represents a capture SAP.

Note: The following default SAPs are not supported on the 7705 SAR:

- *port-id*:*.0
- *port-id*:*.vlan-y

The following list describes how the SAP types in Table 12 process frames. Packet breakdowns are described in Raw and Tagged Modes.

- default SAP (*port-id*:*qtag1*.*)
  - → receives all frames with an explicit outer tag value of *qtag1*, regardless of the inner tag
  - → the outer tag is stripped and the inner tag is passed transparently
  - → example: SAP 1/1/1:10.* only matches a top Q tag of VLAN 10. The SAP may have any bottom tag or no bottom tag at all.
- null SAP (*port-id*:0.*)
  - → receives all untagged frames and/or any frames with a VLAN tag of 0
  - → example: SAP 1/1/1:0.* allows any untagged frames and/or frames with an outer tag of "0" only

    If the outer tag is 10, then the frame is dropped. If the outer tag is 0 and the inner tag is any valid VLAN ID, then the frame is not dropped.
- null inner SAP (*port-id*:*qtag1*.0)
  - → receives all frames with explicit outer tag value *qtag1*, and may have an inner tag of 0 or no inner tag at all
  - → example: SAP 1/1/1:10.0 or SAP 1/1/1:10 will only match "10" as the outer tag, and may have a bottom tag of 0 or no bottom tag at all. The SAP 1/1/1:1/10 will be dropped because its outer tag is not "10".
- invalid SAPs (not supported) (*port-id*:*.*qtag2* and *port-id*:*.0)

QinQ Dot1p Match Behavior

Since a qinq-encapsulated packet has top and bottom Q-tags, you can specify which qtag position provides the dot1p bits (P-bits) when QoS evaluates the P-bits in an ingress packet for a classification match. This is done with the `sap>ingress>match-qinq-dot1p` command. The default setting is to use the P-bits from the inner (bottom) Q-tag, or if the ingress packet has no Q-tags or has null encapsulation, then no match filtering is done.

QinQ Top-Only Mark Option

By default, if dot1p re-marking is configured at SAP egress on a QinQ SAP, the dot1p bits for both the top and bottom tags are re-marked. However, re-marking can be configured such that only the top Q-tag P-bits get remarked by enabling the `sap>egress>qinq-mark-top-only` command. The DEI bit is ignored.

Maximum Number of VLAN Tags

The maximum number of VLAN tags allowed in a packet depends on whether the service requires Layer 3 packet parsing (for example, to access to the IP header), as follows.

- For services that do not require access to the Layer 3 IP header, such as Ethernet PWs and VPLS, there is no limit on the number of VLAN tags in the packet. Any packet with any number of VLAN tags is accepted.
- For services that require access to the Layer 3 IP header, such as IP PWs, IES, and VPRN, there is a limit of three or fewer VLAN tags. When a packet is received that has four VLAN tags, the data following the third VLAN tag is considered as an IP packet. As a result, packets with four VLAN tags will be dropped by these services due to the inability to decode the IP header.

QinQ Configuration Overview

The basic steps to configure QinQ are as follows.

1. Configure the port `encap-type` for qinq and (optionally) set the Ethertype (`qinq-etype`).
2. Create one or more SAPs for the service.
3. Configure the SAP ingress dot1p re-marking behavior (match the top or bottom tag).
4. Enable the SAP egress `qinq-mark-top-only command`.
5. Configure the spoke SDP `vc-type` (ether or vlan).

Service Creation Overview

Figure 26 shows a flowchart that provides an overview of the process to create a service. Service creation can be separated into two main functional areas—core services tasks and subscriber services tasks. Core services tasks are performed prior to subscriber services tasks.

Before starting the process shown in Figure 26, ensure that the 7705 SAR system has been configured with an IP address and (for the 7705 SAR-8 or 7705 SAR-18) has the appropriate adapter cards installed and activated.

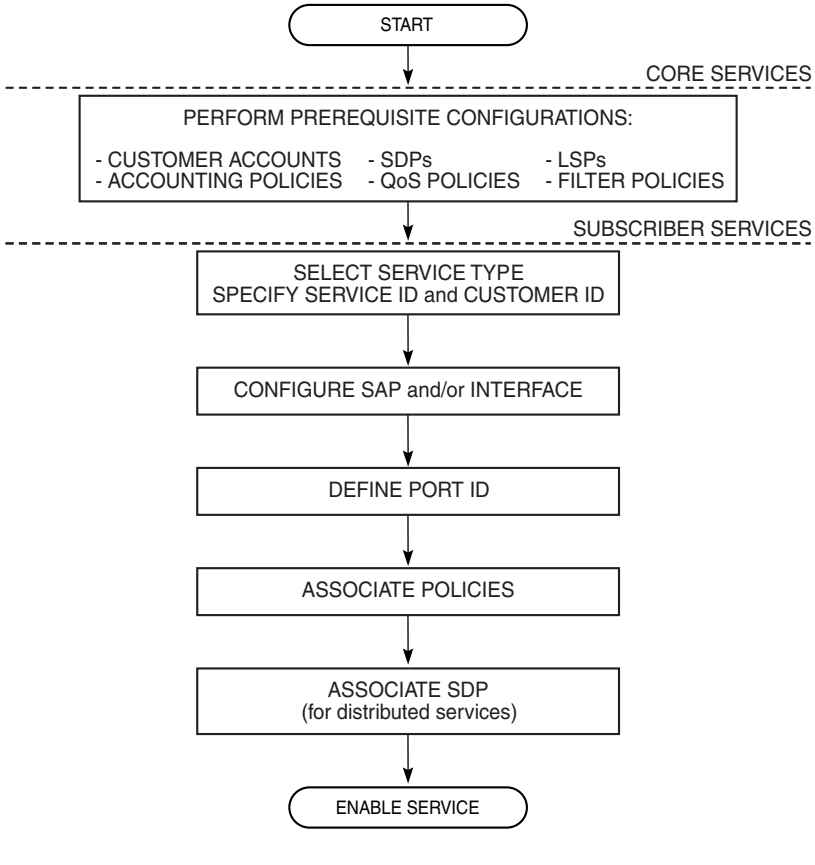
Core services tasks include the following items:

- create customer accounts
- create template QoS and accounting policies
- create LSPs
- create SDPs

Subscriber services tasks include the following items:

- create VLL (Apipe, Cpipe, Epipe, Fpipe, Hpipe, or Ipipe), IES, VPLS, or VPRN services
- configure SAPs
- bind SDPs
- create exclusive QoS policies
- (optionally) assign IP filter policies to Epipe SAPs, Ipipe SAPs, VPLS SAPs, VPRN interface SAPs, IES interface SAPs, and IES Management interface SAPs. IP filter policies can also be applied to VPLS SDPs (ingress mesh and spoke), and VPRN and IES interface ingress spoke SDPs.
- (optionally) assign MAC ingress filter policies to VPLS SAPs and VPLS SDPs (mesh and spoke)

Figure 26: Service Creation and Implementation Flowchart



Port and SAP CLI Identifiers

When typing text in the command line interface (CLI), *port-id* is often displayed to indicate that a port identifier may need to be typed in the command line. Similarly, to identify a SAP, the *port-id* is used, but additional information may need to be appended to indicate a logical sub-element of the port.

On the CLI, a *port-id* is defined using the format *slot*/*mda*/*port*, where *slot* identifies the IOM card slot (always 1), *mda* identifies the physical slot in the chassis for the adapter card, and *port* identifies the physical port on the adapter card.

The value that can be appended to a SAP has the format [:*ID*], [.ID], [:*ID*/*ID*], or [:*ID*.*ID*]. The colon or dot and following ID identify a sub-element of the port (if applicable), such as a TDM channel group for a Cpipe, a VPI/VCI value for an Apipe, or a dot1q or qinq VLAN ID for an Epipe.

For example, a SAP associated with a TDM channel group on port 12 of an 16-port T1/E1 ASAP Adapter card in MDA slot 3 is identified as <1/3/12.3>, where ".3" is the appended value and identifies that for this SAP the channel group begins in timeslot 3.

Reference Sources

For information on standards and supported MIBs, refer to Standards and Protocol Support.

Configuring Global Service Entities with CLI

This section provides information to create subscriber (customer) accounts and to configure service destination points (SDPs) using the command line interface.

Topics in this section include:

- Service Model Entities
- Basic Configuration
- Common Configuration Tasks
- ETH-CFM (802.1ag and Y.1731) Tasks
- Service Management Tasks

Service Model Entities

The Alcatel-Lucent service model uses (logical) service entities to construct a service. Each entity within the model has properties that describe it and influence its behavior. The service model has four main entities to configure a service. The entities are:

- Customers
  → Configuring Customer Accounts
- Service Destination Points (SDPs)
  → Configuring SDPs
- Service Types
  → ATM VLL (Apipe) Services
  → Circuit Emulation VLL (Cpipe) Services
  → Ethernet VLL (Epipe) Services
  → Frame Relay VLL (Fpipe) Services
  → HDLC VLL (Hpipe) Services
  → IP Interworking VLL (Ipipe) Services
  → VPLS
  → Internet Enhanced Service
  → VPRN Services
- Service Access Points (SAPs)
  → Configuring Apipe SAP Parameters
  → Configuring Cpipe SAP Parameters
  → Configuring Epipe SAP Parameters
  → Configuring Fpipe SAP Parameters
  → Configuring Hpipe SAP Parameters
  → Configuring Ipipe SAP Parameters
  → Configuring a VPLS SAP
  → Configuring IES SAP Parameters
  → Configuring VPRN Interface SAP Parameters

Basic Configuration

Before configuring a subscriber service, the QoS, logs, and MPLS LSPs (if applicable) must be configured. Refer to the following guides for more information:

- 7705 SAR OS Quality of Service Guide
- 7705 SAR OS Router Configuration Guide
- 7705 SAR OS System Management Guide
- 7705 SAR OS MPLS Guide

A basic service configuration must have the following items configured:

- a customer ID
- a service type
- a service ID (a *service-id* number is mandatory and a *service-name* is optional)
- a SAP identifying a port and encapsulation value
- an interface (where required) identifying an IP address, IP subnet, and broadcast address
- an associated SDP (for distributed services)

The following example shows an Epipe service configuration displaying the SDP and Epipe service entities. SDP ID 2 was created with the far-end node 10.10.10.104. Epipe ID 6000 was created for customer ID 6, which uses the SDP ID 2.

```
A:ALU-B>config>service# info detail
#------------------------------------------
...
        sdp 2 mpls create
            description "MPLS-10.10.10.104"
            far-end 10.10.10.104
            ldp
            signaling tldp
            no vlan-vc-etype
            no path-mtu
            keep-alive
                shutdown
                hello-time 10
                hold-down-time 10
                max-drop-count 3
                timeout 5
                no message-length
            exit
            no shutdown
        exit
...
...
        epipe 6000 customer 6 vpn 6000 create
           service-mtu 1514
```

```
            service-name "Epipe_6000"
            sap 1/1/2:0 create
                ingress
                    filter ip 1
                    qos 1
                exit
                egress
                     qos 1
                exit
            no shutdown
            exit
            spoke-sdp 2:6111 create
                ingress
                    no vc-label
            exit
                egress
                    no vc-label
            exit
                no shutdown
            exit
            no shutdown
        exit
...
#------------------------------------------
A:ALU-B>config>service#
```

Common Configuration Tasks

This section provides a brief overview of the following common configuration tasks that must be performed to configure a customer account and an SDP:

- Configuring Customer Accounts
- Configuring SDPs
- Configuring Service Names

Configuring Customer Accounts

Use the `customer` command to configure customer information. Every customer account must have a customer ID. Optional parameters include:

- description
- contact name
- telephone number

If special characters are included in the customer description string, such as spaces, #, or ?, the entire string must be enclosed in double quotes.

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
```

Example:

```
config>service# customer 5 create
config>service>cust# contact "Technical Support"
config>service>cust$ description "Alcatel-Lucent
 Customer"
config>service>cust# phone "650 555-5100"
config>service>cust# exit
```

The following example displays the customer account configuration output.

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

Configuring SDPs

Every service destination point (SDP) must have the following items configured:

- a locally unique SDP identification (ID) number
- the system IP address of the far-end router
- an SDP encapsulation type—MPLS, GRE, or IP

SDP Configuration Considerations

Consider the following SDP characteristics when creating and configuring an SDP.

- SDPs can be configured as MPLS, GRE, or IP.
- If an SDP configuration does not include the IP address of the associated far-end router, then VLL, VPLS, and VPRN services to the far-end router cannot be provided.
- A distributed service must be bound to an SDP.

  By default, SDPs are not associated with services. Once an SDP is created, services can be associated with that SDP.
- An SDP can have more than one service bound to it. That is, an SDP is not specific or exclusive to any one service or any type of service.
- When configuring an SDP:
  - → The far-end SDP IP address must be the system IP address of a 7705 SAR or an SR-series router.
  - → For MPLS SDPs, LSPs must be configured before the LSP-to-SDP associations can be assigned. The LSP-to-SDP associations must be created explicitly.
  - → Automatic ingress and egress labeling (targeted LDP) is enabled by default. Ingress and egress VC labels are signaled over a targeted LDP connection between two 7705 SAR routers.



Note: If signaling is disabled for an SDP, then ingress and egress vc-labels for the services using that SDP must be configured 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

When configuring an SDP, consider the following points.

- If you do not specify an encapsulation type, the default is MPLS.
- When configuring a distributed service, you must identify an SDP ID and the far-end IP address. Use the `show>service>sdp` command to display a list of qualifying SDPs.
- When specifying MPLS SDP parameters, you can either specify an LSP or enable an LDP. There cannot be two methods of transport in a single SDP. If an LSP is specified, then RSVP is used for dynamic signaling within the LSP.
- LSPs are configured in the `config>router>mpls` context. See the 7705 SAR OS MPLS Guide for configuration and command information.

Use the following CLI syntax to create an SDP.

CLI Syntax:

```
config>service>sdp sdp-id [gre | mpls | ip] create
     adv-mtu-override
     description description-string
     far-end ip-addr
     keep-alive
          hello-time seconds
          hold-down-time seconds
          max-drop-count count
          message-length octets
          timeout timeout
          no shutdown
     ldp                       (for MPLS SDPs only)
     lsp lsp-name [lsp-name] (for MPLS SDPs only)
     path-mtu octets
     signaling {off|tldp}
     no shutdown
```

Example:

```
config>service# sdp 2 gre create
config>service>sdp# description "GRE-10.10.10.104"
config>service>sdp# far-end "10.10.10.104"
```

```
config>service>sdp# no shutdown
config>service>sdp# exit
config>service# sdp 4 mpls create
config>service>sdp# description “MPLS-10.10.10.104”
config>service>sdp# far-end “10.10.10.104”
config>service>sdp# ldp
config>service>sdp# no shutdown
config>service>sdp# exit
config>service# sdp 8 mpls create
config>service>sdp# description "MPLS-10.10.10.104"
config>service>sdp# far-end “10.10.10.104”
config>service>sdp# lsp "to-104"
config>service>sdp# no shutdown
config>service>sdp# exit
config>service# sdp 104 mpls create
config>service>sdp# description "MPLS-10.10.10.94"
config>service>sdp# far-end “10.10.10.94”
config>service>sdp# ldp
config>service>sdp# no shutdown
config>service>sdp# exit
```

The following example displays the SDP sample configuration output.

```
A:ALU-12>config>service# info
----------------------------------------
...
        sdp 2  create
            description "GRE-10.10.10.104"
            far-end 10.10.10.104
            keep-alive
                shutdown
            exit
            no shutdown
        sdp 4  create
            description "MPLS-10.10.10.104"
            far-end 10.10.10.104
            ldp
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        sdp 8 mpls create
            description "MPLS-10.10.10.104"
            far-end 10.10.10.104
            lsp "to-104"
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        sdp 104 mpls create
            description "MPLS-10.10.10.94"
            far-end 10.10.10.94
```

```
            ldp
            keep-alive
                shutdown
            exit
            no shutdown
        exit
...
-------------------------------------------
A:ALU-12>config>service#
```

Configuring Service Names

After a service has been created, it can be assigned a service name.

Use the following CLI syntax to assign a service name to a service. The syntax example is for an Epipe.

CLI Syntax: `config>service>epipe# service-name` *`service-name`*

Example:
```
config>service# epipe 1 customer 1 create
config>service>epipe# service-name "Epipe_1"
```

The following example displays the service name configuration output.

```
A:ALU-12>config>service>epipe# info
----------------------------------------------
...
            shutdown
            service-name "Epipe_1"
            sdp 2  create
                description "GRE-10.10.10.104"
                far-end 10.10.10.104
                keep-alive
                    shutdown
                exit
                no shutdown
            exit
...
-------------------------------------------
A:ALU-12>config>service#
```

ETH-CFM (802.1ag and Y.1731) Tasks

This section provides a brief overview of the following ETH-CFM tasks:

- Configuring ETH-CFM Parameters (802.1ag and Y.1731)
- Applying ETH-CFM Parameters

Configuring ETH-CFM Parameters (802.1ag and Y.1731)

Configuration commands for both the 802.1ag and the Y.1731 functions are entered in an `eth-cfm` context (either global or Epipe service). For information on Ethernet OAM commands for 802.1ag and Y.1731 OAM, refer to the "Ethernet OAM Capabilities" section in the 7705 SAR OS OAM and Diagnostics Guide.

An 802.1ag MEP and a Y.1731 MEP are similar in function. Configure a MEP to be a Y.1731 MEP by choosing the `format none` keywords in the global `domain` command, and the `format icc-based` keywords in the `global association` command. Configure a MEP to be a Y.1731 MEP that can interoperate with a 802.1ag MEP by choosing the `format none` keywords in the global `domain` command, and the `format string` keywords in the `global association` command.

802.1ag Configuration

The first set of commands occurs at the global level. The second set occurs at the Epipe service level.

```
*A:ALU-1>config>eth-cfm# info
----------------------------------------------
    domain 1 name "kanata_MD" level 5
        association 1 format string name "kanata_MA"
            bridge-identifier 2
            exit
            ccm-interval 60
            remote-mepid 125
        exit
    exit
----------------------------------------------
*A:ALU-1>config>service>epipe# info
----------------------------------------------
    shutdown
    sap 1/5/1 create
        eth-cfm
            mep 1 domain 1 association 1 direction down
                shutdown
            exit
        exit
```

```
    exit
    spoke-sdp 1:11 create
        eth-cfm
            mep 2 domain 1 association 1 direction down
                shutdown
            exit
        exit
        exit
-----------------------------------------------
```

Note: RDI information is carried in the CCM OAMPDU. To be able to transmit and also receive RDI information, a MEP must have CCM enabled. See Applying ETH-CFM Parameters.

Y.1731 Configuration

The following example displays a Y.1731 configuration. The first set of commands occurs at the global level. The second set occurs at the Epipe service level.

```
*A:ALU-1>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:ALU-1>config>service>epipe# info
-----------------------------------------------
    shutdown
    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
...
-----------------------------------------------
```

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

Applying ETH-CFM Parameters

Apply ETH-CFM parameters to the following entities, as shown in the CLI syntax examples below:

- Epipe SAP
- Epipe spoke SDP
- OAM tests (loopback, linktrace, Ethernet test, delay measurement, loss measurement, synthetic loss measurement)

The MAC address for a MEP on an Epipe cannot be changed. For a MEP on a SAP, the MAC address is the port MAC address. For a MEP on a spoke SDP, the MAC address is the system MAC address.

The 7705 SAR supports the following MEPs:

- 802.1ag
  - → SAP Up MEP
  - → SAP Down MEP
  - → spoke SDP Down MEP
- Y.1731
  - → SAP Up MEP
  - → SAP Down MEP

CLI Syntax:

```
config>service>epipe>sap
eth-cfm
     hold-mep-up-on-failure
     mep mep-id domain md-index association ma-index
        [direction {up | down}]
          ccm-enable
          ais-enable
          ccm-ltm-priority priority
          dual-ended-loss-test-enable
          eth-test-enable
          low-priority-defect {allDef|macRemErrXcon|
            remErrXcon|errXcon|xcon|noXcon}
          one-way-delay-threshold seconds
          [no] shutdown
```

CLI Syntax: config>service>epipe>spoke-sdp
eth-cfm
mep *mep-id* domain *md-index* association *ma-index*
[direction {up | down}]
ccm-enable
ccm-ltm-priority *priority*
low-priority-defect {allDef|macRemErrXcon|
remErrXcon|errXcon|xcon|noXcon}
[no] shutdown

CLI Syntax: oam
eth-cfm linktrace *mac-address* mep *mep-id* domain *md-index* association *ma-index* [ttl *ttl-value*]
eth-cfm loopback *mac-address* mep *mep-id* domain *md-index* association *ma-index* [send-count *send-count*] [size *data-size*] [priority *priority*]
eth-cfm eth-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [data-length *data-length*]
eth-cfm one-way-delay-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*]
eth-cfm two-way-delay-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*]
eth-cfm single-ended-loss-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [interval {100ms|1s}] [send-count *send-count*]
eth-cfm two-way-slm-test *mac-address* mep *mep-id* domain *md-index* association *ma-index* [priority *priority*] [send-count *send-count*] [size *data-size*] [timeout *timeout*] [interval *interval*]

Service Management Tasks

This section provides a brief overview of the following service management tasks:

- Modifying Customer Accounts
- Deleting Customers
- Modifying SDPs
- Deleting SDPs
- Deleting LSP Associations

Modifying Customer Accounts

Use the `show>service>customer` command to display a list of customer IDs.

To modify a customer account:

1. Access the specific account by specifying the customer ID.
2. Enter the parameter to modify (`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 typically removes a customer ID and all associated information; however, 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
```

Modifying SDPs

Use the show>service>sdp command to display a list of SDP IDs.

To modify an SDP:

1. Access the specific SDP by specifying the SDP ID.
2. Enter the parameter to modify, such as description, far-end, or lsp, and then enter the new information.



Note: Once the SDP is created, you cannot modify the SDP encapsulation type.

CLI Syntax: config>service# sdp *sdp-id*

Example:
```
config>service# sdp 79
config>service>sdp# description "Path-to-107"
config>service>sdp# shutdown
config>service>sdp# far-end "10.10.10.107"
config>service>sdp# path-mtu 1503
config>service>sdp# no shutdown
```

Deleting SDPs

The no form of the sdp command typically removes an SDP ID and all associated information; however, before an SDP can be deleted, the SDP must be shut down and removed (unbound) from all customer services where it is applied.

CLI Syntax:
```
config>service# no sdp 79
```

Example:
```
config>service# epipe 5 spoke-sdp 79:5
config>service>epipe>spoke-sdp# shutdown
config>service>epipe>spoke-sdp# exit
config>service>epipe 5 no spoke-sdp 79:5
config>service>epipe# exit
config>service# no sdp 79
```

Deleting LSP Associations

The no form of the lsp command removes an LSP ID and all associated information; however, before an LSP can be deleted, it must be removed from all SDP associations.

CLI Syntax:
```
config>service# sdp sdp-id
[no] lsp lsp-name
```

Example:
```
config>service# sdp 79
config>service>sdp# no lsp 123
config>service>sdp# exit all
```

Global Service Command Reference

Command Hierarchies

- Global Service Configuration Commands
  - → Customer Commands
  - → SDP Commands
  - → SAP Commands
  - → ETH-CFM Commands
- Show Commands

Global Service Configuration Commands

Customer Commands

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

SDP Commands

```
config
    — service
        — sdp sdp-id [gre | mpls | ip] [create]
        — no sdp sdp-id
            — [no] adv-mtu-override
            — [no] bgp-tunnel
            — description description-string
            — no description
            — encryption-keygroup keygroup-id direction {inbound | outbound}
            — no encryption-keygroup direction {inbound | outbound}
            — 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
            — [no] lsp lsp-name
            — metric metric
            — no metric
            — path-mtu bytes
            — no path-mtu
            — signaling {off | tldp}
```

```
— [no] shutdown
— vlan-vc-etype 0x0600..0xffff
— no vlan-vc-etype [x0600.0xffff]
```

SAP Commands

```
config
    — service
        — apipe
            — sap sap-id [create]
            — no sap sap-id
        — cpipe
            — sap sap-id [create]
            — no sap sap-id
        — epipe
            — sap sap-id [create]
            — no sap sap-id
        — fpipe
            — sap sap-id [create]
            — no sap sap-id
        — hpipe
            — sap sap-id [create]
            — no sap sap-id
        — ipipe
            — sap sap-id [create]
            — no sap sap-id
        — ies
            — interface ip-int-name [create]
                — sap sap-id [create]
                — no sap sap-id
        — vprn
            — interface ip-int-name [create]
                — sap sap-id [create]
                — no sap sap-id
        — vpls
            — sap sap-id [create]
            — no sap sap-id
```

ETH-CFM Commands

```
config
    — eth-cfm
        — domain md-index [format {dns | mac | none | string}] name md-name level level
        — domain md-index
        — no domain md-index
            — association ma-index [format {icc-based | integer | string | vid | vpn-id}] name
                ma-name
            — association ma-index
            — no association ma-index
                — [no] bridge-identifier bridge-id
                    — vlan vlan-id
                    — no vlan
                — ccm-interval {10ms | 100ms | 1 | 10 | 60 | 600}
                — no ccm-interval
                — [no] remote-mepid mep-id
        — slm
            — inactivity-timer timeout
            — no inactivity-timer
```

Note: For information on configuring ETH-CFM commands, refer to the 7705 SAR OS OAM and Diagnostics Guide.

Show Commands

```
show
    — service
        — customer customer-id
        — sdp [sdp-id | far-end ip-addr] [detail | keep-alive-history]
        — sdp-using sdp-id[:vc-id] | far-end ip-address]
        — service-using [epipe] [fpipe] [hpipe] [ies] [vprn] [vpls] [apipe] [cpipe] [sdp sdp-id]
            [customer customer-id]
```

Command Descriptions

- Global Service Configuration Commands
- Show Commands

Global Service Configuration Commands

- Generic Commands
- Customer Commands
- SDP Commands
- SDP Keepalive Commands

Generic Commands

description

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

Context config>service>customer
config>service>sdp

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

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

Default No description is associated with the configuration context.

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

shutdown

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

Context config>service>sdp
config>service>sdp>keep-alive

Description The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many objects must be shut down before they may be deleted. Many entities must be explicitly enabled using the **no shutdown** command.

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

Services are created in the administratively down state (**shutdown**). 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 are described in the following Special Cases.

Special Cases **Service Admin State —** bindings to an SDP within the service will be put into the out-of-service state when the service is shut down. While the service is shut down, all customer packets are dropped and counted as discards for billing and debugging purposes.

SDP (global) — when an SDP is shut down at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would normally be transmitted using this SDP binding will be discarded and counted as dropped packets.

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

SDP Keepalives — enables SDP connectivity monitoring keepalive messages for the SDP ID. Default state is disabled (shutdown), in which case the operational state of the SDP-ID is not affected by the keepalive message state.

Customer 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 and the **create** keyword must follow each new **customer** *customer-id* entry.

To edit a customer's parameters, enter the existing **customer** *customer-id* without the **create** keyword.

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 to 2147483647

contact

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

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.

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

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

Parameters *contact-information* — the customer contact information entered as an ASCII 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.

phone

Syntax [**no**] **phone** *phone-number*

Context config>service>customer

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

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

Default No telephone number information is associated with a customer.

Parameters *phone-number* — the customer phone number entered as an ASCII 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.

SDP Commands

sdp

Syntax **sdp** *sdp-id* [**gre** | **mpls** | **ip**] [**create**]
no sdp *sdp-id*

Context config>service

Description This command creates or edits an SDP. SDPs must be explicitly configured.

An SDP is a (logical) service entity that is created on the local router. An SDP identifies the endpoint of a logical, unidirectional service tunnel. Traffic enters the tunnel at the SDP on the local router and exits the tunnel at the remote router. Thus, it is not necessary to specifically define far-end SAPs.

The 7705 SAR supports generic routing encapsulation (GRE) tunnels, multiprotocol label switching (MPLS) tunnels, and IP tunnels.

For MPLS, a 7705 SAR 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 LSPs are established in LDP-DU (downstream unsolicited) mode.

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. SDPs are created in the admin down state (**shutdown**). Once all relevant parameters are defined, the **no shutdown** command must be executed before the SDP can be used.

If *sdp-id* exists, the current CLI context is changed to that SDP for editing and modification. If editing an existing SDP, the **gre**, **mpls,** or **ip** keyword is not specified. If a keyword is specified for an existing *sdp-id*, an error is generated and the context of the CLI is not 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 fails, generating an error message specifying the first bound service found during the deletion process. If the specified *sdp-id* does not exist, an error is generated.

Default n/a

Parameters *sdp-id* — the SDP identifier

Values 1 to 17407

gre — specifies that the SDP will use GRE encapsulation tunnels. Only one GRE SDP is supported to a given destination 7705 SAR or 7710/7750 SR.

mpls — specifies that the SDP will use MPLS encapsulation and one or more LSP tunnels to reach the far-end 7705 SAR or 7710/7750 SR. Multiple MPLS SDPs are supported to a given destination service router. Multiple MPLS SDPs to a single destination service router are helpful when they use divergent paths.

ip — specifies that the SDP will use IP encapsulation tunnels. Only one IP SDP is supported to a given destination 7705 SAR because the SDP is tied to the system address of the destination LER.

adv-mtu-override

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

Context config>service>sdp

Description This command overrides the advertised VC-type MTU. When enabled, the 7705 SAR signals a VC MTU equal to the service MTU that includes the Layer 2 header. Under normal operations it will advertise the service MTU minus the Layer 2 header. In the receive direction, it will accept either one.

The **no** form of this command disables the VC-type MTU override.

Default no adv-mtu-override

bgp-tunnel

Syntax [**no**] **bgp-tunnel**

Context config>service>sdp

Description This command allows the use of BGP route tunnels available in the tunnel table to reach SDP far-end nodes. BGP route tunnels are only available with MPLS SDPs. Only one transport method is allowed per SDP: LDP, RSVP-LSP or BGP-Tunnel.

The **no** form of the command disables the use of BGP route tunnels for the SDP far-end.

Default no bgp-tunnel

encryption-keygroup

Syntax **encryption-keygroup** *keygroup-id* **direction** {**inbound** | **outbound**}
no encryption-keygroup direction {**inbound** | **outbound**}

Context config>service>sdp

Description This command is used to bind a key group to an SDP for inbound or outbound packet processing. When configured in the outbound direction, packets egressing the node use the **active-outbound-sa** associated with the key group configured. When configured in the inbound direction, received packets must be encrypted using one of the valid security associations configured for the key group. Services using the SDP will be encrypted.

Encryption is enabled once the outbound direction is configured.

The **no** form of the command removes the key group from the SDP in the specified direction (inbound or outbound).

Default n/a

Parameters *keygroup-id* — the number of the key group being configured

Values 1 to 15 or *keygroup-name* (up to 64 characters)

direction {**inbound** | **outbound**} — mandatory keywords when binding a key group to a service for a particular direction

far-end

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

Context config>service>sdp

Description This command configures the system IP address of the far-end destination 7705 SAR, 7710 SR, 7750 SR, or other router ID platform for the SDP that is the termination point for a service.

The far-end IP address must be explicitly configured. The destination IP address must be a 7705 SAR, 7710 SR, 7750 SR, or other router ID platform system IP address.

If the SDP uses GRE or IP for the destination encapsulation, the local 7705 SAR might not know that the *ip-address* is actually a system IP interface address on the far-end service router.

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 message will be generated.

An SDP cannot be administratively enabled until a **far-end** *ip-address* is defined. The SDP is operational when it is administratively enabled (**no shutdown**).

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 message 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 n/a

Parameters *ip-address* — the system address of the far-end 7705 SAR for the SDP

Values a.b.c.d

ldp

Syntax [**no**] **ldp**

Context config>service>sdp

Description This command enables LDP-signaled LSPs 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, LDP must be disabled. The LSP must have already been created in the **config>router>mpls** context with a valid far-end IP address.

Default no ldp (disabled)

lsp

Syntax [**no**] **lsp** *lsp-name*

Context config>service>sdp

Description This command creates an association between an LSP and an MPLS SDP. This command is implemented only on MPLS-type 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, LDP must be disabled. The LSP must have already been created in the **config>router>mpls** context with a valid far-end IP address. Refer to the 7705 SAR OS MPLS Guide for CLI syntax and command usage.

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 shut down, causing the association with the SDP to be operationally down (the LSP will not be used by the SDP).

LSP SDPs also require that the T-LDP signaling be specified and that the SDP keepalive parameter be enabled and not timed out.

The **no** form of this command deletes an LSP association 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 No LSP names are defined.

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.

metric

Syntax **metric** *metric*
no metric

Context config>service>sdp

Description 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 to select the route with the lower value.

Parameters *metric* — specifies the SDP metric

Values 1 to 17407

path-mtu

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

Context config>service>sdp

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

The default SDP-type **path-mtu** can be overridden on a per-SDP basis.

Dynamic maintenance protocols on the SDP may override this setting.

If the physical **mtu** on an egress interface indicates that 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.

Parameters *bytes* — specifies the number of bytes in the path MTU

Values 576 to 9194

signaling

Syntax **signaling** {**off** | **tldp**}

Context config>service>sdp

Description This command specifies the signaling protocol used to obtain the ingress and egress 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 signaling 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 transport type, MPLS (LDP).

tldp — ingress and egress signaling auto-labeling is enabled

vlan-vc-etype

Syntax **vlan-vc-etype** *0x0600..0xffff*
no vlan-vc-etype [*0x0600..0xffff*]

Context config>service>sdp

Description This command configures the VLAN VC Ethertype. The **no** form of this command returns the value to the default. The etype value populates the Ethertype field in the Ethernet frame. It is used to indicate which protocol is being transported in the Ethernet frame. The default value indicates that the payload is an IEEE 802.1q-tagged frame.

Default no vlan-vc-etype *(0x8100)*

Parameters *0x0600..0xffff* — specifies a valid VLAN etype identifier

SDP Keepalive Commands

keep-alive

Syntax **keep-alive**

Context config>service>sdp

Description This command is the 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, non-error 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.

Table 13 describes keepalive interpretation of SDP Echo Reply response conditions and the effect on the SDP ID operational status.

Table 13: SDP Echo Reply Response Conditions

Result of Request	Stored Response State	Operational State
keepalive request timeout without reply	Request Timeout	Down
keepalive request not sent due to non-existent *orig-sdp-id* [1]	Orig-SDP Non-Existent	Down
keepalive request not sent due to administratively down *orig-sdp-id*	Orig-SDP Admin-Down	Down
keepalive reply received, invalid origination-id	Far End: Originator-ID Invalid	Down

Table 13: SDP Echo Reply Response Conditions (Continued)

Result of Request	Stored Response State	Operational State
keepalive reply received, invalid responder-id	Far End: Responder-ID Error	Down
keepalive reply received, No Error	Success	Up (if no other condition prevents)

Note:

1. This condition should not occur.

hello-time

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

Context config>service>sdp>keep-alive

Description This command 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.

Parameters *seconds* — the time period in seconds between SDP keepalive messages, expressed as a decimal integer

Values 1 to 3600

Default 10

hold-down-time

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

Context config>service>sdp>keep-alive

Description This command 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.

Parameters *seconds* — the time in seconds, expressed as a decimal integer, the *sdp-id* will remain in the operationally down state after an SDP keepalive error 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 to 3600

Default 10

max-drop-count

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

Context config>service>sdp>keep-alive

Description 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-drop-count** 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.

Parameters *count* — the number of consecutive SDP keepalive requests that can fail to be sent or replies missed before the SDP is brought down, expressed as a decimal integer

Values 1 to 5

Default 3

message-length

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

Context config>service>sdp>keep-alive

Description This command configures the size of SDP monitoring keepalive request messages transmitted on the SDP.

The **no** form of this command reverts the **message-length** *octets* value to the default setting.

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

The message length should be equal to the SDP operating path MTU as configured in the path-mtu 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.

Values 72 to 1500

Default 0

timeout

Syntax **timeout** *timeout*
no timeout

Context config>service>sdp>keep-alive

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

Parameters *timeout* — the timeout in seconds, expressed as a decimal integer

Values 1 to 10

Default 5

Show Commands

Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

customer

Syntax **customer** *customer-id*

Context show>service

Description This command displays service customer information.

Parameters *customer-id* — displays only information for the specified customer ID

Values 1 to 2147483647

Default all customer IDs display

Output The following output is an example of customer information, and Table 14 describes the fields.

Sample Output

```
*A:ALU-12# show service customer
============================================================
Customers
============================================================
Customer-ID : 1
Contact : Manager
Description : Default customer
Phone : (123) 555-1212

Customer-ID : 2
Contact : Tech Support
Description : ABC Networks
Phone : (234) 555-1212

Customer-ID : 3
Contact : Fred
Description : ABC Networks
Phone : (345) 555-1212

Customer-ID : 6
Contact : Ethel
Description : Epipe Customer
Phone : (456) 555-1212

Customer-ID : 7
Contact : Lucy
Description : VPLS 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:ALU-12#

*A:ALU-12# show service customer 274
===============================================================================
Customer 274
===============================================================================
Customer-ID : 274
Contact : Mssrs. Beaucoup
Description : ABC Company
Phone : 650 123-4567
-------------------------------------------------------------------------------
Total Customers : 1
-------------------------------------------------------------------------------
*A:ALU-12#
```

Table 14: Show Customer Command Output Fields

Label	Description
Customer-ID	Displays the unique customer identification number
Contact	Displays the name of the primary contact person
Description	Displays generic information about the customer
Phone	Displays the telephone or pager number used to reach the primary contact person
Total Customers	Displays the total number of customers configured

sdp

Syntax **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

Values 1 to 17407

Default all SDPs

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

keep-alive-history — displays the last fifty SDP keepalive events for the SDP

Output The following output is an example of service SDP information, and Table 15 describes the fields.

Sample Output

```
*A:ALU-12# show service sdp
===============================================================================
Services: Service Destination Points
===============================================================================
SdpId    Adm MTU   Opr MTU   IP address       Adm  Opr         Deliver Signal
-------------------------------------------------------------------------------
10       0         0         10.10.10.24      Up   Down        LDP     TLDP
20       0         0         10.10.10.24      Up   Down        MPLS    TLDP
30       4462      1514      10.20.1.21       Up   Up          GRE     TLDP
-------------------------------------------------------------------------------
Number of SDPs : 3
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-12#

*A:ALU-12# show service sdp 10
===============================================================================
Service Destination Point (Sdp Id : 10)
===============================================================================
SdpId    Adm MTU   Opr MTU   IP address       Adm  Opr         Deliver Signal
-------------------------------------------------------------------------------
10       0         0         10.10.10.24      Up   Down        LDP     TLDP
===============================================================================
*A:ALU-12#
```

```
*A:ALU-12# show service sdp 8 detail
===============================================================================
Service Destination Point (Sdp Id : 8) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 8  -(10.10.10.104)
-------------------------------------------------------------------------------
Description           : MPLS-10.10.10.104
SDP Id                : 8                     SDP-Source         : manual
Admin Path MTU        : 0                     Oper Path MTU      : 1550
Delivery              : MPLS
Far End               : 10.10.10.104          Delivery           : MPLS
Admin State           : Up                    Oper State         : Down
Signaling             : TLDP                  Metric             : 0
Last Status Change    : 02/01/2007 09:11:39   Adv. MTU Over.     : No
Last Mgmt Change      : 02/01/2007 09:11:46   VLAN VC Etype      : 0x8100
Flags                 : SignalingSessDown TransportTunnDown
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:ALU-12#

*A:ALU-12>show>service# sdp 5001 keep-alive-history
===============================================================================
Service Destination Point (Sdp Id : 5001)
===============================================================================
Time of Probe Report      RTT(ms) Size   Status
-------------------------------------------------------------------------------
2010/11/30 11:27:32       1011    72     Response Received
2010/11/30 11:27:22       1001    72     Response Received
2010/11/30 11:27:12       1001    72     Response Received
2010/11/30 11:27:02       1001    72     Response Received
2010/11/30 11:26:58       1002    72     Response Received
===============================================================================
*A:ALU-12>show>service#
```

Table 15: Show Service SDP Output Fields

Label	Description
SDP Id	Identifies the SDP
Description	Identifies the SDP by the text description stored in its configuration file
SDP Source	Specifies the SDP source type

Table 15: Show Service SDP Output Fields (Continued)

Label	Description
Adm MTU (Adm Path MTU)	Specifies the desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router
Opr MTU (Opr Path MTU)	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router
Delivery	Specifies the delivery routing protocol
Far End	Specifies the IP address of the remote end of the GRE or MPLS tunnel defined by this SDP
Adm (Admin State)	Specifies the desired state of the SDP
Opr (Oper State)	Specifies the operating state of the SDP
Deliver (Delivery)	Specifies the type of delivery used by the SDP: MPLS, GRE, or IP
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
Metric	Specifies the value used as a tie-breaker by the tunnel table manager to select a route
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
Adv. MTU Over	Specifies the state of the advertised VC-type MTU override command
VLAN VC Etype	Specifies the VLAN VC Ethertype for the SDP
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified
Keepalive Information:	
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 timer expired

Table 15: Show Service SDP Output Fields (Continued)

Label	Description
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state
TX Hello Msgs	Specifies the number of SDP echo request messages transmitted since the keepalive was administratively enabled or the counter was cleared
Rx Hello Msgs	Specifies the number of SDP echo request messages received since the keepalive was administratively enabled or the counter was cleared
Collect Stats.	Specifies that the collection of accounting and statistical data for the SDP is enabled or disabled
Associated LSP LIST:	
Note: If the SDP type is GRE, the following message displays: SDP Delivery Mechanism is not MPLS	
Lsp Name	For MPLS: identifies the name of the static LSP
Time since Last Trans*	For MPLS: specifies the time that the associated static LSP has been in service
Time of Probe Report	Indicates the date and time of the report
RTT (ms)	Indicates round-trip time (RTT), in milliseconds.
Size	Indicates the size of the packet, in bytes
Status	Indicates the status of the response
Group Encryption	
Inbound Keygroup Id	Indicates the key group used to decrypt inbound traffic for the service
Outbound Keygroup Id	Indicates the key group used to encrypt outbound traffic for the 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 to 17407

vc-id — the virtual circuit identifier

Values 1 to 4294967295

ip-address — displays only services matching with the specified far-end IP address

Default services with any far-end IP address

Output The following output is an example of service SDP-using information, and Table 16 describes the fields.

Sample Output

```
*A:ALU-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               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
102         300:102             Spok 10.0.0.13      Up        131067   131067
-------------------------------------------------------------------------------
Number of SDPs : 7
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-1#
```

Table 16: Show Service SDP-Using Output Fields

Label	Description
SvcId	Identifies the service
SdpId	Identifies the SDP
Type	Indicates the type of SDP (mesh or spoke)
Far End	Displays the far-end address of the SDP
Opr State	Displays the operational state of the service

Table 16: Show Service SDP-Using Output Fields (Continued)

Label	Description
I. Label	Displays the ingress label used by the far-end device to send packets to this device in this service by this SDP
E. Label	Displays the egress label used by this device to send packets to the far-end device in this service by this SDP

service-using

Syntax **service-using** [**epipe**] [**fpipe**] [**hpipe**] [**ies**] [**vpls**] [**vprn**] [**apipe**] [**ipipe**] [**cpipe**] [**sdp** *sdp-id*] [**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 services

vpls — displays matching VPLS services

vprn — displays matching VPRN services

apipe — displays matching Apipe services

cpipe — displays matching Cpipe services

fpipe — displays matching Fpipe services

hpipe — displays matching Hpipe services

sdp-id **—** displays only services bound to the specified SDP ID

Values 1 to 17407

Default services bound to any SDP ID

customer-id **—** displays services only associated with the specified customer ID

Values 1 to 2147483647

Default services associated with a customer

Output The following outputs are examples of service-using information, and Table 17 describes the fields.

Sample Output all services used in system

```
*A:ALU-12# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Service Name
-------------------------------------------------------------------------------
1            Cpipe     Down   Down       1                 cpipe_1_global
2            Apipe     Down   Down       1                 apipe_2
103          Epipe     Up     Up         104               epipe_103_sales
104          Epipe     Up     Up         104               epipe_104_marketing
105          Epipe     Up     Up         104               epipe_105_finance
303          Cpipe     Up     Up         104
304          Cpipe     Up     Up         104
305          Cpipe     Up     Up         104
701          Apipe     Up     Down       1
702          Apipe     Up     Down       1
703          Apipe     Up     Down       1
704          Apipe     Up     Down       1
807          Apipe     Up     Down       1
808          Apipe     Up     Down       1
903          Cpipe     Up     Up         1
904          Cpipe     Up     Up         1
5000         VPLS      Down   Down       1                  vpls_5000_sales
5001         VPLS      Down   Down       1
-------------------------------------------------------------------------------
Matching Services : 19
```

Sample Output services used by customer

```
*A:ALU-12# show service service-using customer 1
===============================================================================
Services  Customer 1
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Service Name
-------------------------------------------------------------------------------
1            Cpipe     Down   Down       1                 cpipe_1_global
2            Apipe     Down   Down       1                 apipe_2
701          Apipe     Up     Down       1
702          Apipe     Up     Down       1
703          Apipe     Up     Down       1
704          Apipe     Up     Down       1
807          Apipe     Up     Down       1
808          Apipe     Up     Down       1
903          Cpipe     Up     Up         1
904          Cpipe     Up     Up         1
5000         VPLS      Down   Down       1                  vpls_5000_sales
5001         VPLS      Down   Down       1
-------------------------------------------------------------------------------
Matching Services : 13
*A:ALU-12#
```

Sample Output services by service type (epipe)

```
*A:ALU-12# show service service-using epipe
===============================================================================
Services [epipe]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Service Name
-------------------------------------------------------------------------------
103          Epipe     Up     Up         104               epipe_103_sales
104          Epipe     Up     Up         104               epipe_104_marketing
105          Epipe     Up     Up         104               epipe_105_finance
-------------------------------------------------------------------------------
Matching Services : 3
*A:ALU-12#
```

Table 17: Show Service service-using Output Fields

Label	Description
Service Id	Identifies the service
Type	Specifies the service type configured for the service ID
Adm	Displays the desired state of the service
Opr	Displays the operating state of the service
CustomerID	Displays the ID of the customer who owns this service
Service Name	The service name

VLL Services

In This Chapter

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

Topics in this chapter include:

- ATM VLL (Apipe) Services
- Circuit Emulation VLL (Cpipe) Services
- Ethernet VLL (Epipe) Services
- Frame Relay VLL (Fpipe) Services
- HDLC VLL (Hpipe) Services
- IP Interworking VLL (Ipipe) Services
- Pseudowire Switching
- VLL Service Considerations
- Configuring a VLL Service with CLI
- VLL Services Command Reference

ATM VLL (Apipe) Services

This section provides information about the Apipe service. Topics in this section include:

- ATM VLL for End-to-End ATM Service
- ATM Virtual Trunk Over an IP/MPLS Packet-Switched Network
- ATM SAP-to-SAP Service
- ATM Traffic Management Support
- Control Word

Apipe configuration information is found under the following topics:

- Common Configuration Tasks
- Configuring VLL Components
- Service Management Tasks

See Service Support for information on the adapter cards and chassis that support ATM VLL services.

ATM VLL for End-to-End ATM Service

ATM VLLs (Apipe) provide a point-to-point ATM service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network (see Figure 27). User ATM traffic is connected to a 7705 SAR either directly or through an ATM access network. In both cases, an ATM PVC—for example, a virtual channel (VC) or a virtual path (VP)—is configured on the 7705 SAR. VPI/VCI translation is supported in the ATM VLL.

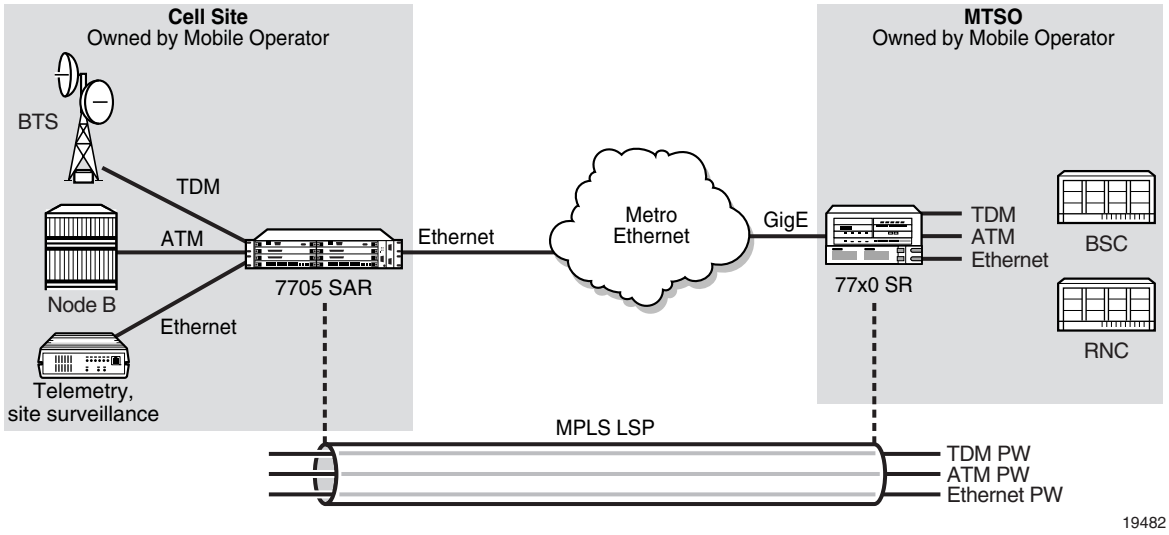
The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

The 7705 SAR receives standard UNI/NNI cells on the ATM SAP, or on a number of SAPs belonging to a SAP aggregation group, which are then encapsulated into a pseudowire packet using N-to-1 cell mode encapsulation in accordance with RFC 4717. See ATM PWE3 N-to-1 Cell Mode Encapsulation for more information about N-to-1 cell mode encapsulation.

In addition to supporting N-to-1 cell mode encapsulation, ATM VLL service supports cell concatenation, control word (CW), SAP-to-SAP (local service), and SAP-to-SDP binding (distributed service). See SAP Encapsulations and Pseudowire Types for more information. ATM SAP-to-SAP service is not supported when N > 1; see ATM SAP-to-SAP Service for information about ATM SAP-to-SAP services.

ATM VLL optimizes the ATM cell from a 53-byte cell to a 52-byte packet by removing the header error control (HEC) byte at the near end. The far end regenerates the HEC before switching ATM traffic to the attached circuit.

Figure 27: ATM VLL for End-to-End ATM Service



ATM Virtual Trunk Over an IP/MPLS Packet-Switched Network

ATM virtual trunks (VT), also known as ATM transparent cell transport in RFC 4446, provide a transparent trunking function for user and control traffic between two ATM switches over an ATM pseudowire. ATM control traffic includes PNNI signaling and routing traffic, ILMI traffic, and F4/F5 OAM traffic. Figure 28 shows two switches labeled ATM 2 and ATM 4 that appears to be directly connected over an ATM link because the virtual trunk emulates a direct connection between the ATM switches.

Virtual trunks are supported on 7705 SAR 4-port OC3/STM1 Clear Channel Adapter cards with ports configured for OC3 or STM1 and 4-port DS3/E3 Adapter cards with ports configured for DS3 and E3.

All cells arriving on a port configured for virtual trunking on the 7705 SAR are fed to a single ATM pseudowire. The VPI range cannot be configured. No VPI/VCI translation is performed on ingress or egress. Cell order is maintained within a VT. SAP to SAP service is supported. The two ATM ports can therefore be on the same PE node.

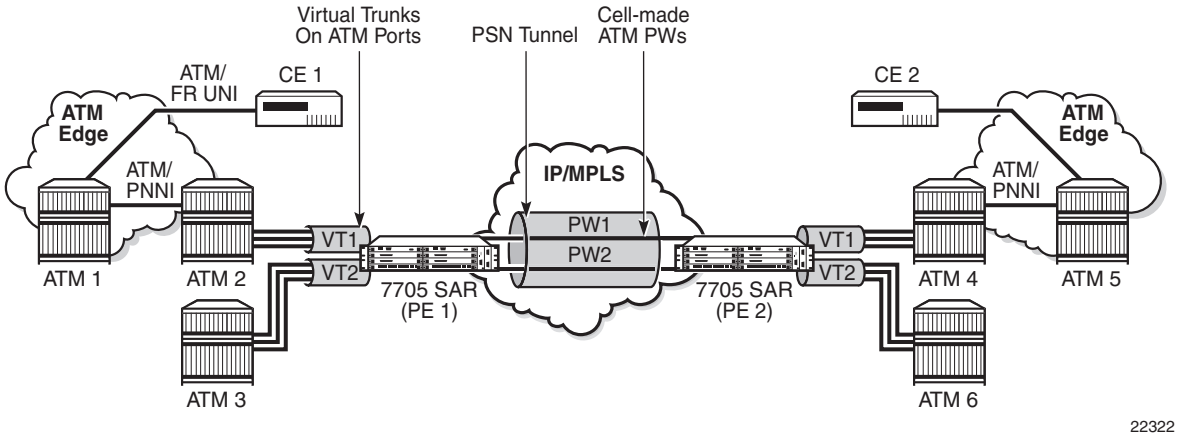
By carrying all cells from all VPIs making up the VT in one pseudowire, a transport solution is provided that is both robust and operationally efficient since the entire VT can be managed as a single entity from the 5620 SAM Network Manager (single point for configuration, status, alarms, statistics).

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

The ATM pseudowire is initiated using targeted LDP (T-LDP) signaling (defined in draft-ietfpwe3-control-protocol-xx.txt, *Pseudowire Setup and Maintenance using LDP*). In this application, there is no ATM signaling on the 7705 SAR gateway nodes since both endpoints of the MPLS network are configured by the network operator. ATM signaling between the ATM nodes is passed transparently over the VT (along with user traffic) from one ATM port on a 7705 SAR PE to another ATM port on a remote (or the same) 7705 SAR PE.

OAM signaling behaves in the same way as user traffic in that OAM cells are transported transparently and no special action is taken when F4 or F5 OAM cells are received at the SAP ingress or egress. As well, all ILMI messages are transported transparently between two endpoints. In the case of a pseudowire failure (for example, label withdrawal, loss of reachability, loss of tunnel to the eLER), the virtual trunk service SAP port is not taken down.

Figure 28: ATM Virtual Trunk Over IP/ MPLS Packet-Switched Network



ATM SAP-to-SAP Service

ATM VLLs can be configured with both endpoints (SAPs) on the same 7705 SAR. This is referred to as an ATM SAP-to-SAP or local ATM service. An ATM SAP-to-SAP service emulates local ATM switching between two ATM endpoints on the 7705 SAR. Both ingress and egress traffic is legacy ATM traffic. The two SAPs can be any combination of supported ATM encapsulation ports, channel-groups, or IMA bundles.

Note: ATM SAP-to-SAP connections are supported on any combination of the following:

- IMA group
- T1/E1 ASAP port
- E3/DS3 port
- OC3/STM1 port. The OC3 port can be configured for clear channel or channelized service.

Note: ATM SAP-to-SAP connections are not supported for pseudowire packets using N-to-1 cell mode encapsulation where N > 1.

ATM Traffic Management Support

The 7705 SAR supports the ATM Forum Traffic Management Specification Version 4.1.

Topics in this section include:

- Network Ingress Classification
- ATM Access Egress Queuing and Shaping

Network Ingress Classification

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

The ingress MPLS packets are mapped to forwarding classes based on EXP bits that are part of the headers in the MPLS packets. The EXP bits are used to ensure an end-to-end QoS application. For PW services, there are two labels: one for the MPLS tunnel and one for the pseudowire itself. Mapping is done according to the outer tunnel EXP bit settings. This ensures that if the EXP bit settings are altered along the path by the intermediate LSR nodes, the newly requested FC selection is carried out properly.

Ingress GRE and IP packets are mapped to forwarding classes based on DSCP bit settings of the IP header.

ATM Access Egress Queuing and Shaping

The 7705 SAR provides a per-SAP queuing architecture on the 16-port T1/E1 ASAP Adapter card, 32-port T1/E1 ASAP Adapter card, 4-port DS3/E3 Adapter card, and 4-port OC3/STM1 Clear Channel Adapter card. After the ATM pseudowire is terminated at the access egress point, all the ATM cells are mapped to default queue 1, and queuing is performed on a per-SAP basis.

Access ingress and access egress traffic management features are identical for SAP-to-SAP and SAP-to-SDP applications.

Note: The ATM access egress shaping configuration in a SAP egress QoS policy is ignored when that policy is assigned to an ATM SAP. The shaping of the egress cell stream is controlled by the `atm-td-profile` command. If the `atm-td-profile` is not configured, the default `atm-td-profile` is in effect.

For more information on ATM access egress queuing and scheduling, refer to the 7705 SAR OS Quality of Service Guide.

Control Word

ATM VLL supports an optional control word (CW). The control word must be used for any pseudowire service that transports packets that are less than 64 bytes; see the Pseudowire Control Word and the control-word command description for more information.

Circuit Emulation VLL (Cpipe) Services

This section provides information about the Cpipe service.

Topics in this section include:

- Cpipe Service Overview
- TDM SAP-to-SAP Service
- Cpipe Service Modes
- TDM PW Encapsulation
- Circuit Emulation Parameters and Options
- Error Situations

Cpipe configuration information is found under the following topics:

- Common Configuration Tasks
- Configuring VLL Components
- Service Management Tasks

See Service Support for information on the adapter cards and chassis that support circuit emulation VLL services.

Cpipe Service Overview

Cpipe service is the Alcatel-Lucent implementation of TDM PW VLL as defined in the IETF PWE3 working group.

The 7705 SAR can support TDM circuit applications that are able to transport delay-sensitive TDM traffic over a packet network. For example, in the case of cell site aggregation, Cpipe services provide transport service for 2G connectivity between the base transceiver station and the base station controller, and for 3G backhaul applications (for example, EVDO traffic from T1/E1 ports with MLPPP). Cpipe services over MPLS or GRE tunnels are supported.

The 2G traffic is transported encapsulated in a TDM VLL over the packet switched network (PSN). The entire T1/E1 frame or part of a frame (n × 64 kb/s) is carried as a TDM VLL over the PSN. At the far end, the transport layer frame structure is regenerated when structured circuit emulation is used, or simply forwarded as part of the payload when unstructured circuit emulation is used. The 3G UMTS R99 traffic uses ATM/IMA as the transport protocol. The IMA sessions are terminated at the site by the 7705 SAR and the 3G ATM traffic is transported across the PSN through the use of ATM VLLs (PWE3).

TDM SAP-to-SAP Service

TDM VLLs can be configured with both endpoints (SAPs) on the same 7705 SAR. This is referred to as TDM SAP-to-SAP or local TDM service. TDM SAP-to-SAP emulates a TDM multiplexing and switching function on the 7705 SAR.

A TDM SAP-to-SAP connection is set up in the 7705 SAR and a pseudowire is configured between the two endpoints.

Note: TDM SAP-to-SAP connections are supported between any T1/E1 port or channel that is configured for access mode and circuit emulation service and another port or channel with the same configuration (endpoint type, bit rate [number of DS0s in a channel group], payload size, CAS enabled/disabled, and RTP enabled/disabled).

Cpipe Service Modes

Cpipe services support unstructured circuit emulation mode (SAToP) for DS3, E3, DS1, and E1 circuits as per RFC 4553 and structured circuit emulation mode (CESoPSN) for n × 64 kb/s timeslots in DS1 and E1 circuits as per RFC 5086.

The 7705 SAR supports MEF 8, which allows both structured and unstructured emulation of TDM services across Epipes, also known as Circuit Emulation Services over Ethernet (CESoETH). See MEF 8 for more information about CESoETH.

Topics in this section include:

- Unstructured Mode (SAToP)
- Structured Mode (CESoPSN)

Unstructured Mode (SAToP)

Structure-agnostic TDM over Packet (SAToP) is an unstructured circuit emulation mode used for the transport of unstructured TDM or structured TDM (where the structure is ignored).

Note: The word “agnostic” is used in RFC 4553, but it is not used in the literal sense. The meaning of agnostic in this case is “unaware or independent”; therefore, structure-agnostic is used to mean structure-unaware or structure-independent.

As a structure-unaware or structure-independent service, SAToP service does not align to any framing; the framing mode for the port is set to unframed. For structured TDM, SAToP disregards the bit sequence and TDM structure in order to transport the entire signal over a PSN as a pseudowire.

Structured Mode (CESoPSN)

Structure-aware circuit emulation is used for the transport of structured TDM, taking at least some level of the structure into account. By selecting only the necessary n × 64 kb/s timeslots to transport, bandwidth usage is reduced or optimized (compared to a full DS1 or E1). Full DS1s or E1s can be transported by selecting all the timeslots in the DS1 or E1 circuit. Framing bits (DS1) or FAS (E1) are terminated at the near end and reproduced at the far end.

The 7705 SAR supports CESoPSN with and without CAS for DS1 and E1.

When CESoPSN with CAS is selected, the ABCD bits are coded into the T1 or E1 multiframe packets, transported within the TDM PW, and reconstructed in the T1 or E1 multiframe at the far end for each timeslot.

Channel Associated Signaling (CAS) includes four signaling bits (A, B, C, and D) in the messages sent over a voice trunk. These messages provide information such as the dialed digits and the call state (whether on-hook or off-hook).

The mechanism for E1 CAS is described in ITU-T G.732. When configured for E1 CAS, timeslot 17 carries the signaling information for the timeslots used for voice trunking. Each channel requires four signaling bits, so grouping 16 E1 frames into a multiframe allows the signaling bits for all 30 channels to be trunked.

As shown in Figure 29, timeslot 1 of all frames within the E1 multiframe is reserved for alignment, alarm indication, and CRC. For Frame 0, timeslot 17 is reserved for multiframe alignment bits. For the remaining 15 frames, timeslot 17 contains ABCD bits for two channels.

Note: For E1 CAS, timeslots are numbered 1 to 32 on the 7705 SAR.

For T1 CAS, the signaling bits are transferred using Robbed Bit Signaling (RBS), where the least significant bit in the channel is used periodically to transport these bits instead of voice data. T1 CAS is supported when ESF or SF framing is configured. ESF framing uses a 24-frame multiframe and transfers all four signaling bits (ABCD). SF framing uses a 12-frame multiframe and transfers only the AB bits. The signaling bits are carried in the least significant bit of the following frames:

- A bit in frame 6
- B bit in frame 12
- C bit in frame 18
- D bit in frame 24

When CESoPSN with CAS is selected, the ABCD bits are decoded from the incoming E1/T1 multiframe, transferred within the TDM PW, and reconstructed in the E1/T1 multiframe at the far end for every DS0 channel. CAS can be configured on a per-T1/E1 basis or on a per-DS0/64 kb/s channel basis.

Figure 29: E1 Framing for CAS Support in an E1 Multiframe

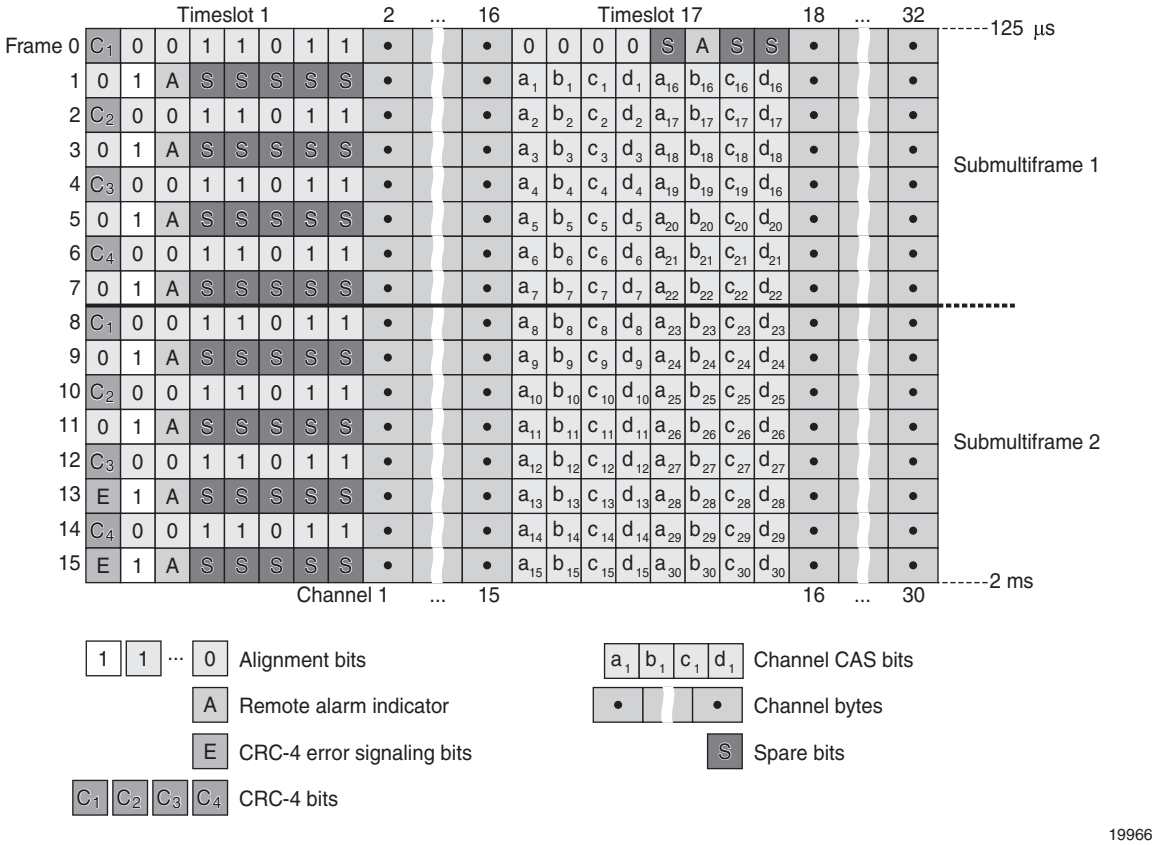


Table 18 shows the structure of a T1 ESF multiframe that uses RBS. The structure of a T1 SF multiframe is based on 12 frames and only the A and B bits are available.

Table 18: T1 Framing for CAS (RBS) Support in a T1 ESF Multiframe

Frame Number	F Bit				Bit Numbers in Each Channel Timeslot		Signaling Channel Designation [4]
	Bit Number within Multiframe	Assignments					
		FAS [1]	DL [2]	CRC [3]	For Character Signal [4]	For Signaling [4]	
1	1	—	m	—	1-8	—	—
2	194	—	—	e1	1-8	—	—
3	387	—	m	—	1-8	—	—
4	580	0	—	—	1-8	—	—
5	773	—	m	—	1-8	—	—
6	966	—	—	e2	1-7	8	A
7	1159	—	m	—	1-8	—	—
8	1352	0	—	—	1-8	—	—
9	1545	—	m	—	1-8	—	—
10	1738	—	—	e3	1-8	—	—
11	1931	—	m	—	1-8	—	—
12	2124	1	—	—	1-7	8	B
13	2317	—	m	—	1-8	—	—
14	2510	—	—	e4	1-8	—	—
15	2703	—	m	—	1-8	—	—
16	2896	0	—	—	1-8	—	—
17	3089	—	m	—	1-8	—	—
18	3282	—	—	e5	1-7	8	C
19	3475	—	m	—	1-8	—	—
20	3668	1	—	—	1-8	—	—
21	3861	—	m	—	1-8	—	—

Table 18: T1 Framing for CAS (RBS) Support in a T1 ESF Multiframe (Continued)

Frame Number	F Bit				Bit Numbers in Each Channel Timeslot		Signaling Channel Designation [4]
	Bit Number within Multiframe	Assignments					
		FAS [1]	DL [2]	CRC [3]	For Character Signal [4]	For Signaling [4]	
22	4054	—	—	e6	1-8	—	—
23	4247	—	m	—	1-8	—	—
24	4440	1	—	—	1-7	8	D

Notes:

1. FAS = frame alignment signal (....001011.....)
2. DL = 4 kb/s data link (m represents message bits)
3. CRC = CRC-6 block check field (e1 to e6 represent check bits)
4. Only applicable for CAS

TDM PW Encapsulation

TDM circuits are MPLS-encapsulated as per RFC 4553 (SAToP) and RFC 5086 (CESoPSN). (see Figure 30 and Figure 31).

For GRE tunnels, the same encapsulations shown in Figure 31 are used, but GRE tunnel headers are used instead of MPLS tunnel headers.

Figure 30: SAToP MPLS Encapsulation

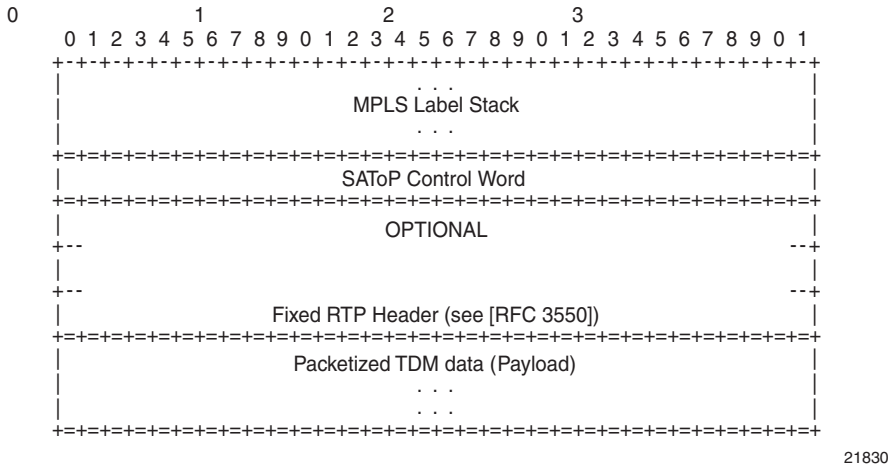


Figure 31: CESoPSN MPLS Encapsulation

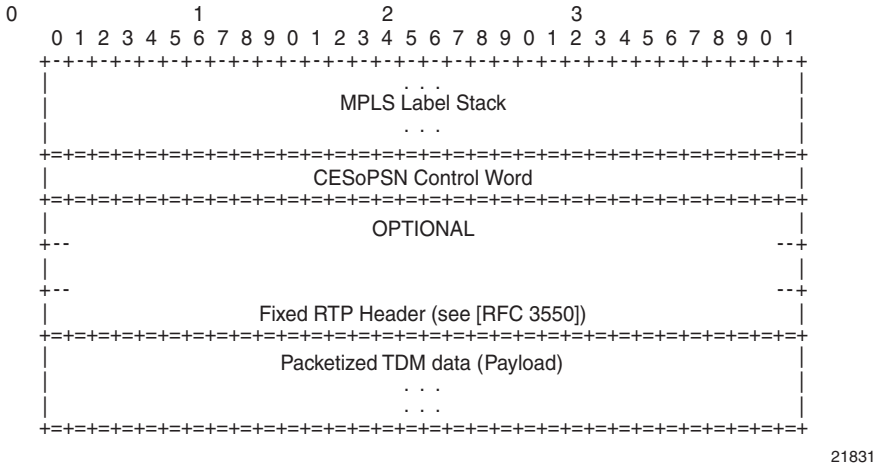
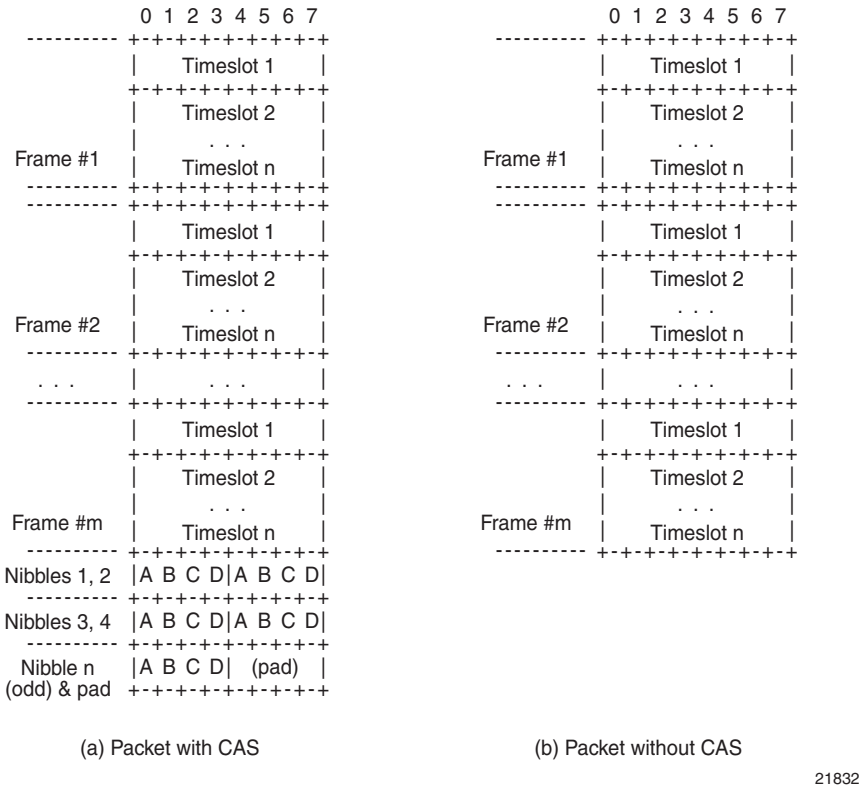


Figure 32 shows the format of the CESoPSN TDM payload (with and without CAS) for packets carrying trunk-specific n × 64 kb/s service.

Figure 32: CESoPSN Packet Payload Format for Trunk-Specific n x 64 kb/s (With and Without CAS Transport)



For CESoPSN without CAS, select the packet size so that an integer number of frames are transported. That is, if *n* timeslots per frame are to be encapsulated in a TDM PW, then the packet size must be a multiple of *n* (where *n* is not equal to 1). For example, if *n* = 4 timeslots, then the packet size can be 8, 12, 16, and so on. If only one timeslot per frame is being transported, the packet size must be an even number.

For CESoPSN with CAS, the packet size is an integer number of frames. A single T1/E1 port can have a mix of CAS and non-CAS traffic in each DS0/64 kb/s channel. You must configure the relevant T1/E1 port or channel group for CAS signal mode before provisioning the TDM PW with CAS or the system will disallow the signal mode configuration. The extra bytes for ABCD (CAS) signaling bits are not included when setting the packet size.

For a single T1/E1 port that contains a mix of CAS and non-CAS signaling, all the non-CAS channel Cpipes inherit the CAS channel restriction concerning 24/16 frames for payload size. For a T1 port, the payload size is equal to the number of CAS and/or non-CAS timeslots × 24 frames/multiframe × *n* multiframes, where *n* = 1 to 8. For an E1 port, the payload size is equal to the number of CAS and/or non-CAS timeslots × 16 frames/multiframe × *n* multiframes, where *n* =1 to 8.

Note: The extra bytes for CAS signaling bits must be included when setting the `service-mtu` size. See Structured T1/E1 CES with CAS for more information.

Circuit Emulation Parameters and Options

Cpipe services support unstructured circuit emulation mode (SAToP) for DS3, E3, DS1, and E1 circuits as per RFC 4553, and structured circuit emulation mode (CESoPSN) for n × 64 kb/s timeslots in DS1 and E1 circuits as per RFC 5086.

Table 19 lists the adapter cards, modules, and chassis that support SAToP and CESoPSN.

Table 19: SAToP and CESoPSN Support on the 7705 SAR

Card/Module/Chassis		SAToP	CESoPSN
Name	**CLI identifier (includes mode and channelization)**		
16-port T1/E1 ASAP Adapter Card	a16-chds1		
	On DS0/64k		✓
	On DS1/E1	✓	
16-port T1/E1 ASAP Adapter Card version 2	a16-chds1v2		
	On DS0/64k		✓
	On DS1/E1	✓	
32-port T1/E1 ASAP Adapter Card	a32-chds1v2		
	On DS0/64k		✓
	On DS1/E1	✓	

Table 19: SAToP and CESoPSN Support on the 7705 SAR (Continued)

Card/Module/Chassis		SAToP	CESoPSN
Name	**CLI identifier (includes mode and channelization)**		
2-port OC3/STM1 Channelized Adapter Card	a2-choc3		
	On DS0/64k		✓
	On DS1/E1	✓	
	On DS3 [1]	✓	
4-port OC3/STM1 / 1-port OC12/STM4 Adapter Card	a4-choc3/12		
	On DS1/E1	✓	
4-port DS3/E3 Adapter Card	a4-chds3		
	On n × DS0		✓
	On DS1 [2]	✓	
	On DS3/E3	✓	
12-port Serial Data Interface Card	a12-sdi		
	On V.35 and X.21 ports		✓
	On RS-232 ports		✓
6-port E&M Adapter Card	a6-em		✓
8-port Voice & Teleprotection Card	a8-vt		✓
8-port FXO Adapter Card	a8-fxo		✓
6-port FXS Adapter Card	a6-fxs		✓
Integrated Services Card	isc		✓ [3]
4-port T1/E1 and RS-232 Combination Module	a4-combo		
	On RS-232 ports		✓
	On T1/E1 ports	✓	✓
7705 SAR-A	i8-chds1	✓	✓
7705 SAR-Hc	i2-sdi		✓
7705 SAR-M	i16-chds1	✓	✓

Table 19: SAToP and CESoPSN Support on the 7705 SAR (Continued)

Card/Module/Chassis		SAToP	CESoPSN
Name	**CLI identifier (includes mode and channelization)**		
7705 SAR-X	i8-chds1-x	✓	✓

Notes:

1. No support for E3
2. No support for E1/64K channel groups
3. The Integrated Services card acts as a Cpipe bridge/multicaster. There are no physical ports on the Integrated Services card.

See Service Support for further information on circuit emulation VLL services.

The following parameters and options are described in this section:

- Unstructured
- Structured DS1/E1 CES without CAS
- Structured T1/E1 CES with CAS
- Packet Payload Size
- Jitter Buffer
- Asymmetric Delay Control
- RTP Header
- Control Word

Unstructured

Unstructured CES is configured by choosing `satop-t1`, `satop-e1`, `satop-t3`, or `satop-e3` as the `vc-type` when creating a Cpipe service. For DS1 and E1 unstructured circuit emulation, the framing parameter of the port must be set to `ds1-unframed` and `e1-unframed` (respectively) because SAToP service ignores the underlying framing. Additionally, channel group 1 must contain all 24 or 32 timeslots, which is configured automatically when channel group 1 is created.

For DS1, E1, DS3, and E3 circuit emulation, the payload packet size is configurable and must be an integer value between 64 and 1514 octets and must be a multiple of 32. The payload packet size affects the packet efficiency and packetization delay. Table 20 shows the default values for packet size and packetization delay. See Packet Payload Size for more information.

Note: When using SAToP to transport DS1 traffic, the framing bit (bit 193) in the DS1 overhead is included and packed in the payload and sent over the PSN. If the underlying framing is ESF, then the Facility Data Link (FDL) channel is transported over the Cpipe as part of the SAToP service. No matter the case, the framing parameter of the port must be set to unframed.

Table 20: Unstructured Payload Defaults

Circuit	Payload Size (Octets)	Packetization Delay (ms)
DS1	192	1.00
E1	256	1.00
DS3	1024	0.183
E3	1024	0.238

Structured DS1/E1 CES without CAS

Structured CES without CAS is configured by choosing `cesopsn` as the `vc-type` when creating a Cpipe service. For n × 64 kb/s structured circuit emulation operation, the framing parameter of the port must be set to a framed setting (such as ESF for DS1). Each channel group contains *n* DS0s (timeslots), where *n* is between 1 and 24 timeslots for DS1 and between 1 and 31 timeslots for E1.

The packet payload size is configurable (in octets) and must be an integer multiple of the number of timeslots in the channel group. The minimum payload packet size is 2 octets (based on two frames per packet and one timeslot per frame). See Table 21 for default and minimum payload size values. The maximum payload packet size is 1514 octets.

If a port on a 16-port T1/E1 ASAP Adapter card, version 2, or 32-port T1/E1 ASAP Adapter card is configured for DCR, the port timing is associated with the service clock of the Cpipe of channel group 1. For a framed T1 port, there is a restriction on the Cpipe payload size of channel group 1:

- for DCR with a timestamp frequency of 77.76 MHz, the payload size = 2 x I x (number of timeslots), where I = 1 to 20
- for DCR with a timestamp frequency of 19.44 MHz, the payload size = 8 x I x (number of timeslots), where I = 1 to 5

This restriction does not apply to framed E1 ports or unframed T1/E1 ports.

Each DS1 or E1 frame contributes a number of octets to the packet payload. That number is equal to the number of timeslots configured in the channel group. Thus, a channel group with four timeslots contributes 4 octets to the payload. The timeslots do not need to be contiguous.

Note that a smaller packet size results in a lower packetization delay; however, it increases the packet overhead (when expressed as a percentage of the traffic).

Calculation of Payload Size

The payload size (S), in octets, can be calculated using the following formula:

S = N x F

where:

N = the number of octets (timeslots) collected per received frame (DS1 or E1)
F = the number of received frames (DS1 or E1) that are accumulated in each CESoPSN packet

For example, assume the packet collects 16 frames (F) and the channel group contains 4 octets (timeslots) (N). Then the packet payload size (S) is:

S = 4 octets/frame x 16 frames

= 64 octets

Calculation of Packetization Delay

Packetization delay is the time needed to collect the payload for a CESoPSN packet. DS1 and E1 frames arrive at a rate of 8000 frames per second. Therefore, the received frame arrival period is 125 µs.

In the previous example, 16 frames were accumulated in the CESoPSN packet. In this case, the packetization delay (D) can be calculated as follows:

D = 125 µs/frame × 16 frames

= 2.000 ms

Table 21 shows the default and minimum values for frames per packet, payload size, and packetization delay as they apply to the number of timeslots (N) that contribute to the packet payload. The default values are set by the operating system as follows:

- for N = 1, the default is 64 frames/packet
- for 2 ≤ N ≤ 4, the default is 32 frames/packet

- for $5 \leq N \leq 15$, the default is 16 frames/packet
- for $N \geq 16$, the default is 8 frames/packet

Table 21: Default and Minimum Payload Size for CESoPSN without CAS

	Default Values			Minimum Values		
Number of Timeslots (N)	**Frames per Packet (F)**	**Payload Size (Octets) (S)**	**Packetization Delay (ms) (D)**	**Frames per Packet (F)**	**Payload Size (Octets) (S)**	**Packetization Delay (ms) (D)**
1	64	64	8.000	2	2	0.250
2	32	64	4.000	2	4	0.250
3	32	96	4.000	2	6	0.250
4	32	128	4.000	2	8	0.250
5	16	80	2.000	2	10	0.250
6	16	96	2.000	2	12	0.250
7	16	112	2.000	2	14	0.250
8	16	128	2.000	2	16	0.250
9	16	144	2.000	2	18	0.250
10	16	160	2.000	2	20	0.250
11	16	176	2.000	2	22	0.250
12	16	192	2.000	2	24	0.250
13	16	208	2.000	2	26	0.250
14	16	224	2.000	2	28	0.250
15	16	240	2.000	2	30	0.250
16	8	128	1.000	2	32	0.250
17	8	136	1.000	2	34	0.250
18	8	144	1.000	2	36	0.250
19	8	152	1.000	2	38	0.250
20	8	160	1.000	2	40	0.250
21	8	168	1.000	2	42	0.250
22	8	176	1.000	2	44	0.250

Table 21: Default and Minimum Payload Size for CESoPSN without CAS (Continued)

	Default Values			Minimum Values		
Number of Timeslots (N)	Frames per Packet (F)	Payload Size (Octets) (S)	Packetization Delay (ms) (D)	Frames per Packet (F)	Payload Size (Octets) (S)	Packetization Delay (ms) (D)
23	8	184	1.000	2	46	0.250
24	8	192	1.000	2	48	0.250
25	8	200	1.000	2	50	0.250
26	8	208	1.000	2	52	0.250
27	8	216	1.000	2	54	0.250
28	8	224	1.000	2	56	0.250
29	8	232	1.000	2	58	0.250
30	8	240	1.000	2	60	0.250
31	8	248	1.000	2	62	0.250

Structured T1/E1 CES with CAS

Structured circuit emulation with CAS is supported for T1 and E1 circuits.

Structured CES with CAS service is configured by choosing `cesopsn-cas` as the `vc-type` when creating a Cpipe service. The DS1 or E1 service on the port associated with the Cpipe SAP should be configured to support CAS (via the `signal-mode {cas}` command) before configuring the Cpipe service to support DS1 or E1 with CAS. Refer to the 7705 SAR OS Interface Configuration Guide for information on configuring signal mode.

For $n \times$ DS0 and $n \times 64$ kb/s structured circuit emulation with CAS, the implementation is almost identical to that of CES without CAS. When CAS operation is enabled, timeslot 16 (channel 17) cannot be included in the channel group on E1 carriers. Since the CAS in-band method is used, separate PW support for CAS is not provided.

When CAS is enabled, the packet size is based on the number of multiframes per packet and whether the circuit is configured for E1 or T1. Payload size is user-configurable to correspond to the desired integer number of multiframes. The 7705 SAR supports up to 8 multiframes, where a multiframe contains 24 frames for T1 and 16 frames for E1. Therefore, the payload size = number of timeslots $\times$ 24 (T1) or 16 (E1) frames per multiframe $\times$ number of multiframes. For example, the payload size for a T1 line (24 frames) using 2 timeslots and 8 multiframes is 384 bytes ($384 = (2 \times 24) \times 8$).

Multiple multiframes are supported on the following cards and platforms:

- 6-port E&M Adapter card (see note below)
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- 7705 SAR-A (variant with T1/E1 ports)
- 7705 SAR-F
- 7705 SAR-M (variants with T1/E1 ports)
- 7705 SAR-X

Note: If the 6-port E&M Adapter card is configured for A-Law companding, the E&M ports support 16 (E1) frames per multiframe. If the card is configured for Mu-Law companding, the ports support 24 (T1) frames per multiframe.

Table 22 shows the default payload sizes based on the number of timeslots.

For CAS, the signaling portion adds (n/2) bytes (n is an even integer) or ((n+1)/2) bytes (n is odd) to the packet, where n is the number of timeslots in the channel group. You do not include the additional signaling bytes when setting the TDM payload size. However, the operating system includes the additional bytes in the total packet payload, and the total payload must be accounted for when setting the `service-mtu` size. Continuing the example above, since n = 4, the total payload is 64 octets plus (4/2 = 2) CAS octets, or 66 octets. Refer to Figure 32 to see the structure of the CES with CAS payload.

CES fragmentation is not supported.

Note: If you configure the `service-mtu` size to be smaller than the total payload size (payload plus CAS bytes), then the Cpipe will not become operational. This must be considered if you change the `service-mtu` from its default value.

Table 22: Default Values for the Payload Size for T1 and E1 CESoPSN with CAS

Number of Timeslots	T1			E1		
	Number of Frames per Packet	Payload Size (Octets)	Packetization Delay (ms)	Number of Frames per Packet	Payload Size (Octets)	Packetization Delay (ms)
1	24	24	3.00	16	16	2.00
2	24	48	3.00	16	32	2.00
3	24	72	3.00	16	48	2.00

Table 22: Default Values for the Payload Size for T1 and E1 CESoPSN with CAS (Continued)

Number of Timeslots	T1			E1		
	Number of Frames per Packet	Payload Size (Octets)	Packetization Delay (ms)	Number of Frames per Packet	Payload Size (Octets)	Packetization Delay (ms)
4	24	96	3.00	16	64	2.00
5	24	120	3.00	16	80	2.00
6	24	144	3.00	16	96	2.00
7	24	168	3.00	16	112	2.00
8	24	192	3.00	16	128	2.00
9	24	216	3.00	16	144	2.00
10	24	240	3.00	16	160	2.00
11	24	264	3.00	16	176	2.00
12	24	288	3.00	16	192	2.00
13	24	312	3.00	16	208	2.00
14	24	336	3.00	16	224	2.00
15	24	360	3.00	16	240	2.00
16	24	384	3.00	16	256	2.00
17	24	408	3.00	16	272	2.00
18	24	432	3.00	16	288	2.00
19	24	456	3.00	16	304	2.00
20	24	480	3.00	16	320	2.00
21	24	504	3.00	16	336	2.00
22	24	528	3.00	16	352	2.00
23	24	552	3.00	16	368	2.00
24	24	576	3.00	16	384	2.00
25	NA	NA	NA	16	400	2.00
26	NA	NA	NA	16	416	2.00
27	NA	NA	NA	16	432	2.00

Table 22: Default Values for the Payload Size for T1 and E1 CESoPSN with CAS (Continued)

Number of Timeslots	T1			E1		
	Number of Frames per Packet	Payload Size (Octets)	Packetization Delay (ms)	Number of Frames per Packet	Payload Size (Octets)	Packetization Delay (ms)
28	NA	NA	NA	16	448	2.00
29	NA	NA	NA	16	464	2.00
30	NA	NA	NA	16	480	2.00

Packet Payload Size

The packet payload size defines the number of octets contained in the payload of a TDM PW packet when the packet is transmitted. Each DS0 (timeslot) in a DS1 or E1 frame contributes 1 octet to the payload, and the total number of octets contributed per frame depends on the number of timeslots in the channel group (for example, 10 timeslots contribute 10 octets per frame).

Jitter Buffer

A circuit emulation service uses a jitter buffer to ensure that received packets are tolerant to packet delay variation (PDV). The selection of jitter buffer size must take into account the size of the TDM-encapsulated packets (payload size). A properly configured jitter buffer provides continuous play-out, thereby avoiding discards due to overruns and underruns (packets arriving too early or too late). The maximum receive jitter buffer size is configurable for each SAP configured for circuit emulation. The range of values is from 1 to 250 ms in increments of 1 ms.

Configuration/design Considerations

Determining the best configuration value for the jitter buffer may require some adjustments to account for the requirements of your network, which can change PDV as nodes are added or removed.

For each circuit, the maximum receive jitter buffer is configurable. Play-out from this buffer must start when the buffer is 50% full, in order to give an operational PDV equal to half the maximum buffer size. The supported range is 1 to 250 ms in increments of 1 ms. The buffer size must be set to at least 3 times the packetization delay and no greater than 32 times the packetization delay. Use a buffer size (in ms) that is equal to or greater than the peak-to-peak PDV expected in the network used by circuit emulation service. For example, for a PDV of ±5 ms, configure the jitter buffer to be at least 10 ms.



Note:

- The jitter buffer setting and payload size (packetization delay) interact such that it may be necessary for the operating system to adjust the jitter buffer setting in order to ensure no loss of packets. Thus, the configured jitter buffer value may not be the value used by the system. Use the `show>service>id` *`service-id`*`>all` command to show the effective PDVT (packet delay variation tolerance).
- If asymmetric delay control is enabled (`asym-delay-control`), it must be enabled on both ends of the Cpipe and the jitter buffer size must match on both ends of the Cpipe; otherwise, a service parameter mismatch state occurs and the service is brought down.

The following values are the default jitter buffer times for structured circuits without CAS, where N is the number of timeslots:

- for N = 1, the default is 32 ms
- for 2 ≤ N ≤ 4, the default is 16 ms
- for 5 ≤ N ≤ 15, the default is 8 ms
- for N ≥ 16, the default is 5 ms

For CESoPSN with CAS, the default jitter buffer is 12 ms for T1 and 8 ms for E1.

Jitter buffer overrun and underrun counters are available for statistics and can raise an alarm (optional) while the circuit is operational. For overruns, excess packets are discarded and counted. For underruns, an all-ones pattern is sent for unstructured circuits and an all-ones or a user-defined pattern is sent for structured circuits (based on configuration).

The circuit status and statistics can be displayed using the `show` command.

Asymmetric Delay Control

If there is high jitter in the network, the last packet for initialization of the circuit emulation service may arrive early or late, resulting in a jitter buffer latency that is different from the expected configured jitter buffer setting (time associated with 50% jitter buffer size). The latency difference between each direction of the TDM PW is known as asymmetric latency, and because some applications (for example, power industry networks) require a very low latency difference, it must be controlled.

The `asym-delay-control` command is used to control the asymmetric latency contributed by the jitter buffer. When the command is enabled, a special startup sequence is triggered when the TDM PW is initially started or is restarted after being brought down (due to faults such as packet overflow, packet underflow, or the port going down).

Upon startup, a configurable number of TDM PW packets are analyzed. During this analysis period, the access port transmits an all-ones pattern (for the 8-port Voice & Teleprotection card) or the configured `idle-payload-fill` value (for the other port types). Refer to the 7705 SAR OS Interface Configuration Guide for information on the `idle-payload-fill` command. The service is considered to be down during this period.

If any packet loss is detected during the analysis period, the analysis is restarted. If no packet loss is detected, the average jitter buffer latency is computed. Based on the difference between the average latency and the expected latency of the jitter buffer size, the network processor will either:

- drop a number of octets based on the difference (if the measured average is higher than expected)
- add a number of dummy octets based on the difference (if the measured average is lower than expected); the dummy octets are based on the `idle-payload-fill` value of the channel or port

Note:

- Asymmetric delay control can only reduce asymmetry in the jitter buffer. It does not reduce any asymmetry that may exist in the network path.
- Because of this, the network must be engineered to maintain symmetrical latency:
  - → use explicit-path LSPs with strict hops
  - → do not use MPLS FRR because it may change the latency characteristics of a single direction without changing the other direction
  - → ensure that both directions of the TDM PW traverse the same path end to end
- With asymmetric delay control, care must be taken when designing the network to prevent a situation where an error recovery mechanism would result in different MPLS paths in the two directions of the Cpipe. If different paths are utilized, latencies may be different and asymmetry introduced. Therefore, in conjunction with asymmetric delay control, hop-by-hop strict LSPs must be used as the spoke SDP for the Cpipe with no alternative paths.

Optionally, the asymmetric delay control analysis can be set to repeat at configured time intervals after the service is up. This analysis is done with live traffic (that is, not with all-ones or the `idle-payload-fill` value). If the difference between the calculated average latency and the expected latency is greater than 250 µs, octets are added or dropped as necessary.

If asymmetric delay control is enabled, it must be enabled on both ends of the Cpipe; otherwise, a service parameter mismatch state occurs and the service is brought down. Jitter buffer size is also included in the set of parameters that will cause a service parameter mismatch if the value is not the same at both ends of the Cpipe. This prevents the operator from changing the jitter buffer size, which would immediately change the latency symmetry of the Cpipe service.

As well, Cpipes using asymmetric delay control must have the same card and port type on both ends of the Cpipe. Mismatched card/port configuration is not blocked in the CLI or in SNMP but must be avoided; otherwise, differential delay will be introduced caused by different framer delays on the cards/ports.

Asymmetric delay control can only be enabled for Cpipes configured as CESoPSN without CAS (no SAToP). If asymmetric delay control is enabled, ACR, DCR, and RTP cannot be enabled on the port.

The following adapter cards, modules, and platforms support asymmetric delay control:

- 4-port T1/E1 and RS-232 Combination module (RS-232 channels) on the 7705 SAR-H
- 8-port Voice & Teleprotection card (G.703 (codir) and C37.94 (TPIF) channels)
- 12-port Serial Data Interface card (RS-232 and X.21 channels)
- 7705 SAR-Hc (RS-232 channels)

RTP Header

For all circuit emulation channels, the RTP in the header is optional (as per RFC 5086). When enabled for absolute mode operation, an RTP header is inserted in the MPLS frame upon transmit. Absolute mode is defined in RFC 5086 and means that the ingress PE will set timestamps using the clock recovered from the incoming TDM circuit. When an MPLS frame is received, the RTP header is ignored. The RTP header mode is for TDM PW interoperability purposes only and should be enabled when the other device requires an RTP header.

RTP cannot be enabled if asymmetric delay control is enabled.

Control Word

The structure of the control word is mandatory for SAToP and CESoPSN and is shown in Figure 33. Table 23 describes the bit fields. See Pseudowire Control Word for more information.

Figure 33: Control Word Bit Structure

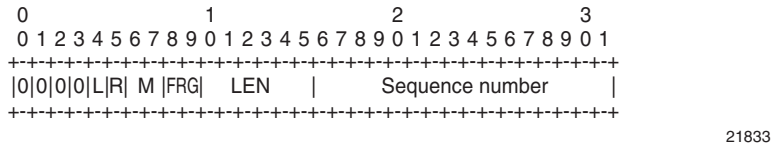


Table 23: Control Word Bit Descriptions

Bit(s)	Description
Bits 0 to 3	The use of bits 0 to 3 is described in RFC 4385. These bits are set to 0 unless they are being used to indicate the start of an Associated Channel Header (ACH) for the purposes of VCCV.
L (Local TDM Failure)	The L bit is set to 1 if an abnormal condition of the attachment circuit such as LOS, LOF, or AIS has been detected and the TDM data carried in the payload is invalid. The L bit is cleared (set back to 0) when fault is rectified.
R (Remote Loss of Frames indication)	The R bit is set to 1 if the local CE-bound interworking function (IWF) is in the packet loss state and cleared (reset to 0) after the local CE-bound IWF is no longer in the packet loss state.

Table 23: Control Word Bit Descriptions (Continued)

Bit(s)	Description
M (Modifier)	The M bits are a 2-bit modifier field. For SAToP, M is set to 00 as per RFC 4553. For CESoPSN, M is set according to RFC 5086, summarized as follows: • When L bit = 0, and M = 00—Normal conditions M = 01—Reserved for future use M = 10—RDI condition for the attachment circuit (AC) M = 11—Reserved for CESoPSN • When L bit = 1, and M = 00—TDM data is invalid M = 01—Reserved for future use M = 10—Reserved for future use M = 11—Reserved for future use
FRG	The FRG bits in the CESoPSN control word are set to 00.
LEN	The LEN bits (bits 10 to 15) carry the length of the CESoPSN packet (defined as the size of the CESoPSN header plus the payload size) if it is less than 64 bytes, and set to 0 otherwise.
Sequence number	The sequence number is used to provide the common PW sequencing function as well as detection of lost packets.

Error Situations

The CE-bound interworking function (IWF) uses the sequence numbers in the control word to detect lost and incorrectly ordered packets. Incorrectly ordered packets that cannot be reordered are discarded.

For unstructured CES, the payload of received packets with the L bit set is replaced with an all-ones pattern. For structured CES, the payload of received packets with the L bit set is replaced with an all-ones or a user-configurable bit pattern. This is configured using the `idle-payload-fill` command. For structured CES with CAS, the signaling bits are replaced with an all-ones or a user-configurable bit pattern. This is configured using the `idle-signal-fill` command. Refer to the 7705 SAR OS Interface Configuration Guide for more information.

All circuit emulation services can have a status of up, loss of packets (LOP) or admin down, and any jitter buffer overruns or underruns are logged.

Ethernet VLL (Epipe) Services

This section provides information about the Epipe service.

Topics in this section include:

- Epipe Service Overview
- Ethernet Access Egress Queuing and Scheduling
- Ethernet SAP-to-SAP
- MEF 8
- Ethernet OAM
- Control Word
- MTU
- Raw and Tagged Modes
- IP Filters (Epipe)

Epipe configuration information is found under the following topics:

- Common Configuration Tasks
- Configuring VLL Components
- Service Management Tasks

See Service Support for information on the adapter cards and chassis that support Ethernet VLL services.

Epipe Service Overview

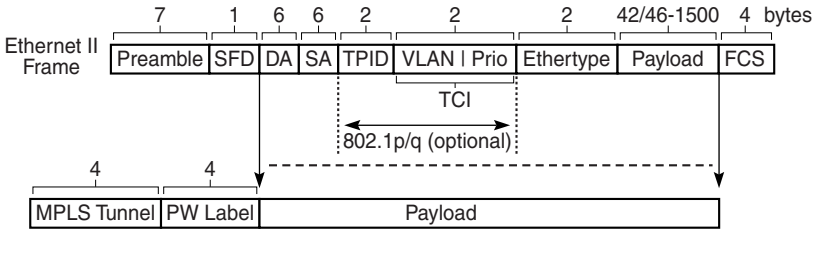
An Ethernet pseudowire (PW) is used to carry Ethernet/802.3 protocol data units (PDUs) over an MPLS or IP network, allowing service providers to offer emulated Ethernet services over existing MPLS or IP networks. For the 7705 SAR, Ethernet emulation is a point-to-point service.

The 7705 SAR uses Ethernet VLLs to carry Ethernet traffic from various sources at a site, including traffic such as e911 locators, power supply probes, and HSPA-dedicated interfaces. Native Ethernet bridging is not supported.

An MPLS Epipe service is the Alcatel-Lucent implementation of an Ethernet VLL based on the IETF RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*.

Figure 34 shows a typical Ethernet VLL frame together with its MPLS tunnel encapsulation.

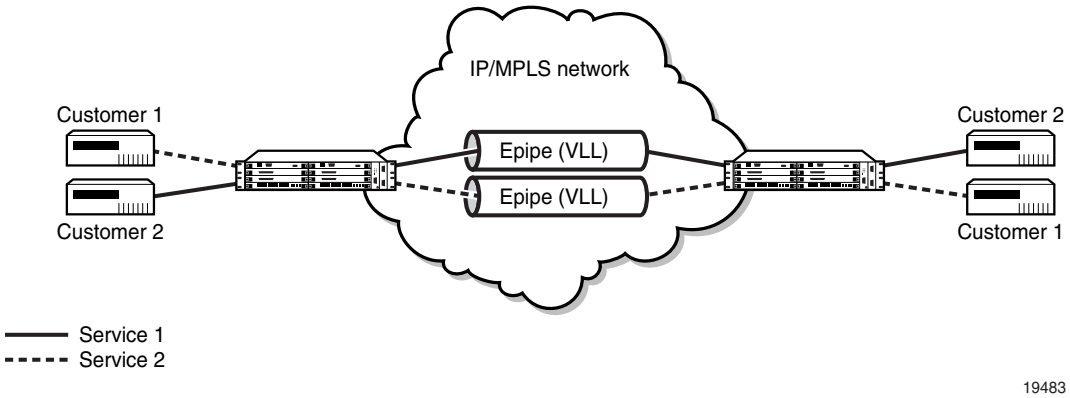
Figure 34: Ethernet VLL Frame with MPLS Encapsulation



An Epipe service is a Layer 2 point-to-point service where the customer data is encapsulated and transported across a service provider's MPLS or IP network. Figure 35 shows a typical Epipe Service. An Epipe service is completely transparent to the subscriber's data and protocols. Like other PW VLL services, Epipe service behaves like a non-learning Ethernet bridge. A distributed Epipe service consists of a SAP and an SDP pair, where one SDP is on same router as the SAP, and the second SDP is on the far-end router.

Each SAP configuration includes a specific port on which service traffic enters the 7705 SAR from the customer side (also called the access side). Each port is configured with an encapsulation type (SAP encapsulation). Thus, a whole Ethernet port can be bound to a single service (that is, the whole Ethernet port is configured as a SAP), or if a port is configured for IEEE 802.1Q or 802.1ad encapsulation (referred to as dot1q or qinq, respectively), then a unique encapsulation value (ID) must be specified.

Figure 35: Epipe Service



Ethernet Access Egress Queuing and Scheduling

Ethernet access egress queuing and scheduling is very similar to the Ethernet access ingress behavior. Once the Ethernet pseudowire is terminated, traffic is mapped to up to eight different forwarding classes per SAP. Mapping traffic to different forwarding classes is performed based on the EXP bit settings of the received Ethernet pseudowire.

For more information on Ethernet access egress queuing and scheduling, refer to the 7705 SAR OS Quality of Service Guide.

Ethernet SAP-to-SAP

Ethernet VLLs can be configured with both endpoints (SAPs) on the same 7705 SAR. This is referred to as Ethernet SAP-to-SAP or local Ethernet service. Ethernet SAP-to-SAP provides local Ethernet switching between two Ethernet endpoints on the 7705 SAR.

An Ethernet SAP-to-SAP connection is set up on the 7705 SAR and a pseudowire is configured between the two endpoints.

When the port encapsulation is null, there is no change to the VLAN tags on the ingress and egress frame headers, if VLAN tags are present.

When the port encapsulation is dot1q, the VLAN tag is removed from the ingress frame header and a new VLAN tag is inserted into the egress frame header. No VLAN tag is inserted into the egress frame header if the SAP has a VLAN ID of 0.

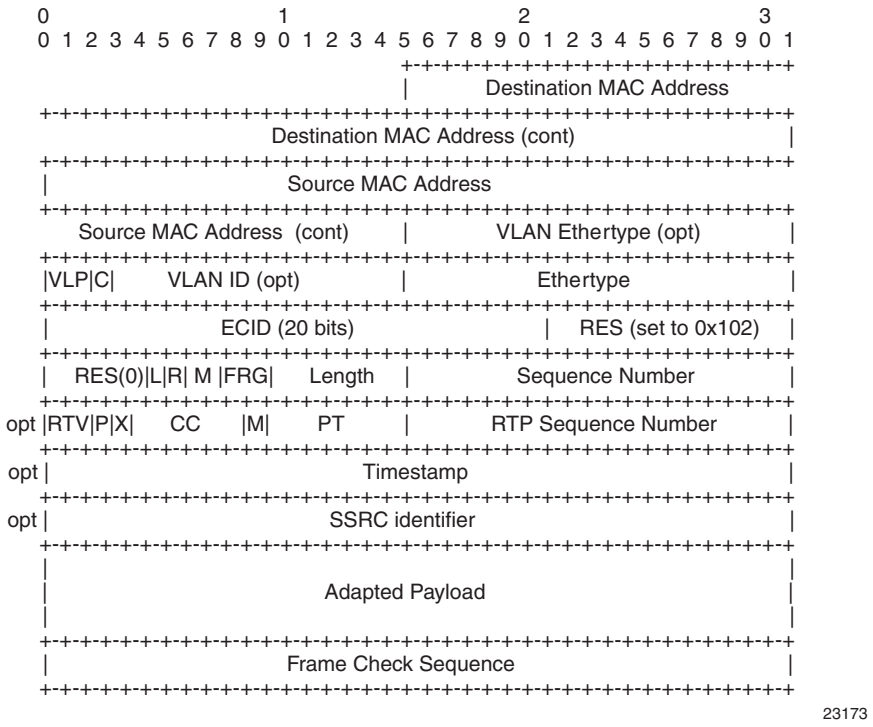
When the port encapsulation is qinq, the VLAN tags are removed from the ingress frame header and a new set of outer and inner VLAN tags are inserted in the egress frame header. No VLAN tags is inserted into the egress frame if the SAP has a VLAN ID of 0 or VLAN IDs of 0.*. Note that 0.0 is not a valid combination.

MEF 8

The 7705 SAR supports MEF 8 as defined in the Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks. Support for the MEF 8 standard allows both structured and unstructured emulation of TDM services across Epipes, also known as Circuit Emulation Services over Ethernet (CESoETH). The MEF 8 feature enables the 7705 SAR to interoperate with equipment that does not support MPLS-based Cpipes but does support MEF 8 Epipes.

Figure 36 shows an Ethernet-encapsulated TDM circuit. See TDM PW Encapsulation for complete information about TDM PW encapsulation.

Figure 36: CESoETH Encapsulation



The 7705 SAR supports the following MEF 8 configuration scenarios:

- TDM SAP to Ethernet SAP (see Figure 37)

  A GPON ONT or other MEF 8-only device is used to encapsulate TDM over a GPON or Ethernet network, and the 7705 SAR is used to terminate the MEF 8.

- TDM SAP to Ethernet SAP via 7705 SAR-M GPON module (see Figure 38)

  The 7705 SAR-M is the GPON ONT and is used to interoperate with a MEF 8-only device. The 7705 SAR-M terminates the MEF 8.

- TDM SAP to spoke SDP (see Figure 39)

  A GPON ONT or other MEF 8-only device is used to encapsulate TDM over an MPLS network, and the 7705 SAR is used to terminate both the LSP and MEF 8.

Figure 37: TDM SAP to Ethernet SAP

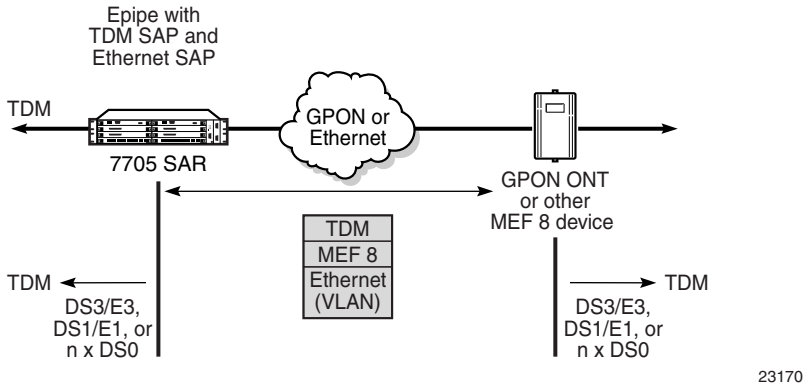


Figure 38: TDM SAP to Ethernet SAP via 7705 SAR-M GPON Module

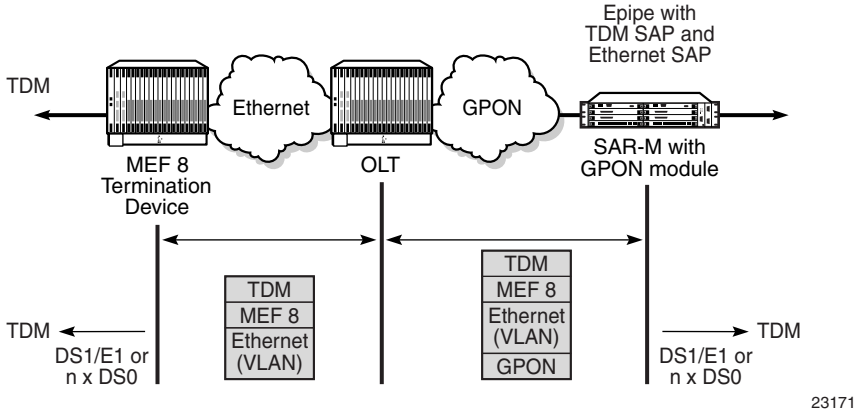


Figure 39: TDM SAP to Spoke SDP

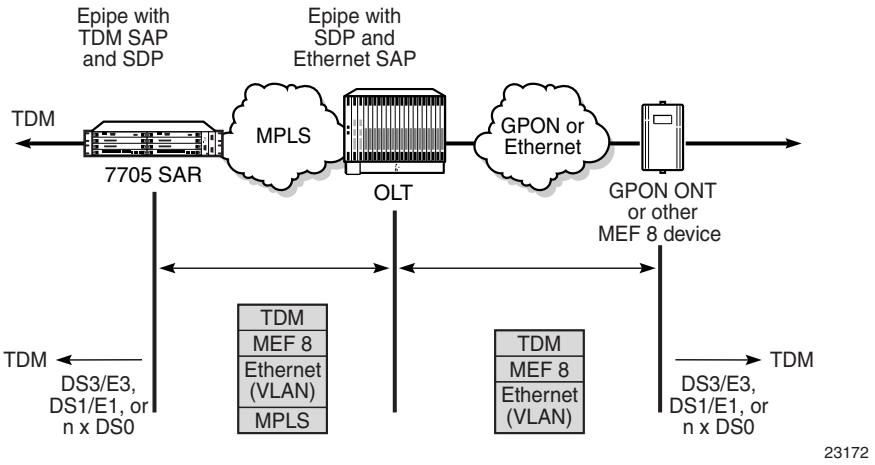


Table 24 shows the platforms, adapter cards, and modules that support MEF 8.

Table 24: MEF 8 Support on the 7705 SAR

	TDM SAP to Ethernet SAP		TDM SAP to Ethernet via Epipe Spoke SDP	
	Structured [1]	Unstructured	Structured [1]	Unstructured
7705 SAR-M All T1/E1 ports	✓	✓ [2]	✓	✓ [2]
7705 SAR-F All T1/E1 ports	✓	✓ [2]	✓	✓ [2]
7705 SAR-A All T1/E1 ports	✓	✓ [2]	✓	✓ [2]
7705 SAR-X All T1/E1 ports	✓	✓ [2]	✓	✓ [2]
16-port T1/E1 ASAP Adapter card	✓	✓ [2]	✓	✓ [2]
32-port T1/E1 ASAP Adapter card	✓	✓ [2]	✓	✓ [2]
2-port OC3/STM1 Channelized Adapter card	✓	✓ [3]	✓	✓ [3]
4-port DS3/E3 Adapter card	✓	✓ [4]	✓	✓ [4]

Table 24: MEF 8 Support on the 7705 SAR (Continued)

	TDM SAP to Ethernet SAP		TDM SAP to Ethernet via Epipe Spoke SDP	
	Structured [1]	Unstructured	Structured [1]	Unstructured
4-port OC3/STM1 / 1-port OC12/STM4 Adapter card		✓ [5]		✓ [5]
4-port T1/E1 and RS-232 Combination module All T1/E1 ports	✓	✓ [2]	✓	✓ [2]

Notes:

1. Supported on n x DS0 channels with or without CAS
2. Supported on DS1/E1 ports
3. Supported on DS1/E1 channels and DS3/E3 channels
4. Supported on DS1/E1 channels and DS3/E3 ports
5. Supported on DS1/E1 channels

Epipe Service Modes

Epipe services support structured circuit emulation mode for nxDS0 and structured or unstructured circuit emulation mode for DS1, E1, DS3, and E3 as defined in the MEF 8 specification.

There are two methods for using MEF 8 to emulate TDM circuits over Ethernet using an Epipe:

- TDM SAP-to-Ethernet SAP
- TDM SAP-to-spoke SDP

Defining one TDM SAP and one Ethernet SAP is known as Circuit Emulation Services over Ethernet (CESoETH). The TDM SAP configured in the Epipe must include a local and remote Emulated Circuit Identifier (ECID) and a far-end destination MAC address. The TDM port's MAC address is used as the source MAC address for the circuit.

TDM can also be encapsulated into Ethernet which is then encapsulated in MPLS (or GRE). This method is known as Circuit Emulation Services over Ethernet over MPLS (CESoETHoMPLS). CESoETHoMPLS is configured with an Epipe with a TDM SAP and a spoke SDP. The TDM SAP configured in the Epipe must include a local and remote ECID and a far-end destination MAC address. The TDM port's MAC address is used as the source MAC address for the circuit.

The 7705 SAR supports unicast MAC addresses and non-IEEE-reserved group multicast MAC addresses.

Note: Users should exercise caution when using multicast MAC addresses, as Ethernet frames with a multicast destination address could be flooded when traversing an Ethernet broadcast domain.

The TDM SAP's framing and CAS settings determine the MEF 8 circuit emulation mode. If the TDM port is framed, MEF 8 is in structured mode. If the TDM port is unframed, MEF 8 is in unstructured mode. If the TDM SAP is configured with CAS enabled, MEF 8 is in structured mode with CAS. See Unstructured, Structured DS1/E1 CES without CAS, and Structured T1/E1 CES with CAS for more information about circuit emulation modes.

Adaptive clock recovery (ACR) is supported for MEF 8 in structured or unstructured mode on the following platforms and adapter cards:

- 7705 SAR-F (T1/E1 ports)
- 7705 SAR-M (variants with T1/E1 ports)
- 7705 SAR-A (variants with T1/E1 ports)
- 7705 SAR-X (T1/E1 ports)
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card

For more information on ACR, refer to the 7705 SAR OS Basic System Configuration Guide, "Adaptive Clock Recovery (ACR)".

Differential clock recovery (DCR) is supported for MEF 8 in structured or unstructured mode on the following platforms, adapter cards, and modules:

- 7705 SAR-M (variants with T1/E1 ports)
- 7705 SAR-A (variant with T1/E1 ports)
- 7705 SAR-X (T1/E1 ports)
- 16-port T1/E1 ASAP Adapter card, version 2
- 32-port T1/E1 ASAP Adapter card
- 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card (DS1/E1 channels)
- 4-port DS3/E3 Adapter card (clear channel DS3/E3 ports and DS1/E1 channels on channelized DS3 ports (E3 ports cannot be channelized)); DCR on DS1/E1 channels is supported only on the first three ports of the card
- T1/E1 ports of the 4-port T1/E1 and RS-232 Combination module

To enable DCR, the network must have a common clock between the two terminating SAPs or SAP/spoke SDP using MEF 8. In each direction, the service clock is compared to the common clock and the difference is encoded into the RTP header in the TDM PW overhead. At the other end of the network, the original service clock is reproduced by comparing the common clock to the frequency difference in the RTP header.

Note: DCR is not supported on DS1 or E1 channels that have CAS signaling enabled.

For more information on DCR, refer to the 7705 SAR OS Basic System Configuration Guide, "Differential Clock Recovery (DCR)".

Ethernet OAM

Ethernet VLL service supports Ethernet OAM functions for ETH-CFM according to the 802.1ag and Y.1731 standards, for Y.1731 Performance Monitoring, and for EFM OAM according to the 802.3ah standard. For more information, see ETH-CFM (802.1ag and Y.1731), and refer to the "Ethernet OAM" section in the 7705 SAR OS Interface Configuration Guide, and the "Ethernet OAM Capabilities" section in the 7705 SAR OS OAM and Diagnostics Guide.

Ethernet ports in access or network mode also support CFM loopback message (LBM) frames for Layer 1 and Layer 2 OAM tests on unlabeled ports. For more information, refer to the "Ethernet OAM" section in the 7705 SAR OS Interface Configuration Guide.

Control Word

Ethernet VLL supports an optional control word (CW). The control word contains protocol control information and must be used for any pseudowire service that transports packets that are less than 64 bytes; see Pseudowire Control Word for more information.

MTU

Network-facing Ethernet ports must support a larger MTU than access-facing Ethernet ports in order to account for the pseudowire headers that are added to the access Ethernet frames.

As an example, the following list gives the worst-case MTU sizes for Ethernet VLLs over Ethernet port(s) under various configurations, where the worst case is the largest MTU size required in order to carry a standard payload (1500 bytes):

- Access, null mode: 1514 bytes
- Access, dot1q mode: 1518 bytes
- Access, qinq mode: 1522 bytes
- Network, null mode: 1572 bytes
- Network, dot1q mode: 1572 bytes

Note: Since it is not practical to split a Layer 2 Ethernet frame into smaller frames, the access port (SAP) MTU must be smaller than the service and network port MTU. If the access port MTU is larger than the tunnel MTU, the Ethernet VLL does not come into service and remains in the inoperative state. See MTU Settings for information on MTU for VLL service.

Raw and Tagged Modes

An Ethernet PW operates in one of two modes: raw or tagged. Raw and tagged modes relate to the way the router handles VLAN tags embedded in the header of an Ethernet frame. Both modes are supported by the 7705 SAR.

Raw and tagged modes are configured using the `vc-type {ether | vlan}` parameter under the `spoke-sdp` command. To configure raw mode, choose the `ether` option; to configure tagged mode, choose `vlan`.

VLAN tags can provide service-delimiting information about a frame. Service-delimiting means that information in the tag affects the forwarding decisions that are made to route the packet. The port connected to the attachment circuit (AC) can be configured for null, dot1q, or qinq operation. When the port is configured for null, the 7705 SAR treats any attached tag received at the SAP (from the AC) as not service-delimiting; when configured for dot1q or qinq, received tags are service-delimiting.

Raw Mode

In raw mode, VLAN tags are not service-delimiting (that is, the port is set to null and the tags do not affect frame forwarding decisions) and are forwarded over the Epipe as part of the payload.

If a service-delimiting tag arrives from the ingress AC (that is, the port is set to dot1q or qinq and a tag is received), the tag(s) is removed (popped) from the payload before the Ethernet frame gets switched over the PSN via the Epipe.

In raw mode, all traffic from the ingress port gets switched to the same endpoint. However, if the MTU (or configured size) of the tunnel is exceeded then service is affected because the frame is dropped.

In raw mode, when the 7705 SAR detects a failure on the Ethernet ingress port or the port is administratively disabled, the 7705 SAR sends a PW status notification message to the remote router.

Tagged Mode

In tagged mode, every frame sent on the Ethernet PW has a service-delimiting VLAN tag. If the frame received by the 7705 SAR from the attachment circuit (AC) does not have a service-delimiting VLAN tag, then the 7705 SAR inserts (pushes) a VLAN tag into the frame header before sending the frame to the SDP and the PW. If the frame received from the AC has a service-delimiting VLAN tag, the tag is replaced.

In tagged mode, when the 7705 SAR detects a failure on the Ethernet physical port or the port is administratively disabled, the 7705 SAR sends a PW status notification message for all PWs associated with the port.

VLAN Translation

VLAN ID translation is supported, as appropriate. Table 27 (see Tagging Rules for Epipe) shows the VLAN ID translation operation for the various packet types. The payload part of the packet is shown in parentheses.

The operations to add, strip (remove), or forward the VLAN headers are performed based on the encapsulation type at the ingress of the attachment circuit (the SAP), in the network, and at the egress circuit.

Tagging Rules for Epipe

Table 25 and Table 26 show the general tagging rules for combinations of interface port type (null, dot1q, or qinq) and Epipe type (Ethernet or VLAN) for SAP ingress and SAP egress directions.

An AC (attachment circuit, ingress or egress) can be configured for one of the following encapsulation types:

- null
- dot1q
- qinq

Note:

- Single-tagged Ethernet packets with VLAN ID set to 0 are handled in the same way as other tagged packets.
- For Ethernet VLL SAP-to-SAP with a VLAN SAP, the dot1p header is not popped at ingress. Instead, it is preserved across the ingress and egress ports to ensure dot1p transparency. Dot1q bits are overwritten with configured VLAN information at egress. If there is a matching non-default QoS policy applied at egress, the dot1p bits are re-marked accordingly.

Table 25: Ingress SAP Tagging Rules

Ingress SAP Type [1]	VC Type (Epipe)	
	Raw (Ethernet)	Tagged (VLAN)
Null	No operation	Push (VC tag)
Dot1q	Pop (outer tag)	Pop (outer tag) Push (VC tag) [2]
QinQ	Pop (outer tag)	Pop (outer tag) Push (VC tag)

Notes:

1. Ingress SAP type is configured at the port level.
2. If the VC tag is not set, then the original tag is preserved. The VC tag is set on the spoke SDP, using the **`vlan-vc-tag`** command.

Table 26: Egress SAP Tagging Rules

Egress SAP Type [1]	VC Type (Epipe)	
	Raw (Ethernet)	**Tagged (VLAN)**
Null	No operation	Pop (VC tag)
Dot1q	Push (SAP tag) [2]	Pop (VC tag) Push (SAP tag) [3]
QinQ	Push (SAP tags)	Pop (VC tag) Push (SAP tags)

Notes:

1. Egress SAP type is configured at the port level.
2. If the SAP tag is 0, then no VLAN tag is pushed.
3. If the SAP tag is 0, then only the pop operation is performed.

Table 27 shows several examples of how VLAN tags are translated as they flow from ingress to egress. The ingress or egress point can be a SAP or an SDP. For a SAP, encapsulation can be null, dot1q, or qinq; for an SDP, encapsulation (vc-type) can be ether raw or VLAN (tagged). When the SAP encapsulation is dot1q or qinq, outer and inner tags are used.

Note: When the SAP type is dot1q or qinq:

- the SAP VLAN tag always delimits the ingress traffic, regardless of the Ethernet VLL type (raw or tagged)
- untagged frames are dropped at the SAP ingress. That is, only the frames with an outer VLAN tag that matches the SAP VLAN tag are forwarded. The exception to this case occurs when the VLAN tag = 0. When a SAP is configured with VLAN ID = 0, any received untagged packets are processed.

Table 27: VLAN Tagging Examples (Epipe)

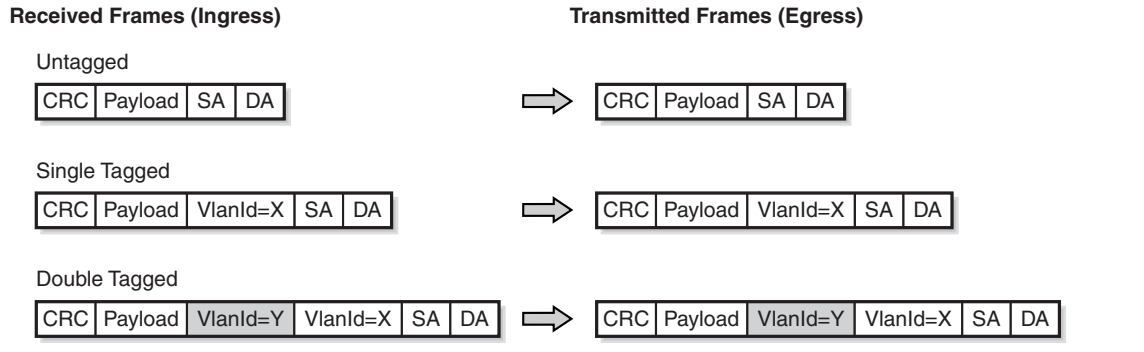
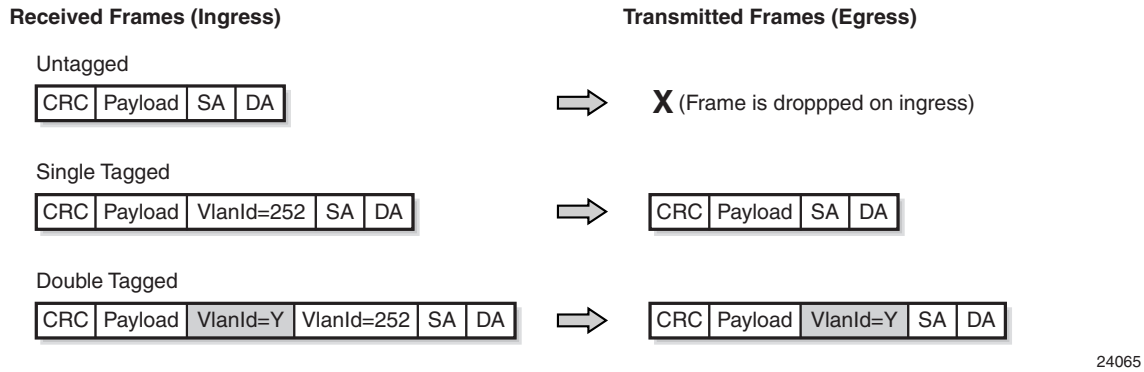
<table>
<tr><th rowspan="2">Example</th><th colspan="6">Configuration Settings</th></tr>
<tr><th>Rx (Ing)/
Tx (Egr)</th><th>SAP/
SDP</th><th>Encap.</th><th colspan="2">SAP VLAN Tag</th><th></th></tr>
<tr><th></th><th></th><th></th><th>Outer</th><th>Inner</th><th></th></tr>
<tr><td rowspan="3">1</td><td>Rx
Tx</td><td>SAP
SAP</td><td>Null
Null</td><td>N/A
N/A</td><td>N/A
N/A</td><td></td></tr>
<tr><td colspan="6">Result: Untagged, single-, and double-tagged frames are accepted. On egress, all frames are transmitted untouched.</td></tr>
<tr><td colspan="6"> 
</td></tr>
<tr><td rowspan="3">2</td><td>Rx
Tx</td><td>SAP
SAP</td><td>Dot1q
Null</td><td>252
N/A</td><td>N/A
N/A</td><td></td></tr>
<tr><td colspan="6">Result: Single- and double-tagged frames are accepted if the outermost VLAN ID matches. On egress, the outermost VLAN tag is popped.</td></tr>
<tr><td colspan="6"> 
</td></tr>
</table>

Table 27: VLAN Tagging Examples (Epipe) (Continued)

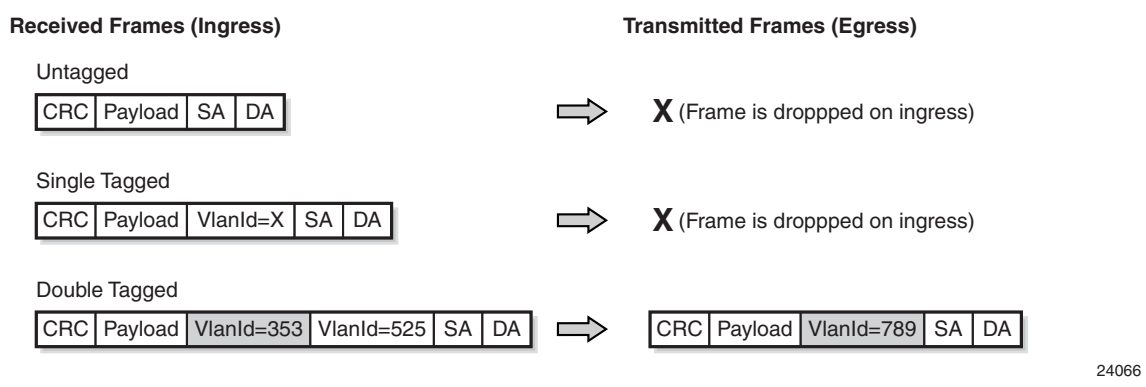
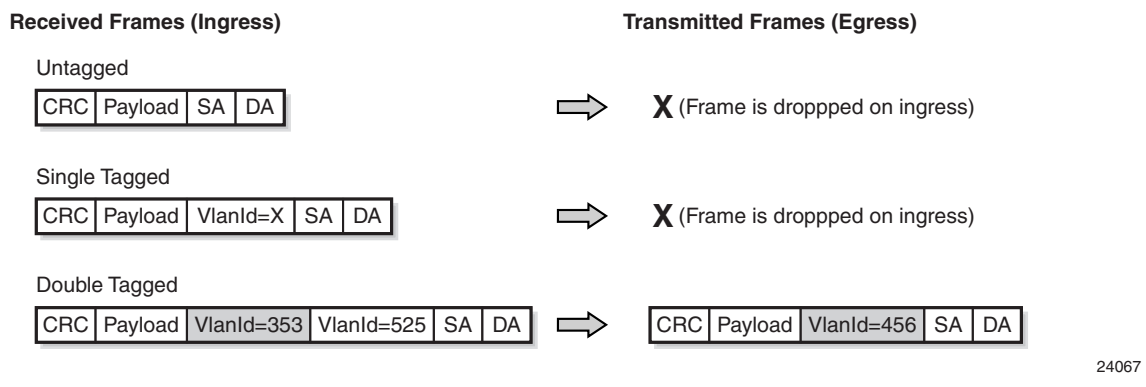
<table>
<tr><th rowspan="2">Example</th><th colspan="6">Configuration Settings</th></tr>
<tr><th>Rx (Ing)/
Tx (Egr)</th><th>SAP/
SDP</th><th>Encap.</th><th colspan="2">SAP VLAN Tag</th><th></th></tr>
<tr><th></th><th></th><th></th><th>Outer</th><th>Inner</th><th></th></tr>
<tr><td rowspan="3">3</td><td>Rx
Tx</td><td>SAP
SAP</td><td>QinQ
Dot1Q</td><td>525
789</td><td>353
N/A</td><td></td></tr>
<tr><td colspan="6">Result: Double-tagged frames are accepted if the outer/inner VLAN ID matches. On egress, the outer VLAN tag is popped and the inner VLAN tag is swapped.</td></tr>
<tr><td colspan="6"> 
</td></tr>
<tr><td rowspan="3">4</td><td>Rx
Tx</td><td>SAP
SDP</td><td>QinQ
VLAN</td><td>525
456</td><td>353
N/A</td><td></td></tr>
<tr><td colspan="6">Result: Double-tagged frames are accepted if the outer/inner VLAN ID matches. On egress, the outer VLAN tag is popped and the inner VLAN tag is swapped.</td></tr>
<tr><td colspan="6"> 
</td></tr>
</table>

Table 27: VLAN Tagging Examples (Epipe) (Continued)

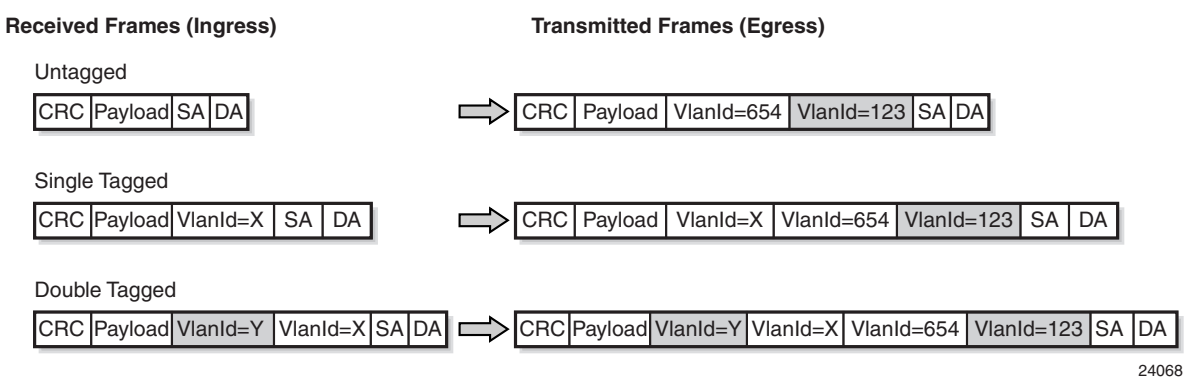
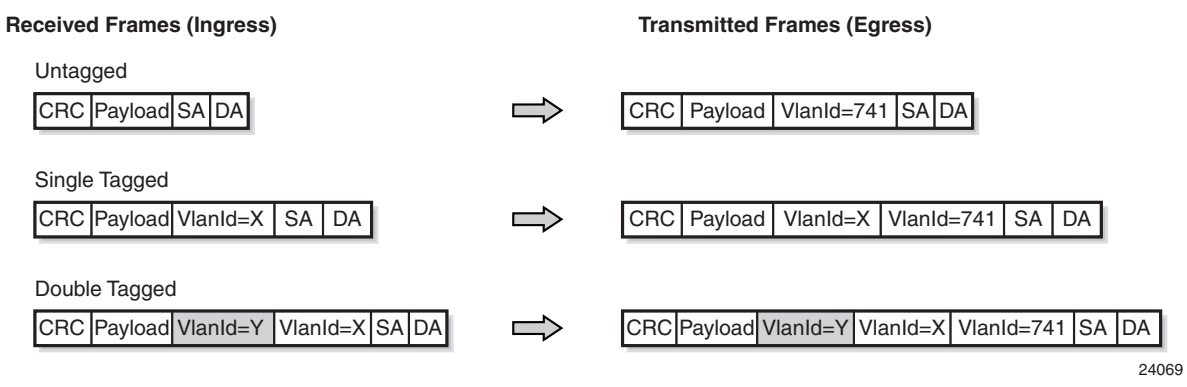
<table>
<tr><th rowspan="2">Example</th><th colspan="6">Configuration Settings</th></tr>
<tr><th>Rx (Ing)/
Tx (Egr)</th><th>SAP/
SDP</th><th>Encap.</th><th>SAP VLAN Tag
Outer</th><th>SAP VLAN Tag
Inner</th><th></th></tr>
<tr><td>5</td><td>Rx
Tx</td><td>SDP
SAP</td><td>Ether
QinQ</td><td>N/A
123</td><td>N/A
654</td><td></td></tr>
<tr><td></td><td colspan="6">Result: Untagged, single-, and double-tagged frames are accepted. On egress, a double-push (outer and inner) is performed.</td></tr>
<tr><td></td><td colspan="6"> 
</td></tr>
<tr><td>6</td><td>Rx
Tx</td><td>SAP
SDP</td><td>Dot1q
VLAN</td><td>Default (*)
741</td><td>N/A
N/A</td><td></td></tr>
<tr><td></td><td colspan="6">Result: Untagged, single-, and double-tagged frames are accepted. On egress, a push is performed.</td></tr>
<tr><td></td><td colspan="6"> 
</td></tr>
</table>

Table 27: VLAN Tagging Examples (Epipe) (Continued)

<table>
<tr><th rowspan="2">Example</th><th colspan="6">Configuration Settings</th></tr>
<tr><th>Rx (Ing)/
Tx (Egr)</th><th>SAP/
SDP</th><th>Encap.</th><th colspan="2">SAP VLAN Tag</th><th></th></tr>
<tr><th></th><th></th><th></th><th>Outer</th><th>Inner</th><th></th></tr>
<tr><td rowspan="3">7</td><td>Rx
Tx</td><td>SAP
SDP</td><td>QinQ
VLAN</td><td>852
963</td><td>0
N/A</td><td></td></tr>
<tr><td colspan="6">Result: Single- and double-tagged frames are accepted if the outermost VLAN ID matches. On egress, the outermost VLAN tag is swapped.</td></tr>
<tr><td colspan="6">Received Frames (Ingress) Transmitted Frames (Egress)
Untagged
CRC | Payload | SA | DA ⇨ X (Frame is droppped on ingress)
Single Tagged
CRC | Payload | VlanId=852 | SA | DA ⇨ CRC | Payload | VlanId=963 | SA | DA
Double Tagged
CRC | Payload | VlanId=0 | VlanId=852 | SA | DA ⇨ CRC | Payload | VlanId=0 | VlanId=963 | SA | DA
24070</td></tr>
<tr><td rowspan="3">8</td><td>Rx
Tx</td><td>SAP
SAP</td><td>QinQ
QinQ</td><td>Default (*)
Default (*)</td><td>Default (*)
Default (*)</td><td></td></tr>
<tr><td colspan="6">Result: Any double-tagged frames are accepted. On egress, all frames are transmitted untouched.</td></tr>
<tr><td colspan="6">Received Frames (Ingress) Transmitted Frames (Egress)
Untagged
CRC | Payload | SA | DA ⇨ X (Frame is droppped on ingress)
Single Tagged
CRC | Payload | VlanId=X | SA | DA ⇨ X (Frame is droppped on ingress)
Double Tagged
CRC | Payload | VlanId=Y | VlanId=X | SA | DA ⇨ CRC | Payload | VlanId=Y | VlanId=X | SA | DA
24071</td></tr>
</table>

IP Filters (Epipe)

IP filters are applied to Epipe SAPs in the ingress direction, as described below. For a full list of entities to which IP filters can be applied, see IP Filter Policies.

Ethernet pseudowires are generally used to transparently switch traffic across an MPLS network to the far end. However, in some cases, the traffic that is switched over the network, consuming valuable bandwidth, is just discarded at the other end of the pseudowire. As well, with the 7705 SAR expanding into areas such as vertical markets, and with local area networks being connected to the 7705 SAR Ethernet ports, an increasing amount of traffic must stay local and not pass through the MPLS network to the far end. By using IP filters at the access ingress, operators can determine what traffic is passed through the pseudowire and therefore use the network links more efficiently.

IP filters can also be used for security purposes, by allowing access only to designated services (for example, allowing e-mail and FTP services while disallowing Telnet services) at the origin of the traffic.

IP filter policies specify either a forward or a drop action for packets, based on information specified in the match criteria. Within each filter policy, you can create entries that define matching criteria.

The same IP filter policy can be assigned to any entity (network interfaces, IP pseudowires, Ethernet pseudowires, IES, and VPRN services), all of which can be configured on the same adapter card. For example, a filter policy defined as “filter-5” can be assigned to multiple Epipe SAPs and, simultaneously, to network interfaces on the same adapter card.

A filter policy assigned to an entity on one adapter card can also be assigned to any entity on another adapter card. For example, a filter policy defined as “filter-2” can be assigned to an Epipe on an Ethernet Adapter card and to a network interface on another Ethernet Adapter card.

Assigning the same filter policy to different entities on a card counts as using one filter policy.

Configuration and assignment of filter policies is similar for network interfaces, IES management SAPs, Ethernet and IP pseudowire SAPs, VPRN and IES SAPs and spoke SDPs, and VPLS SAPs and SDPs (spoke and mesh). This guide describes the assignment of filter policies to SAPs and SDPs. Refer to the 7705 SAR OS Router Configuration Guide, “Filter Policies”, for information on configuring filter policies and assigning them to network interfaces.

Frame Relay VLL (Fpipe) Services

This section provides information about the Fpipe service. Topics in this section include:

- Fpipe Service Overview
- FR SAP-to-SAP Service
- FR Traffic Management
- FR Encapsulation
- Status Signaling and OAM Propagation

Fpipe configuration information is treated under the following topics:

- Common Configuration Tasks
- Configuring VLL Components
- Service Management Tasks

See Service Support for information on the adapter cards and chassis that support frame relay VLL services.

Fpipe Service Overview

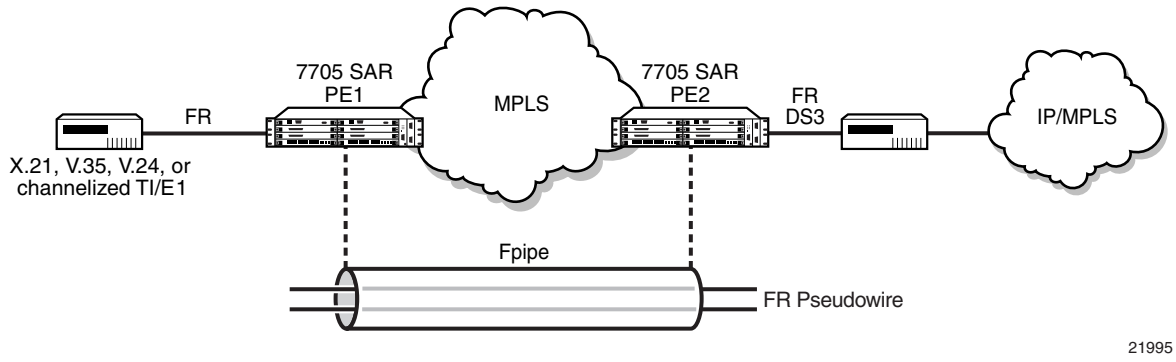
The frame relay VLL (Fpipe) provides a point-to-point frame relay service between users connected to 7705 SAR nodes or other SR routers over an IP/MPLS network. Users are connected to the 7705 SAR nodes using frame relay PVCs.The 7705 SAR receives a standard Q.922 core frame on the frame relay SAP and encapsulates it into a pseudowire packet according to the 1-to-1 frame relay encapsulation mode in RFC 4619 *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*.

The MPLS tunnel label is used by MPLS LSRs to forward a PW packet from one PE to the other. The PW label identifies one PW (that is, one LSP) assigned to an FR VC in one direction. Together the MPLS tunnel label and PW label form an MPLS label stack, as described in RFC 3032.

The control word contains protocol control information and must be used for any pseudowire service that transports packets that are less than 64 bytes; see FR PW Control Word for more information. The payload field corresponds to X.36/X.76 FR frame information field with the following components removed: bit/byte stuffing, FR header, and FCS.The maximum length of the payload field must be agreed on by the two PEs; this agreement can be achieved by using the MTU interface parameter when the PW is established (RFC 4447).

Figure 40 shows an example of the frame relay VLL end-to-end service.

Figure 40: FR VLL for End-to-End FR Service



Fpipe SAPs are supported on the following:

- 16-port T1/E1 ASAP Adapter card, version 2
- 32-port T1/E1 ASAP Adapter card
- 4-port DS3/E3 Adapter card
- 12-port Serial Data Interface card
- T1/E1 ports on the 7705 SAR-M (variants with T1/E1 ports)

Before an Fpipe can be configured on an MDA that supports multiple MDA modes, namely the 16-port T1/E1 ASAP Adapter card, version 2, the 32-port T1/E1 ASAP Adapter card, and the 4-port DS3/E3 Adapter card, the `mda-mode` command must be set to `cem-fr-hdlc-ppp`. Likewise, the AC port that is bound to the Fpipe SAP must have the `encap-type` command set to frame-relay. Refer to the 7705 SAR OS Interface Configuration Guide for more information about how to set the `mda-mode` command at the card level and the `encap-type` command at the port/channel level. See Service Support for more information about the supported chassis and the port- and channel-level configuration requirements.

FR SAP-to-SAP Service

FR VLLs can be configured with both endpoints (SAPs) on the same 7705 SAR, which is referred to as an FR SAP-to-SAP or local FR service. FR SAP-to-SAP emulates local FR switching between two FR endpoints on the same 7705 SAR node.

FR Traffic Management

Traffic management of frame relay VLLs is achieved through the application of ingress and egress QoS policies to SAPs.

Ingress SAP Classification and Marking

If the de-1-out-profile is configured, DE=1 frames are classified as out-of-profile and are not subject to the CIR marking. All received DE=0 frames that exceed the CIR are marked as out-of-profile and have the DE set to 1 regardless of whether the de-1-out-profile command is enabled or disabled. Refer to the 7705 SAR OS Quality of Service Guide for more information about the de-1-out-profile command.

FR Encapsulation

Figure 41 shows the standard FR frame.

Figure 41: FR Frame

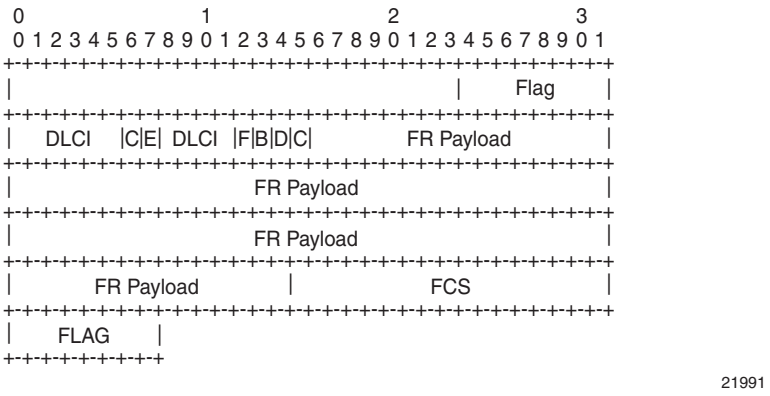
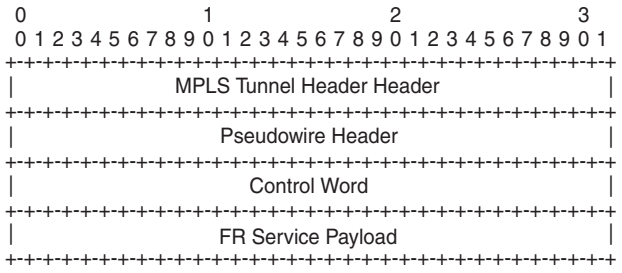


Figure 42 shows the one-to-one mapping mode for FR encapsulation over an MPLS network according to RFC 4619. The FR service payload can be *n* octets.

Figure 42: FR PW 1-to-1 MPLS PSN Encapsulation



The native FR PDU is processed as follows:

- Flag—The FR flags are removed during encapsulation.
- FCS—The FCS can use a 2-byte CRC-CCITT or a 4-byte CRC-32. The CRC is configurable. The FCS is removed during encapsulation.
- Frame header—A 2-byte DLCI frame header is supported. The header is removed during encapsulation.
- The F, B, D, and C control bits are copied into the control word as described in FR PW Control Word.

FR PW Control Word

The control word is required for FR one-to-one mode. Figure 43 shows the FR PW control word.

Figure 43: FR PW Control Word

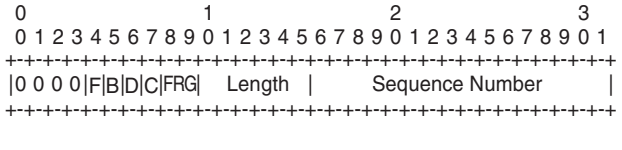


Table 28 describes the bits used in the FR PW control word.

Table 28: Control Word Bit Descriptions

Bit(s)	Description
Bits 0 to 3	4 bits. Bits 0 to 3 are always set to 0 to indicate the presence of the PW.
F	Bit 4; FR FECN bit. The F bit is copied from the FR frame.
B	Bit 5; FR BECN bit. The B bit is copied from the FR frame.
D	Bit 6; FR DE bit. The D bit is copied from the FR frame, but can be reset as a result of the ingress frame policy, as described in FR Traffic Management.
C	Bit 7; FR frame C/R bit. The C bit is copied unchanged from the FR frame.
FRG	2 bits (bits 8 and 9). FRG bits are not supported. The bits must be set to 0.
Length	6 bits (bits 10 to 15). If the frame length (defined as the length of the FR Layer 2 payload plus the length of the control word) is less than 64 octets, the length field must be set to the PW payload length. Otherwise, the length field must be set to zero. The value of the length field, if non-zero, is used to remove the padding characters by the egress PE. See the control-word command description for more information.
Sequence Number	A 16-bit, unsigned circular space. Sequence numbers provide a mechanism to ensure the ordered delivery of PW packets. The sequence number field is not supported. The sequence number value 0 indicates that the sequence number check algorithm is not used.

Status Signaling and OAM Propagation

The 7705 SAR supports the mapping and notification of defect states between an FR PW and an ACs in accordance with *Pseudowire (PW) OAM Message Mapping draft-ietf-pwe3-oam-msg-map-14.txt*, Section 10. Failures in the network are propagated to the customer edge using LMI messages. LMI and AC failures are propagated to the network using PW status signaling.

HDLC VLL (Hpipe) Services

This section provides information about the Hpipe service. Topics in this section include:

- Hpipe Service Overview
- HDLC VLL for End-to-End HDLC Service
- HDLC SAP-to-SAP Service
- HDLC Encapsulation
- Status Signaling

Hpipe configuration information is found under the following topics:

- Common Configuration Tasks
- Configuring VLL Components
- Service Management Tasks

See Service Support for information on the adapter cards and chassis that support HDLC VLL services.

Hpipe Service Overview

An HDLC PW is used to carry HDLC PDUs over an MPLS network. HDLC PWs enable service providers to offer emulated HDLC services over existing MPLS networks.

HDLC mode provides port-to-port transport of HDLC-encapsulated traffic.The HDLC PDU is transported in its entirety, including the HDLC address and control fields, but the HDLC flags and the FCS are excluded. If the optional control word is used, the flag bits in the control word are not used and must be set to 0 for transmitting and must be ignored upon receipt.

HDLC PWs are implemented in accordance with RFC 4618.

Hpipe SAPs are supported on the following:

- 16-port T1/E1 ASAP Adapter card, version 2
- 32-port T1/E1 ASAP Adapter card
- 12-port Serial Data Interface card
- T1/E1 ports on the 7705 SAR-M (variants with T1/E1 ports)

Before HDLC SAPs can be configured on the 16-port T1/E1 ASAP Adapter card, version 2, and the 32-port T1/E1 ASAP Adapter card, the `mda-mode` command must be set to `cem-fr-hdlc-ppp` at the card level. See the 7705 SAR OS Interface Configuration Guide for more information about how to set the `mda-mode` command. See Service Support for more information about the supported chassis and the port- and channel-level configuration requirements.

Only ports that are configured with HDLC encapsulation can be mapped to an Hpipe SAP. HDLC encapsulating ports do not terminate the HDLC. The ports pass the HDLC frames through the Hpipe. HDLC encapsulated ports can pass through any HDLC-framed traffic, such as Cisco-HDLC, FR, PPP, and so on.

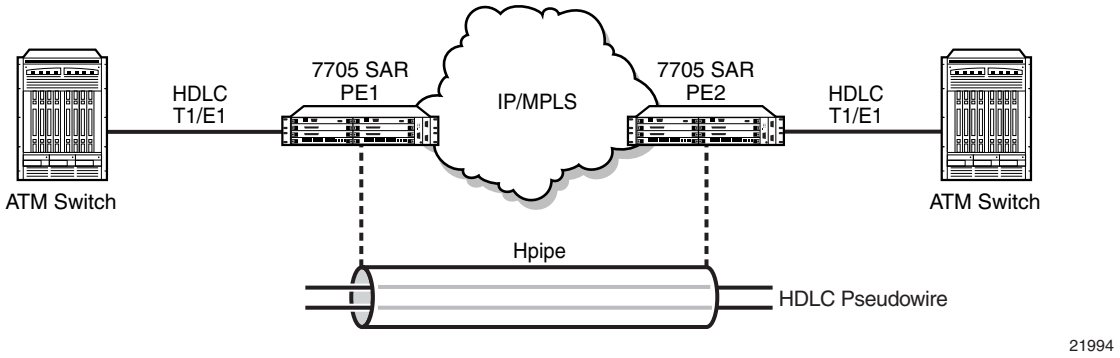
HDLC encapsulation can be used on a port to transmit Cisco-HDLC frames over an H-pipe. To transport Cisco-HDLC traffic over an Hpipe, the attachment circuit (AC) port that is bound to the Hpipe SAP must have the `encap-type` command set to `hdlc`, not `cisco-hdlc`. See the 7705 SAR OS Interface Configuration Guide for more information about how to set the `encap-type` command.

Note: A Cisco-HDLC encapsulating port cannot be bound to an Hpipe SAP. A Cisco-HDLC encapsulating port can only be bound to an Ipipe SAP. See IP Interworking VLL (Ipipe) Services for more information.

HDLC VLL for End-to-End HDLC Service

Figure 44 shows an example of how a mobile operator can deploy end-to-end HDLC services over an MPLS network.

Figure 44: HDLC VLL for End-to-End HDLC Service



In Figure 44, the CE (an ATM switch) transmits HDLC PDUs and receives HDLC PDUs over the physical layer between the CE and a 7705 SAR (PE1). The native service processing (NSP) function in PE1 performs the packet processing, such as bit stuffing, PW-PW bridging, L2 encapsulation, shaping, and policing, for the HDLC packets that are forwarded to the PW termination point in PE1. The PW, which terminates at a logical port in the PE1, delivers the unaltered HDLC packets that are received across the MPLS network to the corresponding logical port on PE2 at the other end of the PW.

The PW termination points on each PE represent the operations that establish and maintain the PW and that encapsulate and decapsulate the HDLC packets. For more information, refer to the PW reference diagram packet processing that supports the HDLC PW as described in RFC 4618.

HDLC SAP-to-SAP Service

HDLC VLLs can be configured with both endpoints (SAPs) on the same 7705 SAR, which is referred to as an HDLC SAP-to-SAP or local HDLC service. HDLC SAP-to-SAP emulates local HDLC switching between two endpoints on the same 7705 SAR node.

HDLC Encapsulation

Figure 45 shows a typical HDLC VLL frame.

Figure 45: HDLC VLL Frame

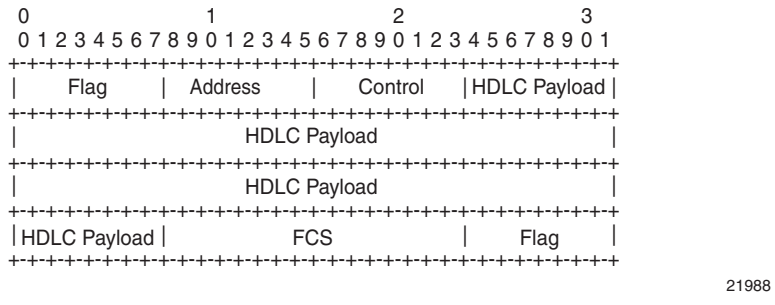
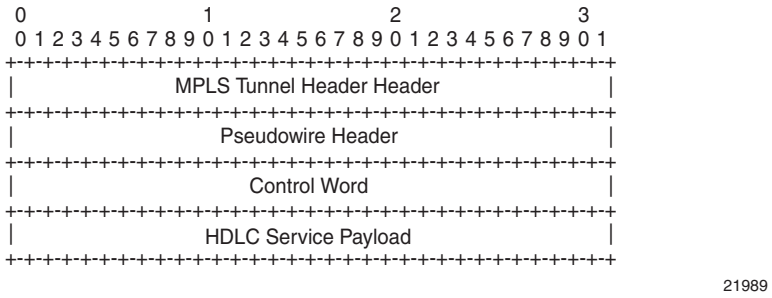


Figure 46 shows a typical HDLC VLL frame together with its MPLS tunnel encapsulation.

Figure 46: HDLC VLL Frame with MPLS Encapsulation



The native HDLC PDU is processed as follows:

- Flag—the HDLC flags are removed during encapsulation
- FCS—the FCS can use a 2-byte CRC-CCITT or a 4-byte CRC-32. The CRC is configurable. The default value is 2-byte. The FCS is removed during encapsulation.
- Address—HDLC address is retained
- Control—HDLC control is retained

The MPLS tunnel is used to transport the encapsulated HDLC across the PSN and the PW header is appended to the modified HDLC PDU as described in RFC 4618. The HDLC control word is inserted in the frame before the HDLC payload. See HDLC PW Control Word and Payload Size for information.

HDLC PW Control Word and Payload Size

To switch frames through the 7705 SAR switch fabric, the frame size must be 64 bytes or larger. If the HDLC frames received at the SAP are not 64 bytes or larger, the 7705 SAR pads the HDLC payload at the access ingress to ensure that the HDLC payload can be passed through the switch fabric and transported properly across the network. When padding occurs, the size of the original payload must be indicated in the length field of the control word so that at the termination of HDLC PW, the far-end device can correctly determine the actual size of the HDLC payload, remove the padding, and forward the original, non-padded HDLC payload to the access egress. See RFC 4618, section 4.1 part-2. Figure 47 shows the HDLC PW control word.

Figure 47: HDLC PW Control Word

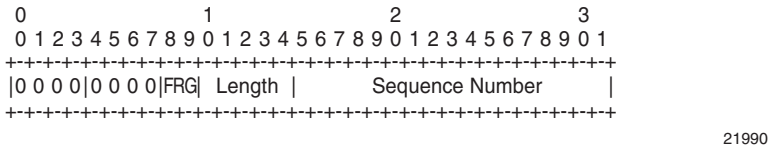


Table 29 describes the bits used in the HDLC PW control word.

Table 29: Control Word Bit Descriptions

Bit(s)	Description
Bits 0 to 3	The use of bits 0 to 3 is described in RFC 4385. The bits are always set to 0 to indicate the presence of the CW.
Flags	4 bits (bits 4 to 7). No flags are defined for the HDLC PW. The bits must be set to 0, and must be ignored by the PE.
FRG	2 bits (bits 8 and 9). FRG bits are not supported. The bits must be set to 0.
Length	6 bits (bits 10 to 15). The length represents the combined size of the CW and the HDLC payload. If the combined size is less than 64 bytes, this field must be populated to indicate the actual size of the HDLC payload that is received at the access ingress. At the access egress, this field is used to strip bytes that are appended for padding purposes. If the combined size of the CW and the HDLC payload exceeds 64 bytes, all bits in this field must be set to 0. See the control-word command description for more information.
Sequence Number	The sequence number is used to provide the common PW sequencing function as well as detection of lost packets. The sequence number field is not supported on the 7705 SAR.

See Pseudowire Control Word for more information.

Status Signaling

The HDLC PW supports status signaling in accordance with RFC 4618, section 5.1. When the PE detects a status change in the attachment circuit (AC) status, such as an AC physical link failure, or if the AC is administratively disabled, the PE sends the appropriate PW status notification message that corresponds to the HDLC AC status. The local PW status is also reflected in a PW status notification message, as described in RFC 4447, section 5.4.

OAM Propagation

The 7705 SAR supports OAM propagation between AC SAPs and the PW and vice-versa. For example, if no viable tunnel exists from the AC to the eLER, the status of the local SAP is set to the down state. Likewise, when a local SAP fails, the 7705 SAR sends a PW status message informing the far end of an AC ingress/egress fault. The far-end eLER then sets the status of the service and the SAP to the down state.

IP Interworking VLL (Ipipe) Services

This section provides information about the Ipipe service.

Topics in this section include:

- Ipipe Service Overview
- IP Interworking VLL Datapath
- CE IP Address Discovery and Distribution
- IP SAP-to-SAP Service
- Hardware Support for Interworking IP PWs
- Control Word
- Termination at Access
- Traffic Management
- Status Signaling
- IP Filters (Ipipe)

Ipipe configuration information is found under the following topics:

- Common Configuration Tasks
- Configuring VLL Components
- Service Management Tasks

See Service Support for information on the adapter cards and chassis that support IP interworking VLL services.

Ipipe Service Overview

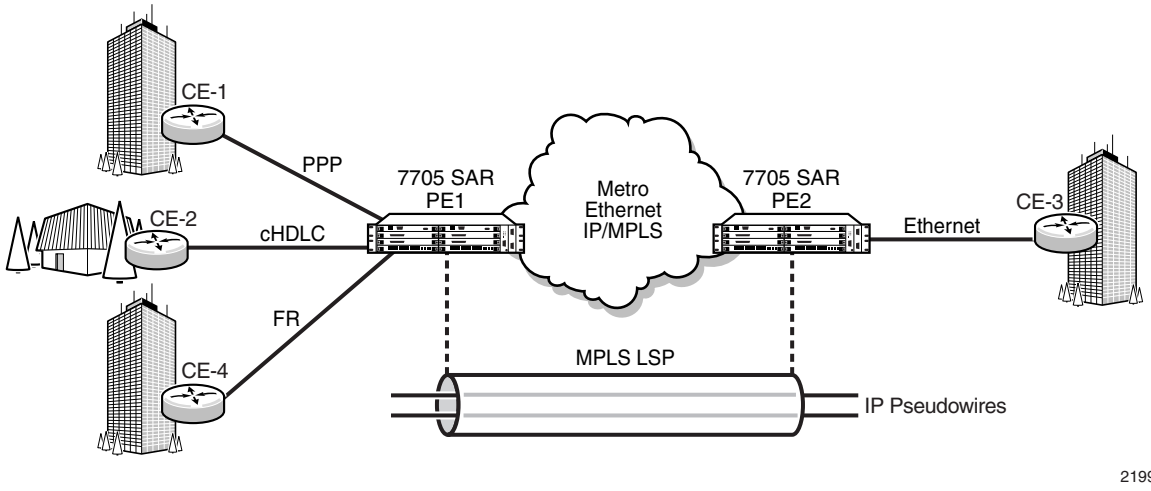
An Ipipe pseudowire (IP PW) enables service interworking between different link layer technologies and network interworking between connections with the same link layer technologies. IP PWs provide an efficient means to connect Layer 3 IP traffic to the IP/MPLS network, even without access to VLANs.

An Ipipe is a point-to-point Layer 2 service where the customer data is encapsulated and transported across an MPLS or IP network. An Ipipe service transparently forwards all packets received on one SAP to the other SAP. No native IP routing of customer packets occurs.

IP interworking allows connections to be created with any combination of PPP, MLPPP, Ethernet, LAG, FR, and Cisco HDLC (cHDLC) SAPs, but the payload must always be IP. Ipipes can be used to transport IP payloads more efficiently than Epipes because an Ipipe service does not need to forward the Ethernet header information.

Figure 48 shows an example of IP connectivity between a host attached to a point-to-point access circuit (FR, cHDLC, and PPP) with routed PDU IPv4 encapsulation and a host attached to an Ethernet interface. Both hosts are on the same LAN segment.

Figure 48: IP Pseudowires Between SAR Nodes



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

IP Interworking VLL Datapath

In order to be able to forward IP packets between CE 1 and CE 3 in Figure 48, PE 2 is configured with both CE 1 and CE 3 IP addresses. These are host addresses and are entered in the /32 format.

Note: Addresses can be configured manually or by using CE Address discovery.

PE 2 maintains an ARP cache context for each IP interworking VLL and responds to ARP request messages received on the Ethernet SAP. PE 2 responds with the Ethernet SAP configured MAC address as a proxy for any ARP request for the CE 1 IP address. PE 2 silently discards any ARP request messages received on the Ethernet SAP for addresses other than CE 1. Likewise, PE 2 silently discards any ARP request messages with source IP addresses other than CE 3. In all cases, PE 2 keeps track of the association of IP to MAC addresses for ARP requests it receives over the Ethernet SAP. All entries are subject to aging.

In order to forward unicast frames destined for CE 3, PE 2 needs to know the MAC address of CE 3. If there is no entry in the ARP cache, PE 2 sends an ARP request message for the CE 3 MAC address over the Ethernet SAP.

IP broadcast and IP multicast packets are sent on the Ethernet SAP using the broadcast or direct-mapped multicast MAC address.

In order to forward unicast frames destined for CE 1, PE 2 validates the MAC destination address of the received Ethernet frame. It should match that of the Ethernet SAP. PE 2 then removes the Ethernet header and encapsulates the IP packet directly into a pseudowire with or without the optional control word. PE 1 removes the pseudowire encapsulation and forwards the IP packet over the SAP using PPP encapsulation.

When a packet reaches the access egress and the configured SAP is over a VLAN, the node pushes (inserts) the appropriate VLAN tag into the Ethernet frame header before forwarding the Ethernet frame out of the port. Ethernet frames at the access egress can also be marked with appropriate dot1 priority bits if the dot1 priority QoS profile is assigned to the forwarding class configuration.

Ethernet frames mapped to an Ipipe service can have a maximum of two VLAN tags. Frames with more than two VLAN tags are dropped at the Ipipe access ingress SAP.

At access ingress, PE 1 performs proxy PPP negotiation and provides the IP address of the remote CE 3 device to CE 1 during IPCP negotiation using the IP-Address option.

A PE does not flush the ARP cache unless the SAP goes administratively or operationally down. The PE with the Ethernet SAP sends unsolicited ARP requests to refresh the ARP cache according to the refresh interval. ARP requests are staggered at an increasing rate if no reply is received to the first unsolicited ARP request. The refresh interval is configurable using the `mac-refresh` CLI command.

CE IP Address Discovery and Distribution

The following subsections describe the customer edge (CE) IP address discovery methods and how the CE IP address is distributed to remote PE nodes. Topics include:

- Manual IP Configuration
- CE IP Address Discovery Using ARP
- CE IP Address Discovery for Frame Relay
- CE IP Address Discovery for cHDLC
- CE IP Address Discovery for PPP/MLPPP
- CE IP Address Distribution to Remote PE Nodes

Manual IP Configuration

Manually configured IP addresses are supported for all attachment circuit types. No further mechanisms for detecting the local or remote CE IP address are required. The PE responds to all ARP requests arriving on Ethernet attachment circuits by replying with the local interface MAC address and the remote CE IP address.

CE IP Address Discovery Using ARP

Ethernet attachment circuits can be configured to use ARP messages that are received from the CE device to determine the local CE IP address. The PE waits for an ARP request from the CE in order to learn the IP address that is associated with the MAC address of the CE. When a valid ARP request is received by the PE from the CE, the ARP cache on the PE is populated with the CE IP/MAC entry. The PE accepts any ARP request message that it receives over the Ethernet SAP and updates the ARP cache entries with no further check. The PE does not validate the source IP address of the ARP request message nor does it try to match the IP address in the ARP request with the programmed IP address.

The 7705 SAR always replies to an ARP request message that is received over the Ethernet SAP. The 7705 SAR replies with the SAP MAC address and a source IP address of the IP address being ARPed without any further check of the latter.

If the SAP status changes to operationally down or if an operator manually clears the ARP cache, the system flushes the ARP cache and the CE address discovered on the SAP is cleared. When the SAP comes into service initially or after a failure, an unsolicited ARP request is not sent over the Ethernet SAP. In the case where multiple ARP messages are received from different CE devices, the last received message prevails and the ARP cache is populated with the newly received information.

An SNMP trap is generated whenever an ARP entry or IPv4 CE address entry is discovered or changed for an Ipipe service.

Prior to Release 7.0, Ipipe IPv4 ARP packets were transmitted by bypassing the service SAPs. These packets are now sent using the service SAP. This enables normal service SAP functions such as QoS, statistics, and filtering to apply to these packets.

CE IP Address Discovery for Frame Relay

Frame relay access circuits use INVARP to learn a local CE address and to propagate the remote CE address.

CE IP Address Discovery for cHDLC

The cHDLC access circuits do not need to discover the IP address of the local and remote CE for point-to-point interfaces. The IP addresses remain 0.0.0.0.

CE IP Address Discovery for PPP/MLPPP

The address of a locally attached CE device can be learned via IPCP. If the Ipipe uses a spoke SDP, when the 7705 SAR sends the label mapping message, this learned address is not included in the address list TLV in the interface parameters field of the pseudowire FEC TLV.

CE IP Address Distribution to Remote PE Nodes

The PE includes an IP address list TLV in the label mapping message of the PW FEC in order to communicate the local CE IP address to the remote PE. If the IP address is set to 0.0.0.0, it is assumed that the connected CE IP address is unknown. For point-to-point connections such as frame relay and cHDLC, an IP address of 0.0.0.0 does not affect the PW status or stop the flow of IP traffic through the Ipipe. Broadcast interfaces such as Ethernet must learn the local CE IP address and MAC relationship before unicast traffic can be sent, but the remote PE IP address can remain as 0.0.0.0. The value of the remote PE IP address is always 0.0.0.0 when the remote PE access circuit is PPP/MLPPP (IPCP) or cHDLC.

If the CE IP address changes, an LDP notification message is sent to the remote PE with the new IP address of the CE.

IP SAP-to-SAP Service

IP VLLs can be configured with both endpoints (SAPs) on the same 7705 SAR, which is referred to as an IP SAP-to-SAP or local IP service. IP SAP-to-SAP emulates local IP switching between two endpoints on the same 7705 SAR node.

Hardware Support for Interworking IP PWs

Table 30 lists the hardware that supports interworking IP PWs.

Table 30: Hardware Support for Interworking PWs

MDA Type	Interworking PWs						
	Eth to IP PW	PPP to IP PW	MLPPP to IP PW	MC-MLPPP to IP PW	FR to IP PW	cHDLC to IP PW	ATM to IP PW
8-port Ethernet Adapter card	✓						
6-port Ethernet 10Gbps Adapter card	✓						
8-port Gigabit Ethernet Adapter card	✓						

Table 30: Hardware Support for Interworking PWs (Continued)

MDA Type	Interworking PWs						
	Eth to IP PW	PPP to IP PW	MLPPP to IP PW	MC-MLPPP to IP PW	FR to IP PW	cHDLC to IP PW	ATM to IP PW
10-port 1GigE/1-port 10GigE X-Adapter card							
— x1-10gb-sf+ mode							
— x10-1gb-sfp mode	✓						
16-port T1/E1 ASAP Adapter card, version 1							
— Channelized		✓					
— Clear channel		✓	✓	✓			
16-port T1/E1 ASAP Adapter card, version 2							
— Channelized		✓			✓	✓	
— Clear channel		✓	✓	✓	✓	✓	
32-port T1/E1 ASAP Adapter card							
— Channelized		✓			✓	✓	
— Clear channel		✓	✓	✓	✓	✓	
4-port DS3/E3 Adapter card							
— on n x DS0							
— on DS1/E1							
— Clear channel (DS3 or E3)					✓		
2-port OC3/STM1 Channelized Adapter card							
— on n x DS0							
— on VT1.5/ VC12		✓	✓				

Table 30: Hardware Support for Interworking PWs (Continued)

MDA Type	Interworking PWs						
	Eth to IP PW	PPP to IP PW	MLPPP to IP PW	MC-MLPPP to IP PW	FR to IP PW	cHDLC to IP PW	ATM to IP PW
12-port Serial Data Interface card							
— on RS-232 ports							
— on V.35/X.21 ports		✓			✓	✓	
4-port SAR-H Fast Ethernet module	✓						
6-port SAR-M Ethernet module	✓						
7705 SAR-A	✓	✓	✓	✓			
7705 SAR-F	✓	✓	✓	✓			
7705 SAR-H	✓	✓	✓	✓			
7705 SAR-M	✓	✓	✓		✓	✓	
7705 SAR-W	✓						
7705 SAR-X	✓	✓	✓				

Control Word

IP interworking VLL supports an optional control word. The control word contains protocol control information and must be used for any pseudowire service that transports packets that are less than 64 bytes; see Pseudowire Control Word for more information.

Termination at Access

The following sections describe the termination requirements and features for interworking pseudowires. Topics include:

- PPP and MLPPP Termination
- FR Termination
- cHDLC Termination
- PPP and MLPPP Termination on 2-port OC3/STM1 Channelized Adapter Cards

PPP and MLPPP Termination

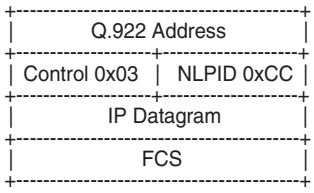
For PPP termination, ports must be configured in access mode with IPCP encapsulation. Access ports on the 12-port Serial Data Interface card or 2-port OC3/STM1 Channelized Adapter card must be configured with IPCP encapsulation to support PPP/MLPPP termination for an interworking pseudowire. The ppp-auto encapsulation type applies only to network-side terminations.

FR Termination

The following features are supported for FR termination:

- ports must be configured for access mode according to Q.922 Annex A frame format
- RFC 2427, *Multiprotocol Interconnect over Frame Relay, routed PDU encapsulation of an IPv4 packet*
- UNI DCE and DTE
- NLPID (0xCC) IPv4 encapsulation
- LMI. See FR LMI.
- configurable service MTU of 1 to 2048 bytes. When encapsulated in the IPv4 PDU, the FR flags, FCS, and FR header are removed. The resulting data is the Ipipe payload represented in the IP Datagram field in Figure 49.

Figure 49: FR Header with NLPID Support for IP Interworking PW

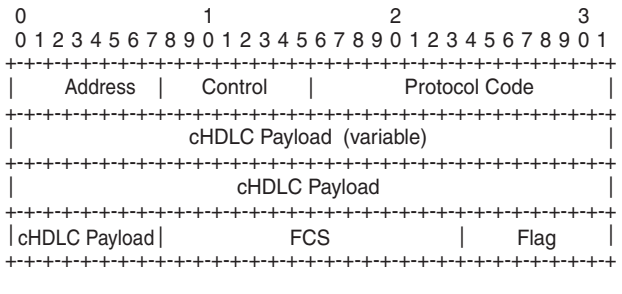


cHDLC Termination

The following features are supported for cHDLC termination:

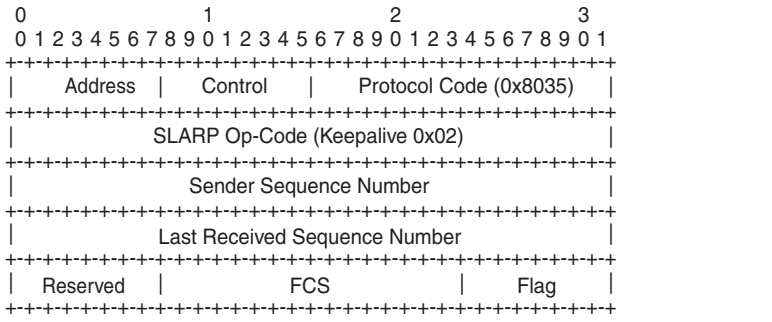
- ports must be configured for access mode according to the cHDLC frame format. Refer to the 7705 SAR OS Interface Configuration Guide for information about how to configure the `encap-type` command to `cisco-hdlc`.
- IP encapsulation (0x0800). Figure 50 shows the cHDLC header frame.
- keepalive support using SLARP (0x8035). Figure 51 shows the SLARP keepalive frame.
- configurable service MTU of 1 to 2048 bytes. When encapsulated in the IPv4 PDU, the cHDLC packet shown in Figure 50 is modified to remove the flag, address, control, protocol ID, and FCS.

Figure 50: cHDLC Header Frame



Note: cHDLC encapsulation cannot be used to transmit HDLC frames into an Ipipe.

Figure 51: SLARP Keepalive Frame



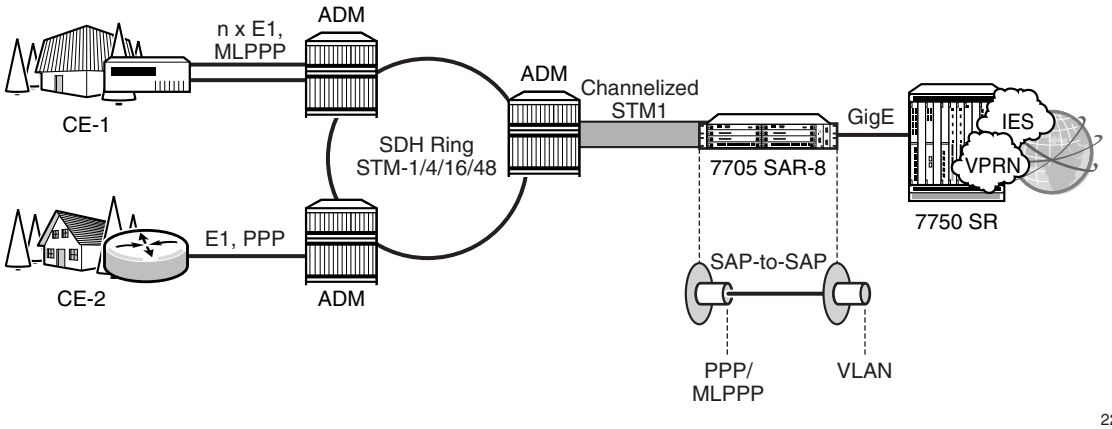
Refer to the 7705 SAR OS Interface Configuration Guide for more information about how to configure access ports for interworking IP PWs.

PPP and MLPPP Termination on 2-port OC3/STM1 Channelized Adapter Cards

One use for PPP and MLPPP termination to IP pseudowires is in fixed backhaul networks with customers connected to the network via DS1/E1 or *n* x DS1/E1 rates.

Figure 52 shows an example of a network using PPP and MLPPP encapsulation to IP PWs on a 2-port OC3/STM1 Channelized Adapter card.

Figure 52: PPP and MLPPP to IP PWs on 2-port OC3/STM1 Channelized Adapter Cards



Customer traffic aggregation remains on the available SONET/SDH infrastructure and is sent to a 2-port OC3/STM1 Channelized Adapter card on the 7705 SAR, where the PPP/MLPPP encapsulation termination takes place. Using an IP PW with PPP/MLPPP encapsulation from a SAP-to-VLAN SAP, the 7705 SAR passes the data to a 7750 SR. The 7750 SR offers the necessary scale of IP services, such as IES or VPRN, needed to handle the large amount of traffic coming from each OC3/STM1 port.

After terminating PPP/MLPPP encapsulated traffic on a channelized OC3/STM1 port, the 7705 SAR switches the IP traffic, without performing any IP lookup, to another SAP over an Ipipe. IP traffic is then passed on to the 7750 SR with a unique VLAN identifier on a per-connection basis.

Refer to the 7705 SAR OS Interface Configuration Guide for more information about how to configure 2-port OC3/STM1 Channelized Adapter card ports for interworking IP PWs and PPP/MLPPP termination.

Traffic Management

Traffic Management for FR IW SAPs

Traffic management of interworking SAPs is achieved by applying ingress and egress QOS policies to SAPs. Access ingress forwarding classes are determined by inspecting the DSCP marking of the IP packets to determine the queuing and the EXP bit setting of packets. If all traffic for a DLCI must be classified as a single forwarding class, you can define a policy to classify all DSCP values from the IP header to one forwarding class.

Status Signaling

FR LMI

LMI enables the exchange of information between the CE and the LER about the status of the link, device, and logical circuit. The frame relay CE device (the DTE) begins an LMI exchange by sending a Status Enquiry message. DTE status enquiry message supports link integrity verification requests. The frame relay network (the DCE) completes the exchange by responding with a Status message.

The DCE status messages support the following response types:

- full status
- link integrity verification
- single PVC status

The LMI protocol operates over a dedicated PVC of a frame relay link. Since LMI occupies its own PVC, its link signaling cannot congest or interfere with traffic on the PVCs that carry subscriber data.

The following standards are supported:

- ANSI T.617 Annex D
- ITU-T Q.922 Annex A
- Cisco Rev.1

IP Filters (Ipipe)

IP filters are applied to Ipipe SAPs in the ingress direction, as described below. For a full list of entities to which IP filters can be applied, see IP Filter Policies.

IP pseudowires are generally used to transparently switch traffic across an MPLS network to the far end. However, in some cases, the traffic that is switched over the network, consuming valuable bandwidth, is just discarded at the other end of the pseudowire. As well, with the 7705 SAR expanding into areas such as vertical markets, and with local area networks being connected to the 7705 SAR Ethernet ports, an increasing amount of traffic must stay local and not pass through the MPLS network to the far end. By using IP filters at the access ingress, operators can determine what traffic is passed through the pseudowire and therefore use the network links more efficiently.

Another use for IP filters is in cases where a customer router is connected to an access port on the 7705 SAR with PPP/MLPPP encapsulation. The service provider may want to filter incoming traffic from the customer at the boundaries of the network.

IP filters can also be used for security purposes, by allowing access only to designated services (for example, allowing e-mail and FTP services while disallowing Telnet services) at the origin of the traffic.

IP filter policies specify either a forward or a drop action for packets, based on information specified in the match criteria. Within each filter policy, you can create entries that define matching criteria.

The same IP filter policy can be assigned to any entity (network interfaces, IP pseudowires, Ethernet pseudowires, IES, and VPRN services), all of which can be configured on the same adapter card. For example, a filter policy defined as “filter-5” can be assigned to multiple Ipipe SAPs and, simultaneously, to network interfaces on the same adapter card.

A filter policy assigned to an entity on one adapter card can also be assigned to any entity on another adapter card. For example, a filter policy defined as “filter-2” can be assigned to an Ipipe on an Ethernet Adapter card and to a network interface on another Ethernet Adapter card.

Assigning the same filter policy to different entities on a card counts as using one filter policy.

Configuration and assignment of filter policies is similar for network interfaces, IES management SAPs, Ethernet and IP pseudowire SAPs, VPRN and IES SAPs and spoke SDPs, and VPLS SAPs and SDPs (spoke and mesh). This guide describes the assignment of filter policies to SAPs and SDPs. Refer to the 7705 SAR OS Router Configuration Guide, “Filter Policies”, for information on configuring filter policies and assigning them to network interfaces.

Pseudowire Switching

Topics in this section include:

- Overview
- Pseudowire Switching with Pseudowire Redundancy
- Pseudowire Switching Behavior
- Pseudowire Switching TLV

Overview

The pseudowire switching feature provides the user with the ability to create a VLL service by cross-connecting two spoke SDPs.

Services with one SAP and one spoke SDP are created normally on the PE; however, when a pseudowire originates at customer equipment and not the 7705 SAR, the target destination of the SDP is the 7705 SAR pseudowire switching node instead of the remote PE. In such cases, the user must configure a VLL service on the pseudowire switching node (the 7705 SAR, in this case) using the two SDPs and no SAP. This creates a VLL service that travels over two different types of tunnels. The first pseudowire segment connects the Node B to the 7705 SAR, and the second segment connects the 7705 SAR to a 7750 SR node. Figure 53 shows an example of a VLL pseudowire switching service.

For the 7750 SR node, there is no implementation change required. The pseudowire segment is treated as a 7705 SAR-initiated dynamic ATM pseudowire. The 7705 SAR signals for the ATM pseudowire and negotiates all required parameters, including the pseudowire label, control word, and VCCV type.

The pseudowire switching node acts in a passive role with respect to signaling of the pseudowires. The node waits until one or both of the PEs send 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. For example, the 7705 SAR assumes the pseudowire payload from the Node B, then signals for the ATM pseudowire to the far end and negotiates all required parameters, including the pseudowire label, control word, and VCCV type.

The switching pseudowire inserts a pseudowire switching point TLV indicating the system address in the label mapping message. This TLV is useful because it allows for troubleshooting of the path of the pseudowire, especially if multiple pseudowire switching points exist between the two PEs. The TLV also helps in loop detection of T-LDP signaling messages, where a switching point receives a label mapping message that it already relayed.

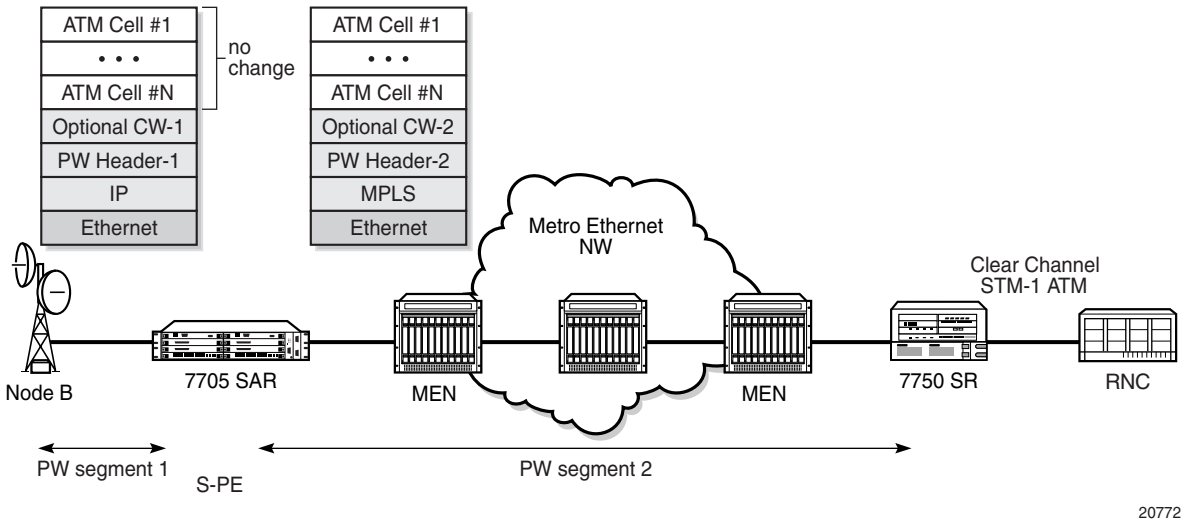
The switching point TLV is inserted in pseudowire status notification messages when they are sent end-to-end or from a pseudowire switching node towards a destination PE. See Pseudowire Switching TLV for more information.

Pseudowire OAM is supported for dynamic switching pseudowires and allows the 7705 SAR pseudowire switching node to relay end-to-end pseudowire status notification messages between the two PEs. The 7705 SAR can generate a pseudowire status and send it to one, or both, of the PEs by including its system address in the pseudowire switching point TLV. This allows a 7705 SAR PE to identify the origin of the pseudowire status notification message.

The pseudowire segment between the 7705 SAR and the 7750 SR supports OAM tools as well as BFD; however, since Node Bs are not dual-homed, OAM tools operating on this pseudowire segment can only provide informative messages. The 7705 SAR responds to Node B-initiated ping packets if the destination IP address is the system IP or interface IP address of the 7705 SAR. If the IP packet does not contain a SAR destination IP address, the 7705 SAR does not respond, and instead, forwards the packet.

Figure 53 shows an example of static-simplex to dynamic-simplex pseudowire switching. This service consists of a SAP and a spoke SDP. However, the target destination of the SDP is not the remote PE but the pseudowire switching node. In addition, the user configures a VLL service on the pseudowire switching node using the two SDPs.

Figure 53: Simplex to Simplex Pseudowire Switching



Configuration examples can be found in Configuring PW Switching.

Table 31 shows the pseudowire switching options supported on the 7705 SAR.

Table 31: Supported Pseudowire Switching Options

PW Segment 1		PW Segment 2								
Tunnel	**PW**	**Tunnel**	**PW**	**E pipe**	**F pipe**	**H pipe**	**I pipe**	**A pipe**	**Cpipe -ces**	**Cpipe -satop**
RSVP-TE, LDP, Static MPLS, IP, or GRE/IP	Static	RSVP-TE, LDP, Static MPLS, IP, or GRE/IP	T-LDP (with or without pseudowire redundancy)	✓	✓	✓	✓	✓	✓	✓
RSVP-TE, LDP, Static MPLS, IP, or GRE/IP	T-LDP	RSVP-TE, LDP, Static MPLS, IP, or GRE/IP	T-LDP	✓	✓	✓	✓	✓	✓	✓
IP	Static [1]	MPLS-Static	Static [1]	✓	✓	✓	✓	✓	✓	✓

Note:

1. A static- to-static pseudowire switching service will fail if there is more than one S-PE node between the T-PE nodes.

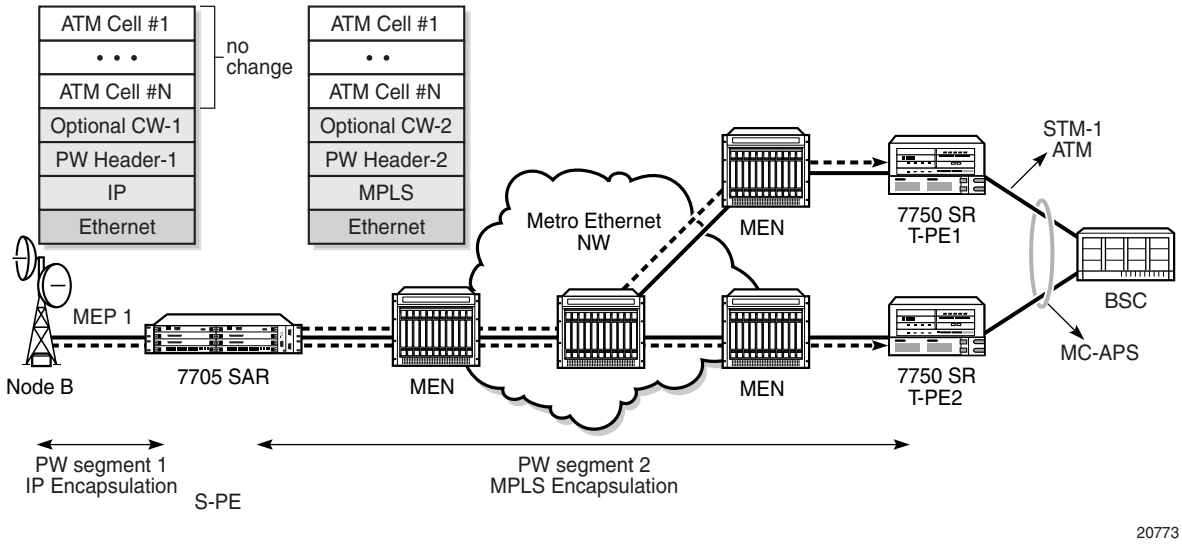
The 7705 SAR supports use of an optional control word on both pseudowire segments. Use of a control word is negotiated by the 7705 SAR and 7750 SR during the signaling phase. The 7705 SAR and 7750 SR negotiate during the signaling phase even if a control word is not used. If a control word is used, the 7705 SAR generates it and configures it with all 0s.

When the 7705 SAR appends or strips the control word to support VCCV ping type 1, the TTL value of the switched pseudowire is reset to 255. If the control word is present on the ingress pseudowire packet and it is not removed because it is on an end-to-end service, the 7705 SAR reduces the pseudowire TTL by 1 at the time of pseudowire switching.

Pseudowire Switching with Pseudowire Redundancy

Pseudowire switching with pseudowire redundancy supports one redundant pseudowire with up to four redundant spoke SDPs. Figure 54 shows an example of a network with simplex to redundant pseudowire switching.

Figure 54: Simplex to Redundant Pseudowire Switching

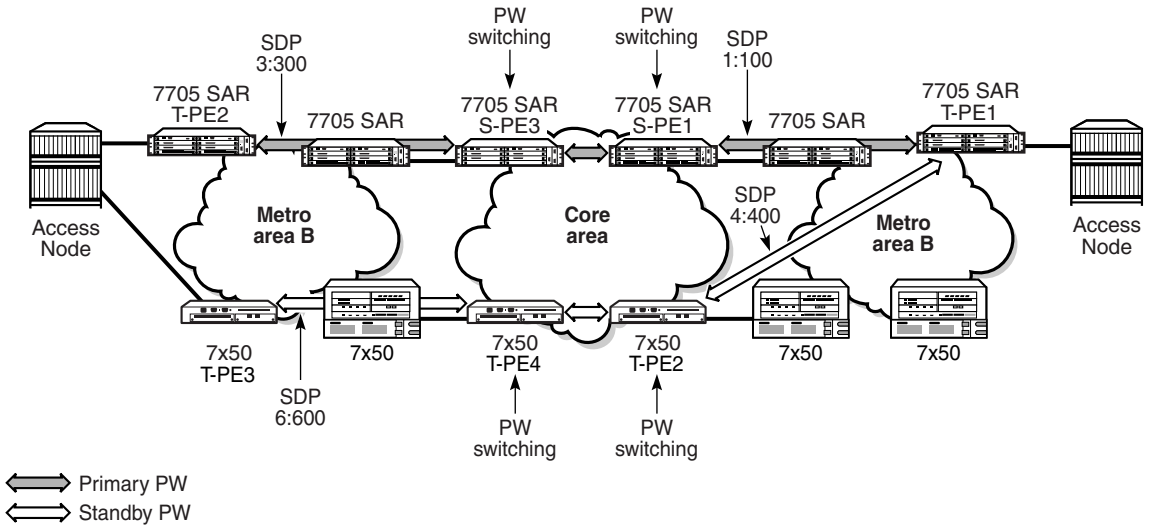


To enable pseudowire redundancy, the first pseudowire segment must be a static pseudowire (that is, T-LDP disabled). The second pseudowire segment can then be configured with up to four redundant spoke SDPs. See Pseudowire Redundancy for instructions on configuring redundancy. Pseudowire switching with pseudowire redundancy also supports standby signaling. See Active/Standby Mode for Pseudowire Redundancy (Standby Signaling) for more information.

Pseudowire Switching Behavior

In the network in Figure 55, T-PE nodes act as masters and pseudowire switching nodes (S-PE nodes) act as slaves for the purpose of pseudowire signaling. Switching nodes need to pass the SAP interface parameters of each PE to the other PE. T-PE1 sends a label mapping message for the Layer 2 FEC to the peer pseudowire switching node; for example, S-PE1. It includes the SAP interface parameters, such as MTU, in the label mapping message. S-PE1 checks the FEC against the local information, and if a match exists, it appends the optional pseudowire switching point TLV to the FEC TLV in which it records its system address. T-PE1 then relays the label mapping message to S-PE2. S-PE2 performs a similar operation and forwards a label mapping message to T-PE2.

Figure 55: Pseudowire Switching Network



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 affect the spoke SDP cross-connect only when both directions of the pseudowire have been signaled and matched.

Pseudowire status notification messages can be generated by the T-PE nodes and/or the S-PE nodes. Pseudowire status notification messages received by a switching node are processed and then passed on to the next hop. An S-PE node appends the optional pseudowire switching TLV, with its system address added to it, to the FEC in the pseudowire status notification message only if it originated the message or the message was received with the TLV in it. If the S-PE node is not the originator or if there is no TLV with the system address in the message, this means that the message was originated by a T-PE node. In this case, the S-PE processes and passes 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 7705 SAR 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 the value of the pseudowire segment to the next hop.
- When the local operational status for any of the spokes is down, the S-PE always sends SDP-binding down status bits regardless if the received status bits from the remote node indicated SAP up or down or SDP-binding up or down.

When one segment of the pseudowire cross-connect at the S-PE is static while the other is signaled using T-LDP, the S-PE operates much like a T-PE from a signaling perspective and as an S-PE from a data plane perspective. The S-PE signals a label mapping message as soon as the local configuration is complete. The control word C-bit field in the pseudowire FEC is set to the value configured on the static spoke SDP.

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

In a static-to-dynamic pseudowire switching service, it is possible for the end nodes of the static pseudowire segment to be misconfigured. In this case, an S-PE or T-PE node may receive packets with the wrong encapsulation. If this happens, an invalid payload might be forwarded over the pseudowire or the SAP respectively.

Furthermore, if the S-PE or T-PE node is expecting the control word in the packet encapsulation, and the received packet arrives with no control word but the first nibble below the label stack is 0x0001, the packet may be mistaken for a VCCV OAM packet and may be forwarded to the CSM. In that case, the CSM will perform a check of the IP header fields. If any of the fields fail the check, the VCCV packet will be discarded.

Pseudowire Switching With IP Tunnels

You cannot enable T-LDP dynamic pseudowire establishment on pseudowire switching segments using IP tunnels if the second pseudowire segment is configured for pseudowire redundancy. The pseudowire label, control word, VCCV type, and so on, must be configured manually.

On the first pseudowire segment, ATM pseudowires are natively transported over IP using GRE encapsulation with the IP type set to 0x8847, or IP encapsulation. The destination IP address of pseudowire packets received from the Node B must be set to the system IP or interface IP address of the 7705 SAR. Similarly, the destination IP address of pseudowire packets received from the 7750 SR must be set to the system IP or interface IP address of the Node B. Node B management traffic is transported over the same Ethernet link between the 7705 SAR and the Node B. The 7705 SAR forwards the management IP traffic to its destination based on longest prefix match.

On the second pseudowire segment, ATM pseudowires are transported over MPLS tunnels. The MPLS tunnels can also be used to transport additional cell site traffic, such as BTS traffic using TDM pseudowires, or LTE base station traffic using IP or Ethernet pseudowires.

Pseudowire Switching TLV

Figure 56 shows the format of the pseudowire switching TLV and Table 32 describes the fields.

Figure 56: Pseudowire Switching TLV

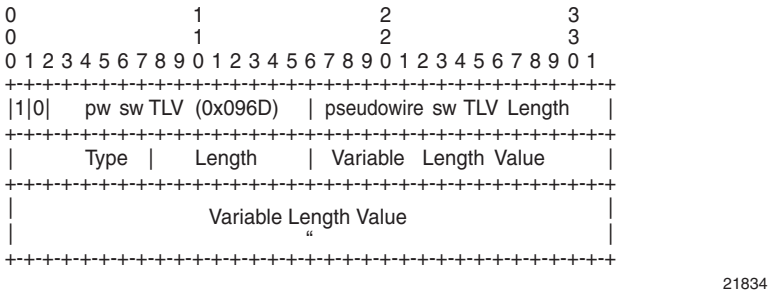


Table 32: Pseudowire Switching TLV Field Descriptions

Field	Description
PW sw TLV Length	Specifies the total length of all the following pseudowire switching point TLV fields in octets
Type	Encodes how the Value field is to be interpreted
Length	Specifies the length of the Value field in octets
Value	Octet string of Length octets that encodes information to be interpreted as specified by the Type field

The following list describes details specific to pseudowire switching point sub-TLVs:

- pseudowire ID of last pseudowire segment traversed
- pseudowire switching point description string—an optional description string of text up to 80 characters long
- IP address of pseudowire switching point—an optional sub-TLV containing the IPv4 address of the pseudowire switching point
- MH VCCV capability indication

VLL Service Considerations

This section describes the general 7705 SAR service features and any special capabilities or considerations as they relate to VLL services.

Topics in this section include:

- Service Support
- SDPs
- SAP Encapsulations and Pseudowire Types
- ATM PWE3 N-to-1 Cell Mode Encapsulation
- SAP Aggregation Groups
- QoS Policies
- IP Filter Policies
- MTU Settings
- QinQ (VLL Service)
- Pseudowire Control Word
- Pseudowire Redundancy
- Active/Standby Mode for Pseudowire Redundancy (Standby Signaling)

Service Support

The section describes hardware support for the following VLL services:

- ATM
- Ethernet
- Frame relay
- TDM
- HDLC
- IP interworking

ATM

ATM VLL service is supported on the following:

- T1/E1 ports on the 2-port OC3/STM1 Channelized Adapter card (when the port is configured for ATM or IMA)
- T1/E1 ports on the 4-port DS3/E3 Adapter card (when the port is configured for ATM or IMA)
- 4-port OC3/STM1 Clear Channel Adapter card (when the port is configured for ATM)
- 16-port T1/E1 ASAP Adapter card (when the port is configured for ATM or IMA)
- 32-port T1/E1 ASAP Adapter card (when the port is configured for ATM or IMA)
- T1/E1 ports on the 7705 SAR-F
- T1/E1 ports on the 7705 SAR-M (variants with T1/E1 ports)

Ethernet

Ethernet VLL service is supported on the following:

- 6-port Ethernet 10Gbps Adapter card
- 8-port Ethernet Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card
- Ethernet ports on the 7705 SAR-A (both variants)
- Ethernet ports on the 7705 SAR-F
- Ethernet ports on the 7705 SAR-H
- Ethernet ports on the 7705 SAR-Hc
- Ethernet ports on the 7705 SAR-M (all variants)
- 7705 SAR-W
- Ethernet ports on the 7705 SAR-Wx (all variants)
- Ethernet ports on the 7705 SAR-X
- 6-port DSL Combination module
- 8-port xDSL module
- GPON module
- 4-port SAR-H Fast Ethernet module
- 6-port SAR-M Ethernet module

Frame Relay

Frame relay VLL service is supported on the following:

- DS3/E3 clear channel or channelized DS1/E1 ports on the 4-port DS3/E3 Adapter card
- V.35 and X.21 serial ports on the 12-port Serial Data Interface card
- 16-port T1/E1 ASAP Adapter card, version 2
- 32-port T1/E1 ASAP Adapter card
- T1/E1 ports on the 7705 SAR-M (variants with T1/E1 ports)

TDM

TDM VLL service is supported on the following:

- T1/E1 ports and DS3 channels on the 2-port OC3/STM1 Channelized Adapter card
- T1/E1 ports (DS3 ports only) and DS3/E3 ports on the 4-port DS3/E3 Adapter card
- T1/E1 ports on the 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card
- 6-port E&M Adapter card (when the port is configured for cem encapsulation)
- 8-port Voice & Teleprotection card
- 12-port Serial Data Interface card
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- 4-port T1/E1 and RS-232 Combination module
- T1/E1 ports on the 7705 SAR-A (variant with T1/E1 ports)
- T1/E1 ports on the 7705 SAR-F
- RS-232 serial ports on the 7705 SAR-Hc
- T1/E1 ports on the 7705 SAR-M (variants with T1/E1 ports)
- T1/E1 ports on the 7705 SAR-X

HDLC

HDLC VLL service is supported on the following:

- V.35 and X.21 serial ports (super-rate speeds only) on the 12-port Serial Data Interface card
- 16-port T1/E1 ASAP Adapter card, version 2
- 32-port T1/E1 ASAP Adapter card

- T1/E1 ports on the 7705 SAR-M (variants with T1/E1 ports)

IP Interworking

IP interworking VLL service is supported on the following:

- 2-port OC3/STM1 Channelized Adapter card (when the payload is configured as vt1.5/vc12)
- DS3/E3 clear channel ports on the 4-port DS3/E3 Adapter card (when the port is configured for frame-relay encapsulation)
- 6-port Ethernet 10Gbps Adapter card
- 8-port Ethernet Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card
- V.35 and X.21 serial ports (super-rate speeds only) on the 12-port Serial Data Interface card (when the port is configured for ipcp, frame-relay, or cisco-hdlc encapsulation)
- 16-port T1/E1 ASAP Adapter card (on PPP/MLPPP connections)
- 32-port T1/E1 ASAP Adapter card (on PPP/MLPPP connections)
- all ports on the 7705 SAR-A (on PPP/MLPPP connections on the T1/E1 ports)
- all ports on the 7705 SAR-F (on PPP/MLPPP connections on the T1/E1 ports)
- 7705 SAR-H
- Ethernet ports on the 7705 SAR-Hc
- all ports on the 7705 SAR-M (on PPP/MLPPP connections on the T1/E1 ports; variants with T1/E1 ports also support frame relay and HDLC SAPs on the T1/E1 ports)
- 7705 SAR-W
- Ethernet ports on the 7705 SAR-Wx (all variants)
- 7705 SAR-X
- 4-port SAR-H Fast Ethernet module
- 6-port SAR-M Ethernet module



Note: MPLS and VLL service over MPLS are not supported on access ports.

SDPs

The most basic SDPs must have the following characteristics:

- a locally unique SDP identification (ID) number and a VC-ID
- the system IP address of the far-end 7705 SAR routers
- an SDP encapsulation type—GRE, IP, or MPLS

SDP Statistics for VLL Services

The 7705 SAR supports local CLI-based and SNMP-based statistics collection for each VC used in the SDPs. This allows for traffic management of tunnel usage by the different services and, with aggregation, the total tunnel usage.

SAP Encapsulations and Pseudowire Types

The section describes encapsulations and PW types for the following VLL services:

- Apipe
- Cpipe
- Epipe
- Fpipe
- Hpipe
- Ipipe

Apipe

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

Note: ATM SAP-to-SAP connections are not supported for pseudowire packets using N-to-1 cell mode encapsulation where N > 1.

As an alternative to signaled pseudowires, manual configuration of pseudowires is also supported.

The Apipe service supports virtual trunking, VP connections, and VC connections, which are identified by specifying the `vc-type` when provisioning the Apipe. When `vc-type` is `atm-cell`, ATM transparent cell transport mode is used for VT connections. The N-to-1 cell transport mode is supported for VC and VP services (see ATM PWE3 N-to-1 Cell Mode Encapsulation). For VCCs, the value of N can be 1 (N = 1) or greater than 1 (N > 1). The value of N is always 1 for VPCs.

The supported PW service types are 0x0009 (for ATM N-to-1 VCC cell mode, where N is ≥ 1), 0x000A (for ATM N-to-1 VPC cell mode, where N = 1) and 0x0003 (for ATM transparent cell transport mode). Refer to RFC 4717 and RFC 4446 for more information.

Cpipe

Cpipe service supports CESoPSN and SAToP encapsulation over MPLS or GRE tunnels to connect to the far-end circuit. Cpipes support SAP-to-SAP and SAP-to-spoke SDP binding with a default service MTU of 1514 bytes.

The supported PW service types are 0x0011 (SAToP E1), 0x0012 (SAToP T1), 0x0013 (SAToP E3), 0x0014 (SAToP T3), 0x0015 (CESoPSN basic mode), and 0x0017 (CESoPSN TDM with CAS).

Epipe

Epipe service is designed to carry Ethernet frame payloads, so it can provide connectivity between any two SAPs on different nodes that pass Ethernet frames. The following SAP encapsulations are supported on the 7705 SAR Epipe service:

- Ethernet null
- Ethernet dot1q
- Ethernet qinq

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

The supported PW service types are 0x0004 (Ethernet tagged mode), and 0x0005 (Ethernet raw).

Fpipe

Fpipe service supports frame relay services over an MPLS PSN. MPLS label switched paths—also referred to as MPLS tunnels—are used to forward PW packets between two PEs.

The 7705 SAR supports one-to-one mapping of FR VCs to PWs. An MPLS tunnel can contain several PWs, but each PW encapsulates the traffic of one FR VC.

Fpipes support SAP-to-SAP and SAP-to-spoke SDP binding.

Fpipe service supports the 0x0019 (frame relay DLCI) PW service type.

Hpipe

The 7705 SAR supports many-to-one mapping of HDLC PDUs to PWs, which is also known as port mode encapsulation.The Hpipe provides port-to-port transport of HDLC-encapsulated traffic. The HDLC PDU is transported from PE port to PE port in its entirety, including the HDLC address and control fields, but excluding HDLC flags and the FCS.

Hpipes support SAP-to-SAP and SAP-to-spoke SDP binding.

Hpipe service supports the 0x0006 (HDLC) PW service type.

Ipipe

Ipipe service supports Ethernet null, Ethernet dot1q, Ethernet qinq, IPCP, PPP/MLPPP, FR, and cHDLC SAP encapsulation over IP or MPLS. Ipipes support SAP-to-SAP and SAP-to-spoke SDP binding with a default service MTU of 1500 bytes.

Ipipe service supports 0x000B (IP Layer 2 Transport) PW service type.

ATM PWE3 N-to-1 Cell Mode Encapsulation

The 7705 SAR supports N-to-1 cell mode encapsulation for ATM VPCs and VCCs (per RFC 4717), where N represents the number of VCs or VPs that can be multiplexed onto a single ATM VLL.

For VCCs, N is a configurable value where N can be greater than or equal to 1 ($N \geq 1$). VCC cell mode supports the 0x0009 (ATM N-to-1 VCC Cell Mode) PW service type. The $N > 1$ cell mode encapsulation enables service providers to multiplex multiple ATM VCs onto a single VLL to optimize the use of PWs in the network, to reduce the associated overhead of maintaining the PWs, and to increase the bandwidth available to transport user data. $N > 1$ cell mode encapsulation is implemented on the 7705 SAR using SAP aggregation groups. See SAP Aggregation Groups for more information about how to configure a SAP aggregation group.

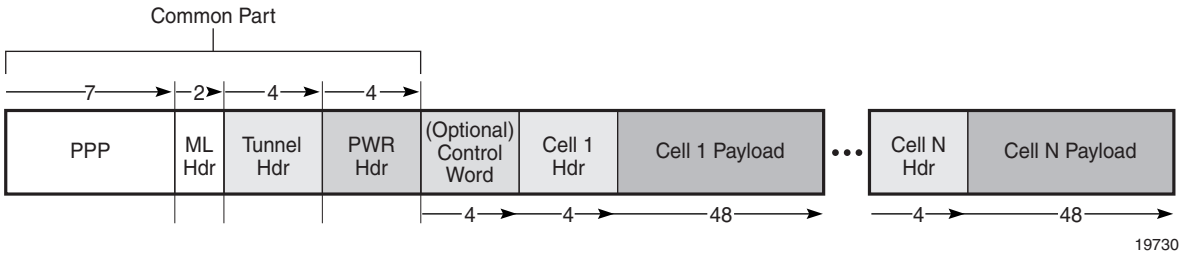
For VPCs, N is not user-configurable and must be equal to 1 ($N = 1$). VPC cell mode supports the 0x000A (ATM N-to-1 VPC Cell Mode) PW service type.

In N-to-1 mode, OAM cells are transported through the VLL in the same way as any other cell.

An optional control word (CW) is supported for ATM VLLs. See Pseudowire Control Word for more information.

Figure 57 shows the structure of an N-to-1 cell mode frame.

Figure 57: N-to-1 Cell Mode Encapsulation



N-to-1 Cell Mode Encapsulation (N = 1)

The following sections describe ATM cell mode encapsulation where N = 1.

Deployment Scenario for N = 1 Cell Mode Encapsulation

In a typical network multiple cell sites are aggregated to individual 7705 SAR nodes. Each cell site has one or more Node Bs. The Node Bs are typically from the same vendor and deployed on a regional basis, and it is common for carriers to simplify provisioning by using the same VPI and VCI values for specific types of traffic handled at different cell sites.

Such a scenario applies, for example, to a service provider that deploys multiple Node Bs in a specific region to support voice, low-speed, and high-speed packet data. In this case, the 7705 SAR grooms the three traffic types from each of the node Bs onto the network. Each traffic type is transported over a dedicated VC as follows:

- circuit-switched VC for voice services
- packet-switched VC for low-bandwidth data services
- HSxPA VC for high-bandwidth data services

In this N-to-1 scenario where N = 1, three VCs and three ATM PWs are required for each Node B.

VPI/VCI Translation (N = 1)

Before traffic from different cell sites can be switched to an RNC, VPI and VCI translation may be required in order to uniquely identify the site and the far-end equipment. If overlapping VPI/VCI values, as described in Deployment Scenario for N = 1 Cell Mode Encapsulation, are not used, VPI/VCI translation is not necessary.

The endpoints of a PWE3 N-to-1 cell mode ATM VLL can be:

- ATM VCs—VPI/VCI translation is supported

  In this case, when the VPI/VCI used at the endpoints (NodeB endpoint and RNC endpoint) are different, the VPI/VCI value can be modified at the endpoint of the far-end PE node, before the cells are switched to the ATM interface.
- ATM VPs—VPI translation is supported (the VPI at each endpoint need not be the same, but the original VCI will be maintained)

  In this case, when the VPI and VCI used at the endpoints are different, only the VPI can be modified at the endpoint (VPI can only be changed by the far-end PE node, before the cells are switched to the ATM interface).

See VPI/VCI Translation for SAP Aggregation Group Members (N>1) for information about how VPI/VCI translation functions with N > 1 mode.

Cell Concatenation (N = 1)

The 7705 SAR supports the concatenation of ATM VP and VC cells into a pseudowire packet payload. Cells are packed on ingress to the VLL and unpacked on egress.

The number of cells in the payload of a single VLL packet is user-configurable, which ensures proper transport of traffic sensitive to delay and jitter. For example, for voice traffic in 3G/WCDMA, delay is a crucial factor and the time spent for concatenation should be minimized. The payload is extremely delay-sensitive and should be transported with only a small amount of bandwidth optimization.

In all cases, the number of cells in a VLL packet must be less than the MTU size, where the MTU maximum is 1514 bytes and the maximum N-to-1 mode payload is 29 cells (52 ATM bytes per cell (no HEC byte)).

While cells are being packed, the concatenation process may be terminated and the packet sent by any one of the following conditions. Each condition has a configurable attribute associated with it:

- reaching a maximum number of cells per packet
- expiring of a timer
- changing of the cell loss priority (CLP) bit

The CLP bits are untouched, even if VPI/VCI translation occurs at egress.



Note: Configuring the attributes that provide the best compromise between minimizing delay (low number of cells concatenated) and maximizing bandwidth (high number of cells concatenated) requires careful planning.

QoS and Traffic Descriptor Profiles (N = 1)

QoS is configured on individual SAPs at ingress and egress using the SAP configuration hierarchy.

Individual SAPs support VC-based characteristics that include ingress and egress ATM (Layer 2) traffic descriptor profiles. Apipe rate-limiting on egress is controlled by the traffic-descriptor profile, not the QoS policy queue rates.

OAM (N = 1)

The ATM PW N-to-1 mode supports OAM operations in non-terminating mode for N = 1 services. The far-end PE node translates the incoming VPI/VCI values of the ATM OAM cells in the same way as the user data cells.

N-to-1 Cell Mode Encapsulation (N > 1)

The following sections describe ATM cell mode encapsulation where N > 1.

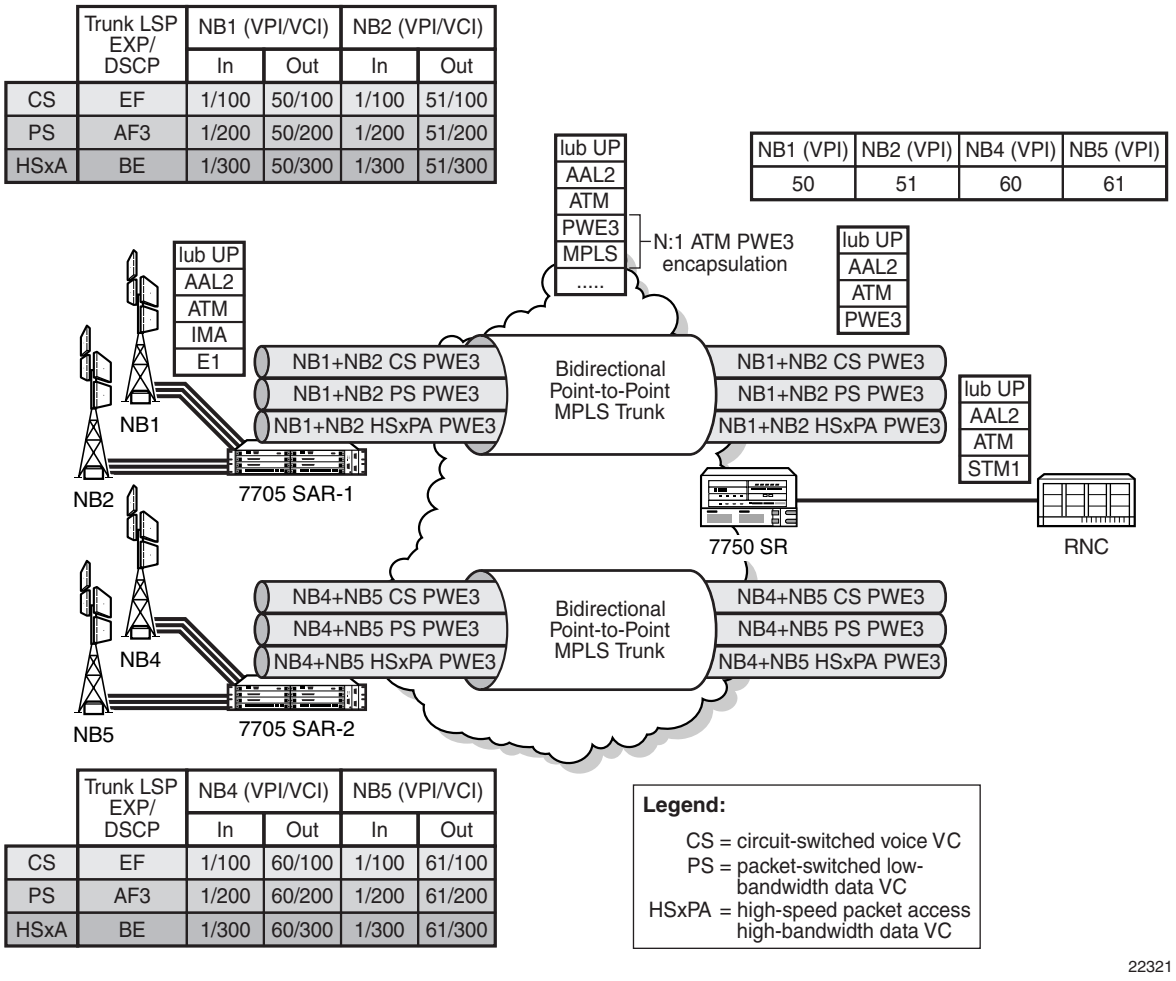
Deployment Scenario for N > 1 Cell Mode Encapsulation

N-to-1 mode where N >1 can be used by wireless service providers to optimize the PWs that are deployed in a network. By multiplexing the VCs of the same type that derive from different ATM base stations (Node Bs), the number of PWs can be decreased to one PW per service type.

With N-to-1 cell mode encapsulation where N > 1 traffic with the same characteristics and QoS requirements (such as delay, jitter, and loss) can be carried over the payload of a single ATM PW.

Figure 58 shows a typical network where N > 1 is deployed.

Figure 58: VCC Cell Mode Encapsulation with N > 1



PW Optimization (N > 1)

A network that uses a single PW for each type of traffic significantly reduces the number of ATM PWs that need to be configured to groom the traffic from multiple Node Bs. The result is a more efficient network backhaul configuration and traffic management strategy.

Overhead Optimization (N > 1)

A VC carrying signaling information between a Node B and an RNC must be configured with minimal packetization delay. N-to-1 cell mode encapsulation where N > 1 can be employed to:

- optimize the overhead associated with ATM PW headers because one PW packet header is used to transport multiple cells
- improve performance for delay and throughput efficiency

Typically, operators configure packetization delay to be 1 ms. Every 1 ms, an ATM PW packet is required to switch the signaling information from the 7705 SAR to the 7750 SR—irrespective of the number of ATM cells that are received to transport signaling information. In most cases, the ATM PW packet carries very few ATM cells.

When signaling VCs from multiple Node Bs are mapped to the same ATM PW in N-to-1 fashion, the number of ATM cells that are received every 1 ms potentially increases. For example, if signaling from five Node Bs is aggregated for each ATM PW, a much larger number of ATM cells can be transported per ATM PW header.

Cell Concatenation (N > 1)

The maximum N-to-1 mode payload is 29 cells for N-to-1 cell mode where N > 1. This value applies to all of the SAPs that are members of the same SAP aggregation group. When the cell concatenation is configured for a specified number of cells, the 7705 SAR counts the cells it has received from all of the SAPs in the same SAP aggregation group and transmits the ATM PW packet when the limit is reached.

For example, for an ATM N >1 VCC service with four SAPs bound to the same SAP aggregation group, where the maximum number of cells for cell concatenation is set to 10, SAP1 might have 2 cells, SAP2 might have 5 cells, SAP3 might have 0 cells, and SAP4 might have 3 cells.

In all cases, the number of cells in a VLL packet must be less than the MTU size, where the MTU maximum is 1514 bytes and the maximum N-to-1 mode payload is 29 cells (52 ATM bytes per cell (no HEC byte)).

The `max-delay` configuration also applies at the SAP aggregation group level.

Note: CLP change-based termination of payload concatenation is not supported for N > 1 cell mode.

OAM (N > 1)

The ATM PW N-to-1 mode supports OAM operations in non-terminating mode for N > 1 services. The 7705 SAR forwards the Layer 2 AIS cells that it receives from the SAP aggregation group SAPs over the PW. The 7705 SAR translates the incoming VPI/VCI values of the AIS cells in the same way as any other type of traffic.

Hardware Support for N > 1 Cell Mode Encapsulation

The following hardware supports N > 1 cell mode encapsulation:

- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- 7705 SAR-F
- 7705 SAR-M (variants with T1/E1 ports)

SAP Support for N > 1

The SAP aggregation group can include SAPs on:

- different ATM ports
- different IMA ports
- a combination of ATM and IMA ports

Note: All SAPs that are bound to the same N > 1 ATM PW must be configured on the same adapter card. An N > 1 ATM PW cannot span multiple adapter cards.

SAP Aggregation Groups

N > 1 cell mode encapsulation is implemented on the 7705 SAR using SAP aggregation groups. The following sections describe how to configure and manage SAP aggregation groups:

- Configuration
- QoS and Traffic Descriptor Profiles (N > 1)
- Statistics and Counters
- Fault Management



Note: An Apipe can have either regular SAPs or SAPs that belong to SAP aggregation groups, but not both.

See Apipe Service with SAP Aggregation Configuration Commands for more information about the commands and parameters that are required to configure a SAP aggregation group.

Configuration

The following section describes the objects required to configure a SAP aggregation group and group members. See VLL Services Configuration Commands for more information.

Sap-aggregation-group Keyword

The `sap-aggregation-group` keyword is used to associate multiple SAPs with a single ATM VCC Apipe service. The SAP aggregation group is a high-level object under which general features are defined. These features include accounting, statistics, and the packet layer QoS profile. All common access parameters are configured under the `sap-aggregation-group`.

Sap-aggregation-group Group Identifier

The `sap-aggregation-group` *`group-id`* is used for two purposes:

- to identify a SAP aggregation group. The *`group-id`* is an alphanumeric identifier. For example:

  ```
  sap-aggregation-group group_id_1
  ```

- to identify a SAP as a member of an aggregation group. SAPs that are to be bound to the same ATM PW payload must be tagged with the same `sap-aggregation-group` *`group-id`*. For example:

  ```
  sap 1/2/3.1:3/300 sap-aggregation-group group_id_1
  ```

The `sap-aggregation-group` *`group-id`* can be up to 32 alphanumeric characters.

VPI/VCI Translation for SAP Aggregation Group Members (N>1)

The `vcid-translation` *vpi/vci* keyword is an optional configuration item that is used only for SAPs that are members of a SAP aggregation group. The `vcid-translation` *vpi/vci* keyword translates the VPI and VCI values of the incoming ATM cells before the cells are mapped to an ATM PW payload. That is, at ingress, the VPI/VCI values for a SAP that is a member of a SAP aggregation group are rewritten by the VPI/VCI values of the `vcid-translation` keyword.

In the reverse direction, when the 7705 SAR receives the cells with translated VPI/VCI values from its peer (such as a 7750 SR), another translation to SAP-configured VPI/VCI values is required before cells are sent to the SAP.

The `vcid-translation` keyword applies to user and OAM cells. When the `vcid-translation` keyword is configured, all cells are translated.

If the `vcid-translation` keyword is not configured for any ATM SAP aggregation group, the ingress VPI/VCI values are retained.

Note: The 7705 SAR performs a check to ensure the uniqueness of the translated VPI/VCI values for all of the SAPs of the same ATM PW service, that is, within the same SAP aggregation group. If there are duplicate VC identifiers, the status of the VCs are set to operationally down and flagged as `ApipeSapVcIdNotUnique`. It is the sole responsibility of the user to ensure uniqueness of VPI/VCI values after translation.

QoS and Traffic Descriptor Profiles (N > 1)

Packet layer, N > 1 ATM PW QoS functions are configured using the SAP aggregation group hierarchy. The QoS profiles for ingress and egress are configurable for each N > 1 service. Any existing QoS profile can be applied to an N > 1 service. The QoS policy determines the QoS offering, including the classification and queuing for the whole PW, irrespective of the number of SAPs that are bound to the aggregated service.

Although a single SAP egress policy is configured for a SAP aggregation group, a separate egress queue is created for each SAP. The MBS and CBS values for each of these egress SAP queues are set to equal the MBS and CBS values configured in the SAP egress QoS policy for the SAP aggregation group. The SAP egress QoS policy causes *n* times the values of the CBS buffers to be committed, where *n* is the number of SAPs in the SAP aggregation group. Refer to the 7705 SAR OS Quality of Service Guide for more information.



Note: ATM layer policing on individual SAPs within a SAP aggregation group is not supported. You cannot apply ATM QoS traffic descriptor profiles on ingress to a SAP in a SAP aggregation group; the profile is set to the default (1).

Rate-limiting on egress for aggregated SAPs is controlled by the ATM traffic-descriptor profile, not the QoS policy queue rates. Each SAP that is a member of a SAP aggregation group has its own egress Layer 2 traffic descriptors. These descriptors are used for shaping and scheduling priority at egress.

Statistics and Counters

The statistics for aggregation groups are maintained on a per-group basis. Statistics for SAP aggregation groups are available using the following commands:

- `show service id` *`n`* `sap-aggregation-group` *`group-id`* `stats`
- `show service id` *`n`* `sap-aggregation-group` *`group-id`* `group-stats`

The `stats` keyword shows statistics for the aggregation group on a per-queue basis. The `group-stats` keyword shows the SAPs that are members of a specified group and the corresponding SAP-level statistics, including aggregated queue statistics. See Show Commands for more information.

Because SAP aggregation groups can span multiple ports, the 7705 SAR does not support port-level packet and discard counters for N > 1 SAPs. When only N > 1 services are configured on an adapter card, the 7705 SAR shows a value of 0 for port packet and discard statistics.

Octet counters are available under each SAP group member. The port and bundle ATM PVC statistics are recorded for the SAPs in a SAP aggregation group. These statistics are helpful for debugging an individual SAP.

A statistics counter tracks all unconfigured or unknown VPI/VCI values that are received in an ATM PW payload from the network. The received VPI/VCI values are compared to the `vcid-translation` *`vpi/vci`* keyword values, and if a value is detected that does not match, the counter for unknown VPIs/VCIs is incremented. The values appear in the Dropped Egress Cells (unconfigured vpi/vci) field and are available using the `show service id` *`n`* `all` and `show service id` *`n`* `sap-aggregation-group` *`group-id`* `detail` commands. See Show Commands for more information.

The `monitor service id` *`n`* `sap-aggregation-group` *`group-id`* provides user-configurable controls for the interval and rate of statistics collection. Refer to the 7705 SAR OS Basic System Configuration Guide for more information.

Statistics are cleared using the `clear service statistics sap` *`sap-id`* and `clear service statistics sap-aggregation-group` *`group-id`* commands. The `sap` *`sap-id`* clears ATM Layer 2 counters; QoS counters are not applicable for SAPs that are members of an aggregation group. The `sap-aggregation-group` *`group-id`* clears network Layer 3 QoS queue statistics. See Clear Commands for more information.

Fault Management

A PW failure results in the transmission of a pw-status TLV message to the far-end node. If the 7705 SAR receives a pw-status TLV message, the 7705 SAR sends an AIS to all of the ATM SAPs.

If an individual SAP in a SAP aggregation group of an N > 1 service fails or is disabled, the 7705 SAR inserts OAM cells into the PW for the failing VC, using the translated VPI/VCIs. If all SAPs in a SAP aggregation group fail or are shut down, the 7705 SAR generates a pw-status TVL message that designates the SAP as sap-down.

In the case of a card failure for N = 1 and N > 1 services, if all the SAPs of the same service fail or if the service is shut down, the 7705 SAR transmits a pw-status TLV message to the far-end node. If the 7705 SAR receives a pw-status TLV from the far-end node (that is, a lacIngressFault or lacEgressFault), AIS messages are generated and sent to all of the SAPs that are part of the SAP aggregation group.

If a port fails, the 7705 SAR sends AIS cells over the ATM PW to the far-end node.

For end-to-end resiliency, an N > 1 ATM PW service supports PW redundancy.

QoS Policies

When applied to 7705 SAR Apipe, Cpipe, Epipe, and Ipipe services, service ingress QoS policies only create the unicast queues defined in the policy.

With Apipe, Cpipe, Epipe, and Ipipe services, egress QoS policies function as with other services where the class-based queues are created as defined in the policy.

Both Layer 2 and Layer 3 criteria can be used in the QoS policies for traffic classification in a Cpipe, Epipe, or Ipipe service. QoS policies on Apipes cannot perform any classification.

IP Filter Policies

The 7705 SAR supports IPv4 and IPv6 filter policies on the following entities:

- IPv4 ingress
  - → network interfaces
  - → Ethernet SAPs on all services
  - → ATM SAPs on a 4-port OC3/STM1 Clear Channel Adapter card in conjunction with bridged llc-snap ATM SAP to VPLS
  - → ATM SAPs on a 16- or 32-port T1/E1 ASAP Adapter card in conjunction with routed VC-Mux ATM SAP to IES
  - → FR and cHDLC SAPs on a 16- or 32-port T1/E1 ASAP Adapter card in conjunction with IP PWs
  - → FR SAPs on a 4-port DS3/E3 Adapter card in DS3, clear channel mode in conjunction with IP PWs
  - → V.35 and X.21 cHDLC SAPs on a 12-port Serial Data Interface card in conjunction with IP PWs
  - → spoke SDPs on VPLS, VPRN interfaces, and IES interfaces
  - → mesh SDPs on VPLS
- IPv4 egress
  - → network interfaces
  - → Ethernet SAPs on VPRN and IES interfaces
  - → Ethernet SAPs on VPLS

- IPv6 ingress and egress
  - → Ethernet network interfaces (with null or dot1q encapsulation)
  - → network interfaces on the 4-port OC3/STM1 Clear Channel Adapter card (with POS encapsulation)
  - → Ethernet SAPs on IES

Configuration and assignment of IP filter policies is similar for network interfaces, IES management SAPs, Ethernet and IP pseudowire SAPs, VPRN and IES SAPs and spoke SDPs, and VPLS SAPs and SDPs (spoke and mesh). Refer to the 7705 SAR OS Router Configuration Guide, "Filter Policies", for information on configuring IP filters.

MTU Settings

There are several MTU values that must be set properly for a VLL service (Apipe, Cpipe, Epipe, Fpipe, Hpipe, or Ipipe) to work from end to end. Figure 59 locates the MTU point for each value. Table 33 describes the MTU points. The MTU points are:

- access port MTU
- SAP MTU
- service MTU
- path MTU
- network port MTU

In order for a VLL service to be declared "up" without any MTU-related error messages, the following rule must be true:

SAP MTU ≥ Service MTU ≤ Path MTU

Figure 59: MTU Points on the 7705 SAR

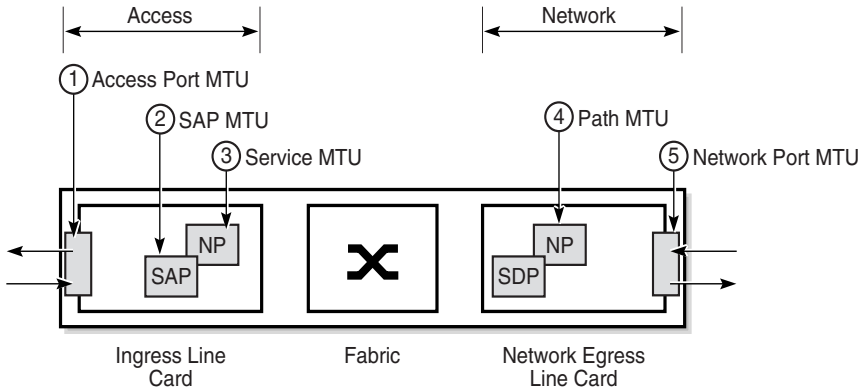


Table 33: MTU Points and Descriptions

	MTU Point	Description	Command
1	Access port MTU	The access port MTU value is a configurable value that accounts for the Layer 2 header and the payload. The default access port MTU value for the following Fast Ethernet port SAP encapsulations is: • **null:** 1514 bytes (payload = 1500 bytes, L2 header = 14 bytes) • **dot1q:** 1518 bytes (payload = 1500 bytes, L2 header = 18 bytes) • **qinq:** 1522 bytes (payload = 1500 bytes, L2 header = 22 bytes)	`mtu`, under the `config>port` context, where the port type can be Ethernet, DSL, GPON, TDM, serial, or SONET/SDH
2	SAP MTU	The SAP MTU value is not a configurable value. It is set at the SAP by the 7705 SAR operating system. It defines the service payload capability of the service and is automatically set to be the same value as the access port MTU.	Not user-configurable
3	Service MTU	The service MTU value is a configurable value and is the same size as the VLL payload. The service MTU is sometimes called the VC-type MTU in the 7705 SAR documentation set. In Figure 59, NP stands for network processor. For CESoPSN with CAS service, ensure that the service MTU is set to a value large enough to account for the extra bytes appended to the packet payload for CAS bits. See Structured T1/E1 CES with CAS for more information.	`service-mtu`, under the `config>port` context, where the port type can be Ethernet, DSL, GPON, TDM, serial, or SONET/SDH
		IP MTU is specific to Layer 3 spoke-SDP termination. Layer 3 spoke-SDP termination can be configured for interfaces under IES and VPRN services. IP MTU is used to signal the service MTU of Layer 3-spoke SDP termination to a peer PE node.	`ip-mtu`, under the `config>service>vprn` or `ies` context for an interface
4	Path MTU	The path MTU is configured at the SDP. It is the maximum that the SDP can transmit without rejecting and discarding the packet. The path MTU value is derived from the network port MTU value by subtracting the Layer 2 and Layer 2.5 overhead values (for MPLS) and the Layer 2 and Layer 3 overhead values (for GRE). If the network port SDP binding is Ethernet, then the following equations hold: • For MPLS: **Path MTU = Port MTU - (Ethernet header [14 bytes or 18 bytes or 22 bytes] + Tunnel header + PW header)** • For GRE: **Path MTU = Port MTU - (Ethernet header [14 bytes or 18 bytes or 22 bytes] + IP header [20 bytes] + Tunnel header [4 bytes] + PW header [4 bytes])**	`path-mtu`, under the `config>service>sdp` context

Table 33: MTU Points and Descriptions (Continued)

	MTU Point	Description	Command
5	Network port MTU	The network port MTU is a configurable value equal to the payload plus all headers (L2, IP (for GRE), tunnel and PW), up to the maximum supported value (hardware limit) of 9728 bytes.	Same as access port MTU (above)

Table 34 displays the default, minimum, and maximum service MTU values for Ethernet ports. These values are dependent upon the port type, mode, encapsulation type, and service type.

Table 35 and Table 36 aid in calculating MTU values for various configurations and operating scenarios.

Table 34: Service MTU Default and Maximum Values

Port Type	Mode	Encap Type	Service Type	Service MTU (bytes)			
				Default	Minimum	Maximum (SAP to SDP)	Maximum (SAP to SAP)
10/100 Ethernet [1]	Access	null, dot1q, qinq [3]	Epipe, VPLS	1514	1	9670 [2]	9724 [2]
			Ipipe	1500			
GigE SFP [1]	Access	null, dot1q, qinq [3]	Epipe, VPLS	1514	1	9670 [2]	9724 [2]
			Ipipe	1500			

Notes:

1. The maximum MTU value is supported only on cards that have buffer chaining enabled; therefore, it is not supported on the 7705 SAR-F and the 8-port Ethernet Adapter card, version 1.
2. On the Packet Microwave Adapter card, the MWA ports support 4 bytes less than the Ethernet ports. Thus, MWA ports support a maximum MTU of 9720 bytes (null) or 9724 bytes (dot1q). MWA ports do not support QinQ.
3. QinQ is supported only on access ports.

For more information on port MTU, refer to “MTU Configuration Guidelines” in the 7705 SAR OS Interface Configuration Guide.

Table 35: MTU Calculator—Service Creation (Worst Case) Access Ports and SAPs

Service Creation									
		Access Port Default MTU		SAP					
		TDM/ ATM	Eth	Epipe	Ipipe	Apipe	Cpipe	Fpipe	Hpipe
	Max Payload			9732	9732	1514	1514	2048	2048
	RTP Header						12		
	Control Word			4	4	4	4	4	4
SDP Encap: GRE/MPLS	IP Header for			20	20	20	20	20	20
	GRE/MPLS			4	4	4	4	4	4
	PW Header			4	4	4	4	4	4
	VCCV Type 2			4	4	4	4	4	4
	Fast Reroute Label								
	LDP over RSVP								
Physical Media(T1/E1 ASAP and Ethernet Adapter cards)	Eth Null		1514						
	Eth Dot1q		1518						
	Eth QinQ		1522						
	Eth Type								
	Eth-SA								
	Eth-DA								
	TDM/ATM	1572	1572						
	PPP Protocol								
	ML Sequence								
	ML Preamble								
	Total			**9768**	**9768**	**1550**	**1562**	**2084**	**2084**

Table 36: MTU Calculator—Service Creation (Worst Case) Network Ports

Service Creation									NW
		Network Port Default MTU							
							Epipe over MPLS Encap	Epipe /Ipipe over GRE	LSR
		PPP	ML/ PPP	Eth Null	Eth Dot1q	MPLS Label	Best Case	Worst Case	Worst Case
	Max Payload						40	2048	2084
	RTP Header								
	Control Word							4	
SDP Encap: GRE/MPLS	IP Header for							20	
	GRE/MPLS						4	4	4
	PW Header						4	4	
	VCCV Type 2					4		4	
	Fast Reroute Label					4			
	LDP over RSVP					4			
Physical Media(T1/E1 ASAP and Ethernet Adapter cards)	Eth Null								
	Eth Dot1q				4			4	4
	Eth QinQ				8			8	8
	Eth Type			2	2			2	2
	Eth-SA			6	6			6	6
	Eth-DA			6	6			6	6
	TDM/ATM								
	PPP Protocol	2	2				2		
	ML Sequence		3						
	ML Preamble		1						
	Total	**2**	**6**	**14**	**26**		**50**	**2110**	**2114**

Note:

- In order to accommodate current and future services (including overhead), the MTU values for Gigabit Ethernet and PPP/MLPPP ports have the default value set to 1572 bytes. For 10/100 Ethernet ports, the MTU value is set to 1514 or 1518 or 1522 bytes, depending on the encapsulation setting (null or dot1q or qinq).
- The default service MTU value is 1514 bytes; the maximum value is 1522 bytes.

Targeted LDP and MTU

The extended discovery mechanism for Label Distribution Protocol (LDP) sends LDP Targeted Hello messages to a specific address. This is known as targeted LDP or T-LDP. Refer to RFC 5036 for detailed information about the extended discovery mechanism.

During the VLL service creation process (that is, using targeted LDP signaling), the MTU or payload size of a service is signaled to the far-end peer. MTU settings at both ends (near and far peers) must match in order for the VLL service to operate. Table 37 shows the values that are expected to match.

Table 37: Matching MTU or Payload Values for Signaled VLL Services

	Apipe	Cpipe	Epipe	Fpipe	Hpipe	Ipipe
Payload size (bytes)		Yes				
Bit rate		Yes				
Maximum number of ATM cells	Yes					
Service MTU			Yes	Yes	Yes	Yes
Must match at both ends	Yes	Yes	Yes	Yes	Yes	Yes

QinQ (VLL Service)

Epipe and Ipipe VLL services support QinQ functionality. For details, see QinQ Support.

Pseudowire Control Word

The PW control word (CW) is a 32-bit field that is inserted between the VC label and the Layer 2 frame. The presence of the control word is indicated by the C bit of the FEC element used in LDP signaling. The PW control word is described in RFC 4385.

The PW control word is supported for all implemented PW types (ATM N-to-1 cell mode, Ethernet VLLs, SAToP, CESoPSN, FR PW, HDLC PW, and IP PW). The control word must be used for any pseudowire service that transports packets that are less than 64 bytes; see the control-word command description for more information.

The following points describe the behavior of the 7705 SAR when it receives a Label Mapping message for a PW. It is assumed that no Label Mapping message for the PW has been sent to the next PW router yet. The 7705 SAR operating system does the following.

- If the received Label Mapping message has C = 0 (where C refers to the C bit of the FEC element), a Label Mapping message with C = 0 is sent forward to the next router (or hop). In this case, the control word is not used.
- If the received Label Mapping message has C = 1 and the PW is locally configured such that the use of the control word is mandatory, then the 7705 SAR sends a Label Mapping message with C = 1. In this case, the control word is used. (Note: SAToP and CESoPSN are the only services that require the control word.)
- If the received Label Mapping message has C = 1 and the locally configured PW does not support use of an optional control word (that is, Ethernet or ATM N-to-1 cell mode PWs), then the 7705 SAR sends a new Label Mapping message in which the C bit is set to correspond to the locally configured preference for use of the control word (that is, C = 0).

Pseudowire Redundancy

Pseudowire (PW) redundancy protects a PW and any services on the PW against endpoint failures. This differs from LSP redundancy and FRR, which offer protection against link and node failures within the backhaul network.

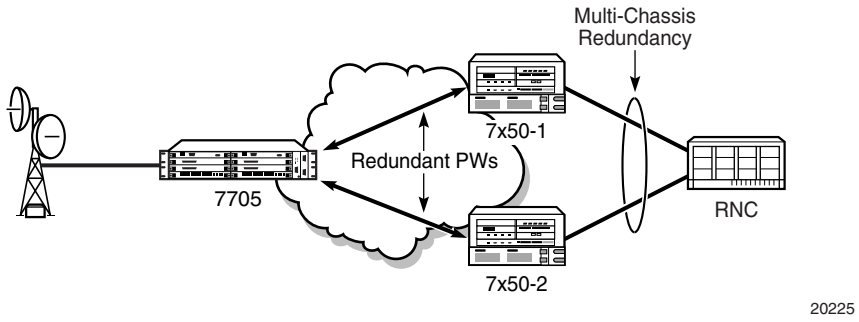
As shown in Figure 60, in order to provide redundant PWs, the 7705 SAR must signal PWs to two endpoints at the MTSO (7x50-1 and 7x50-2), which is done using two spoke SDPs on the 7705 SAR. This configuration removes any single point of failure from a given network. If 7x50-1 loses all of its connectivity to the network or to the RNC, then the 7705 SAR can reroute the PW traffic to 7x50-2, which switches traffic to the RNC.

For end-to-end protection, PW redundancy is supported in the following scenarios:

- PW redundancy with multi-chassis LAG (MC-LAG)/multi-chassis APS (MC-APS)
- PW redundancy with VPLS or PW switching

For more information on MC-APS, refer to the 7705 SAR OS Interface Configuration Guide, "Automatic Protection Switching". For more information on MC-LAG, refer to the SAR OS Interface Configuration Guide, "Multi-Chassis LAG".

Figure 60: Pseudowire Redundancy



PW redundancy applies to all VLL services available on the 7705 SAR: Apipe, Cpipe, Epipe, Fpipe, Hpipe, and Ipipe.

PW Redundancy Operation

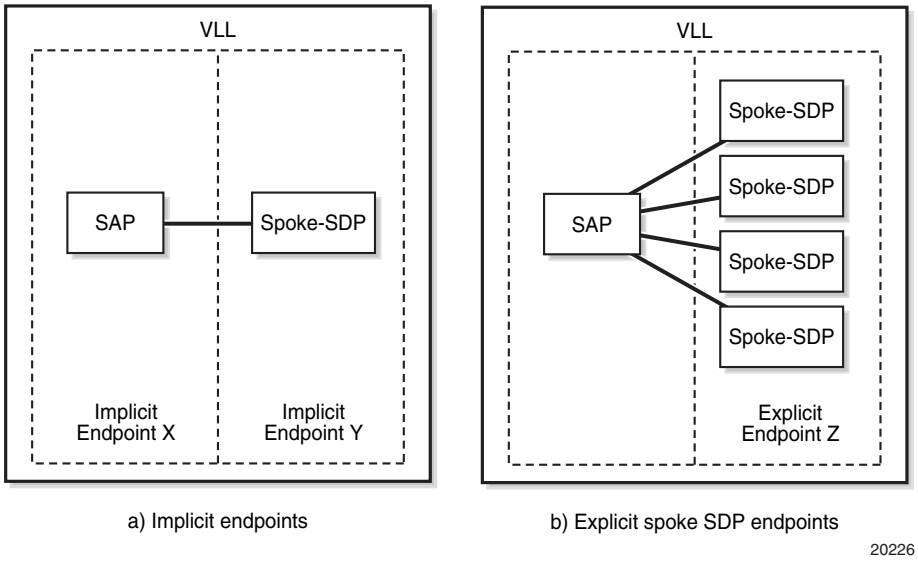
PW redundancy on the 7705 SAR is similar to a point-to-multipoint implementation for PWs (in the ingress to the egress direction). A single SAP can be bound to more than one spoke SDP, and traffic from multiple spoke SDPs can all be switched to the same SAP. To implement PW redundancy, a PW service on the 7705 SAR must be able to accommodate more than one spoke SDP on the spoke SDP side. This is achieved using the concept of endpoints.

An endpoint can be thought of as a container for a single SAP, a single spoke SDP, or multiple spoke SDPs. Figure 61 illustrates the model for a redundant VLL service based on the endpoints. Endpoints are implicit or explicit objects.

Implicit endpoints are transparent to the user and are not user-configurable. As shown in Figure 61a, implicit endpoints mean that one endpoint is a SAP and another endpoint is a spoke SDP. Endpoints are considered implicit if the `endpoint` command is not used in the `config>service>xpipe>spoke-sdp` context, where `xpipe` refers to any of the VLL services.

Explicit endpoints are user-configurable and apply when there are multiple spoke SDPs. As shown in Figure 61b, explicit endpoints mean that there can be multiple spoke SDPs associated with the endpoint. An endpoint created explicitly can have up to four spoke SDPs associated with it. The explicit endpoint method is used for PW redundancy. Explicit endpoints are user-configurable.

Figure 61: Implicit and Explicit Endpoint Objects



The 7705 SAR supports the following types of endpoint objects:

- SAP—there can be only one SAP per PW endpoint (Endpoint X in Figure 61a)
- Spoke SDP—from the perspective of a 7705 SAR, if there is only one SDP endpoint, then it is a spoke SDP endpoint and it is implicitly defined. In other words, there can be only one implicitly defined spoke SDP per PW endpoint (Endpoint Y in Figure 61a).
- Primary spoke SDP—there can be only one explicitly defined primary spoke SDP per PW endpoint (one of the spoke SDPs at Endpoint Z in Figure 61b). If a primary spoke SDP is defined, then there can be up to three secondary spoke SDPs per endpoint and the service can be revertive.

- Secondary spoke SDP—there can be up to four explicitly defined secondary spoke SDPs per endpoint if no primary spoke SDP is defined; otherwise, there can be up to three. Secondary spoke SDPs are assigned a precedence value that is used by the 7705 SAR OS to determine which secondary PW becomes active when the currently active PW fails (see Selecting the Active Spoke SDP for PW Redundancy Configuration).

Multiple spoke SDPs can be established between a 7705 SAR and any SR platform. For example, multiple spoke SDPs on a 7705 SAR can connect to a 7750 SR. In this case, the 7750 SR must be configured to use multi-chassis backup in conjunction with multi-segment PWs; that is, the 7750 SR nodes at the far end must support multi-chassis redundancy.

A PW service endpoint can only use a single active spoke SDP for transmission at any given time. A PW SAP can receive traffic from any of the endpoint spoke SDPs assigned to the service.

7705 SAR nodes support user-initiated manual switchover of the VLL path to the primary path or any of the secondary paths using the `force-switchover` command under the `tools>perform>service-id` context. A manual switchover is useful during planned outages such as node upgrade procedures.

Selecting the Active Spoke SDP for PW Redundancy Configuration

There are two main scenarios for configuring PW redundancy. One scenario uses a primary spoke SDP and provides revertive behavior. The other scenario uses only secondary spoke SDPs for non-revertive behavior.

Note: Non-revertive behavior is not supported on Cpipes.

Primary and Secondary Spoke SDPs

If a primary spoke SDP is defined, up to three secondary spoke SDPs can also be defined. The VLL service always uses the primary endpoint PW and only switches to a secondary PW when the primary PW is down. The PW service switches the path back to the primary PW when the primary PW is back up. The user can configure a timer to delay reverting back to the primary path or to never revert back. When the primary PW goes down, the 7705 SAR OS selects the secondary spoke SDP that is operationally up and has the highest precedence setting.

Secondary Spoke SDPs Only

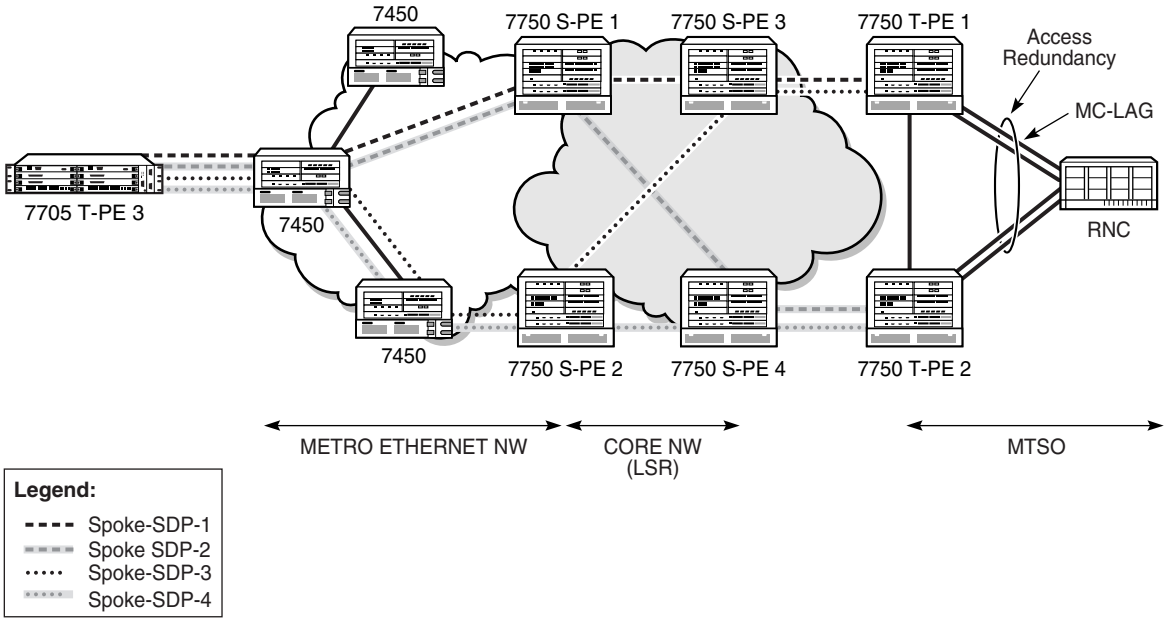
If a primary spoke SDP is not defined, up to four secondary spoke SDPs can be defined. The user can configure the precedence of each secondary PW to indicate the order in which secondary PWs are activated. The secondary PW with the highest precedence is selected first. If two or more secondary spoke SDPs are assigned the same precedence, the 7705 SAR OS selects the secondary path that is operationally up and has the lowest spoke SDP identifier. There is no revertive behavior between secondary paths, which means that a secondary path will not switch to another secondary path of higher precedence if one becomes available.

The use of four secondary spoke SDPs is illustrated in Figure 62, where:

- spoke SDP-1 goes over S-PE-1 to T-PE1 (red path) (S-PE is a switching PE and T-PE is a terminating PE)
- spoke SDP-2 goes over S-PE-1 to T-PE2 (green path)
- spoke SDP-3 goes over S-PE-2 to T-PE1 (violet path)
- spoke SDP-4 goes over S-PE-2 to T-PE2 (orange path)

.

Figure 62: Pseudowire Redundancy with Four Spoke SDPs



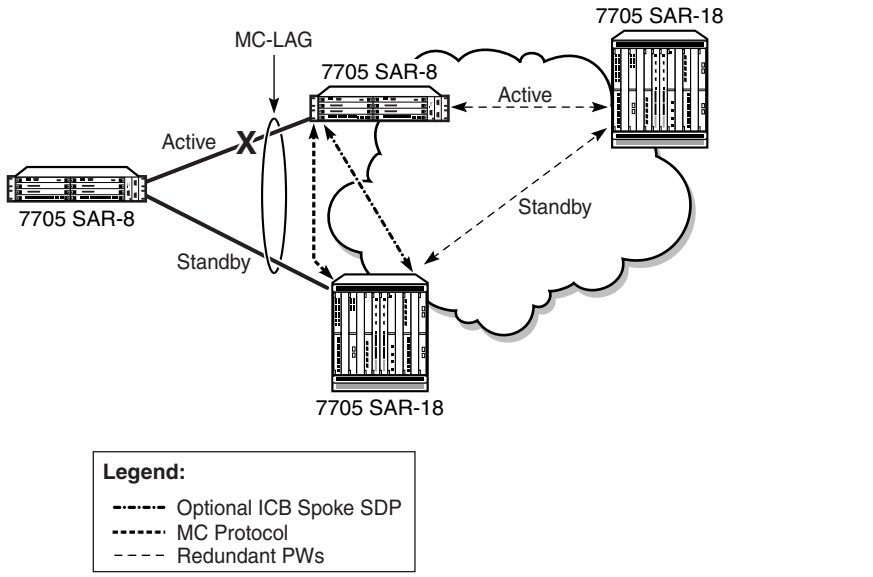
PW Redundancy and Inter-Chassis Backup

ICB (Inter-Chassis Backup) spoke SDPs are supported for use with Cpipe services in an MC-APS configuration and with Epipe services in an MC-LAG configuration. ICB improves switch times, provides additional protection in case of network failures, and reduces packet loss when an active endpoint is switched from a failed MC-APS or MC-LAG node to a protection node.

An MC-APS or MC-LAG switchover consists of two switchovers, first on the access side and then on the network side. A failure on the access side (an MC-APS or MC-LAG switchover), triggers a network-side (pseudowire) switchover. However, a pseudowire switchover does not trigger an MC-APS or MC-LAG switchover.

Figure 63 shows a network experiencing an access side failure.

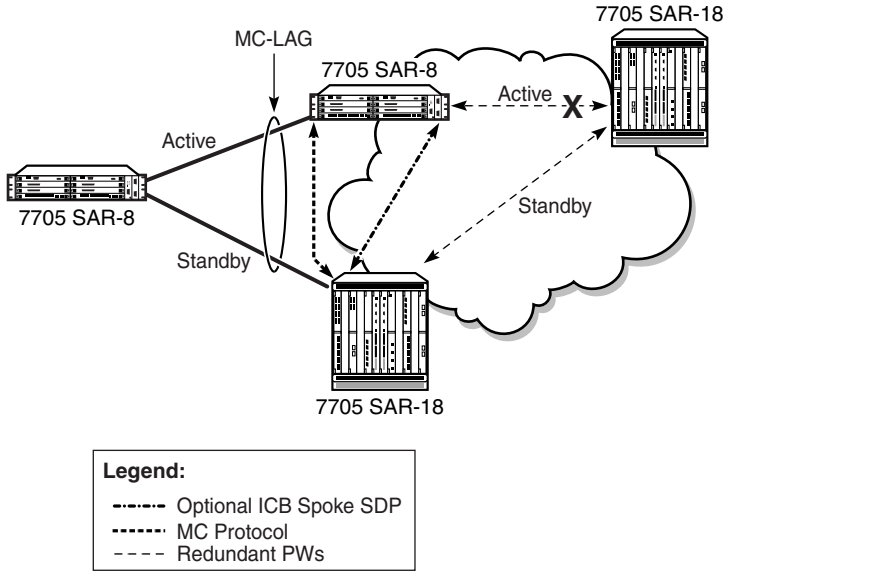
Figure 63: Access Side Failure with ICB Protection



If the active access link of an MC-APS or MC-LAG group fails, the MC-APS or MC-LAG protocol causes the access side to switch to the protection (standby) link. If ICB is configured, the in-flight packets coming from the network side will be sent over the ICB to the newly active access link. Shortly after the access-side switchover, pseudowire redundancy causes the network side to switch as well. In this scenario, ICB reduces the switch time by carrying the in-flight packets during the access and network switchovers.

Figure 64 shows a network experiencing a network side failure.

Figure 64: Network Side Failure with ICB Protection



If the active network link of an MC-APS or MC-LAG group fails, pseudowire redundancy causes the network side to switch to the standby pseudowire. This does not trigger an MC-APS or MC-LAG switchover. If there is no ICB, then the end-to-end transmission will be lost as the access link never switches. If ICB is configured, the packets coming in from the newly active network link are sent across the ICB to the active access side and vice versa. In this scenario, ICB provides protection against network failures.

An endpoint can have only one ICB spoke SDP, which must be identified as ICB in the `spoke-sdp` command. If an ICB spoke SDP is added to an endpoint, then a SAP can be added only if it is part of an MC-LAG or MC-APS group. Similarly, if an MC-LAG or MC-APS SAP is added to an endpoint, the only other possible addition to that endpoint is an ICB spoke SDP.

AIS Fault Propagation

When the 7705 SAR is operating in an MPLS network with TDM pseudowires using PW redundancy, it can interoperate with SDH networks that use subnetwork connection protection (SNCP). SNCP is a path-based protection mechanism for T1/E1 services. When the 7705 SAR interoperates with SDH networks that use SNCP, it can make PW redundancy switching decisions based on SDH signaling, which keeps the active data paths in the SDH and the MPLS networks synchronized.

This functionality is only available on unframed E1 channels and unframed DS1 channels on the following cards:

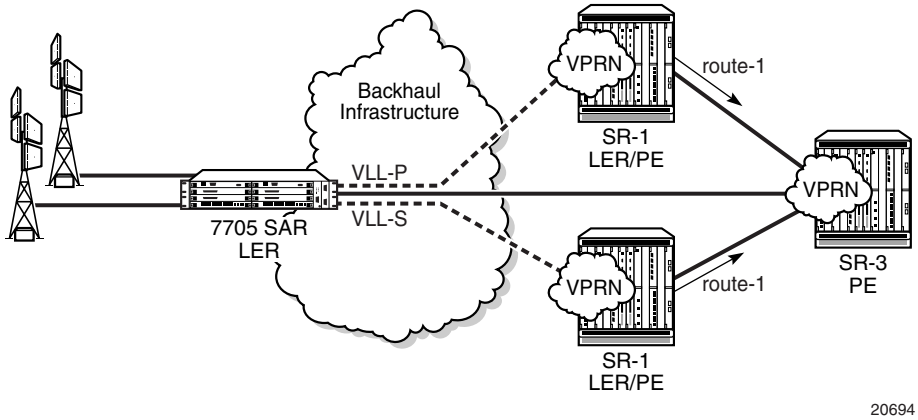
- 2-port OC3/STM1 Channelized Adapter card
- 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card

If SNCP-protected equipment detects a failure in an SDH network, it inserts an AIS into the TU-12 overhead of the SDH frame so that the appropriate activity switch can occur in the SDH network. When the 7705 SAR detects a TU12-AIS for a specific VC-12 path in an SDH network and AIS propagation is enabled using the command `config>card>mda>ais-propagation`, the 7705 SAR generates a TU12-AIS for the corresponding VC-12 at the other end of the Cpipe. If the VC-12 path is involved in PW redundancy, a PW activity switch occurs, which signals the SDH node to do an SNCP switch.

Active/Standby Mode for Pseudowire Redundancy (Standby Signaling)

Pseudowire redundancy as described in the previous section operates in active/active mode; that is, the primary pseudowire is up and ready to transmit and receive traffic, and the secondary pseudowire is up and ready to receive traffic. In Figure 65, if both pseudowires were active, this mode of operation would offer seamless redundancy in most cases. However, this mode could also potentially stress the IGP; for example, in active/active mode the number of routes advertised is greater than in active/standby mode. Another example is the duplication of Ethernet control frames to the 7705 SAR from a VPLS or VPRN service on an SR node through both primary and secondary VLLs.

Figure 65: Active/Standby Mode for Redundant Pseudowires



Active/standby mode is introduced to address these issues. Active/standby mode is also referred to as standby signaling. Standby signaling is supported on all VLLs: Apipes, Cpipes, Epipes, Ipipes, Fpipes, and Hpipes.

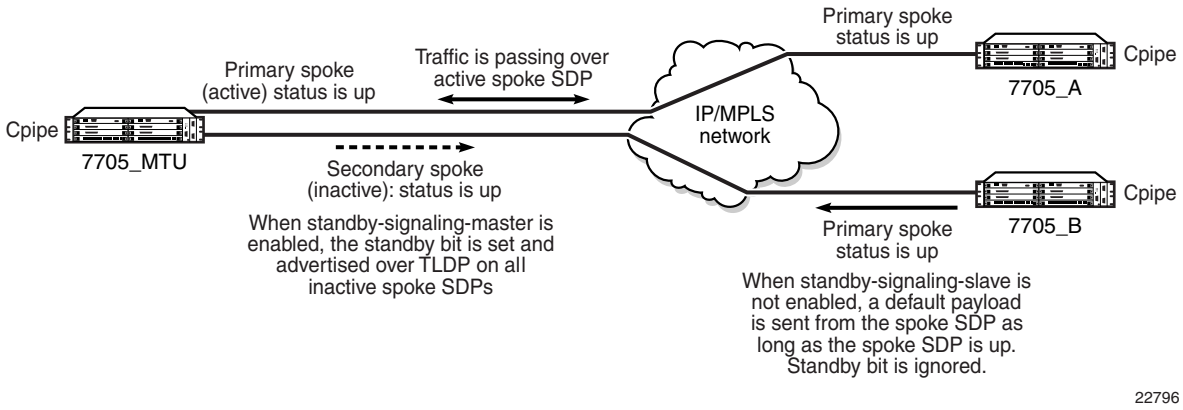
Standby signaling has two components: standby-signaling master and standby-signaling slave. Figure 66 shows an example where `standby-signaling-master` has been enabled on the 7705_MTU node. A Cpipe has been configured between the 7705_MTU and the 7705_A node and between the 7705_MTU and the 7705_B node. The spoke SDP towards 7705_A is the active PW (precedence primary), while the spoke SDP towards 7705_B is the standby PW (precedence 1).

An example of the CLI syntax to configure the 7705_MTU is:

CLI Syntax:

```
config>service# cpipe 555 customer 1 vc-type cesopsn
 create
     endpoint "stdbyMaster" create
          standby-signaling-master
     exit
     spoke-sdp 5:555 endpoint "stdbyMaster" create
          precedence 1
     exit
     spoke-sdp 131:555 endpoint "stdbyMaster" create
          precedence primary
     exit
```

Figure 66: Standby-signaling-master Enabled



Note: Standby-signaling-master should be enabled on only one endpoint of the pseudowire; otherwise, the pseudowire could bounce.

Traffic passes over the active PW. If the status of the active PW changes, the standby PW becomes active and starts passing traffic.

Because the 7705_MTU is configured for `standby-signaling-master`, the standby PW sends its status to the far end by sending a standby bit over T-LDP. However, the receiving end (in Figure 66, 7705_B) ignores the bit and continues to transmit data towards the 7705_MTU as long as the PW is up; therefore, pseudowire redundancy does not work.

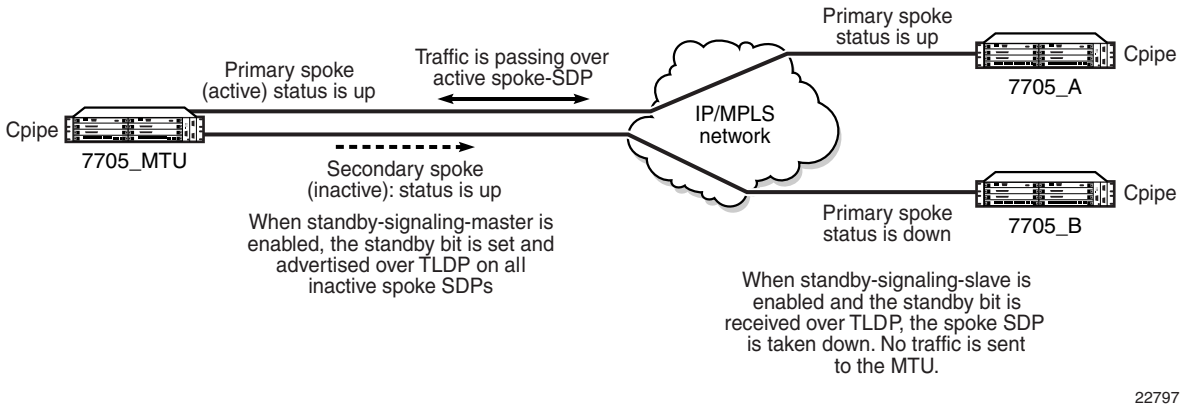
In order to stop data from being transmitted along standby spoke SDPs, the far-end endpoints must be enabled for `standby-signaling-slave`. Figure 67 shows a scenario where 7705_A and 7705_B have been enabled for `standby-signaling-slave`.

An example of the CLI syntax to configure the 7705_A and 7705_B nodes is:

CLI Syntax:
```
config>service# cpipe 555 customer 1 vc-type cesopsn
 create
     endpoint "FarendA" create
          standby-signaling-slave
     exit
     spoke-sdp 130:555 endpoint "FarendA" create
          precedence primary
     exit
```

CLI Syntax:
```
config>service# cpipe 555 customer 1 vc-type cesopsn
 create
     endpoint "FarendB" create
          standby-signaling-slave
     exit
     spoke-sdp 130:555 endpoint "FarendB" create
          precedence primary
     exit
```

Figure 67: Standby-signaling-slave Enabled



With `standby-signaling-slave` enabled, when the standby PW sends the standby bit to the far end over T-LDP, the SDP flags bit is set to block traffic being transmitted back to the 7705_MTU (flags bit is set to StandbySigSlaveTxDown). The spoke SDP remains up. If the standby PW becomes active, the flags bit is reset and traffic resumes on the PW.

Note: Standby-signaling-master and standby-signaling slave cannot be enabled at the same time on the same endpoint.

PW Status Signaling Label Withdrawal Option

The 7705 SAR supports PW status signaling or label withdrawal for signaling PW status.

Signaling PW status based on label withdrawal requires the PW label to be released, whereas PW status signaling can mark the PW as unusable based on local-end and far-end status and on status messages exchanged between endpoints.

As indicated in RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol*, PW status signaling is the preferred method for exchanging state information between two endpoints and should be used as long as both endpoints support it.

However, the PW label withdrawal method for exchanging PW status can be configured even if the far end supports PW status signaling and the PW status TLV for designating the operational (forwarding) state of a PW.

This configuration option allows the PEs at both ends of the spoke SDP to use PW label withdrawal rather than the PW status TLV. This is necessary when a 7705 SAR must interoperate with PEs that do not support the PW forwarding standby bit, or when multiple remote PEs are connected to an r-VPLS instance via spoke-SDP termination.

Interoperation with PEs that do not Support the PW Forwarding Standby Bit

Not all PEs support the PW forwarding standby bit as part of pseudowire status signaling. If `standby-signaling-master` is enabled on the 7705 SAR, then it signals standby on all but one of the PWs, and blocks the transmit direction for all standby PWs. However, if a PE does not support processing of that bit, it will not block its end of the PW and will forward traffic onto that PW. As a result, traffic received on a standby PW at the 7705 SAR end is forwarded to the associated SAP. Master-slave PW redundancy prevents this from occurring only if all of the PEs support PW active/standby mode in a master-slave configuration and block their transmit directions for the standby PW. Previously, the 7705 SAR only used the PW label withdrawal method if a PE did not support the PW status signaling bit.

In order to allow PW redundancy on the 7705 SAR to interoperate with third party PEs that do not support the PW forwarding standby bit, an option to disable PW status signaling is provided. This will ensure that forwarding on a standby PW will be bidirectionally blocked in order to prevent PWs in standby mode from transmitting traffic via label withdrawal.

Multiple Remote PEs Connected to an r-VPLS Instance Via Spoke-SDP Termination

When multiple remote PEs are hooked up to an r-VPLS instance via spoke-SDP termination, label withdrawal can be used to control the operational status of the associated IP interface at the r-VPLS.

If all of the spoke SDPs of an r-VPLS enter the forwarding standby state, the spoke SDP is locally blocked if the `no ignore-standby-signaling` option is configured. In its default mode of operation, the 7705 SAR automatically negotiates the use of PW status signaling on the spoke SDP. This means that a standby spoke SDP will not go operationally down at the endpoint. As a result, there is no change in the operational status of the VPLS to which the spoke SDPs are bound, and consequently, no change to the operational status of the r-VPLS or the IP interface.

The ability to disable PW status signaling makes it possible to configure the use of PW label withdrawal on a node, rather than allowing the automatic negotiation of PW status signaling. When pseudowire status signaling is disabled, a 7705 SAR does not include the PW status TLV in the initial label mapping message of the pseudowire that is used for a spoke SDP. This forces both 7705 SAR nodes to use the pseudowire label withdrawal method for exchanging pseudowire status. When the remote endpoint determines that a particular PW should be in standby mode, it will withdraw the PW label. This causes VPLS to go operationally down if the label is withdrawn for all PWs on a VPLS. The IP interface associated with r-VPLS goes down as a result.

Configuring a VLL Service with CLI

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

Topics in this section include:

- Common Configuration Tasks
- Configuring VLL Components
- Service Management Tasks

Common Configuration Tasks

The following list provides a brief overview of the tasks that must be performed to configure a VLL service.

- Associate the service with a customer ID.
- Define SAP parameters.
  - → Optional—select egress and ingress QoS policies (configured in `config>qos` context)
  - → Optional—select ingress IP filter policies (for Epipes and Ipipes only)
- Define spoke SDP parameters.
  - → Optional—select egress and ingress vc label parameters
  - → Optional—explicitly assign spoke SDP endpoints for pseudowire (PW) redundancy applications
- Enable the service.

Configuring VLL Components

This section provides configuration examples for components of VLL services. Each component includes some or all of the following: introductory information, CLI syntax, a specific CLI example, and a sample CLI display output. Included are the following VLL components:

- Creating an Apipe Service
- Creating a Cpipe Service
- Creating an Epipe Service
- Creating an Fpipe Service
- Creating an Hpipe Service
- Creating an Ipipe Service
- Configuring PW Switching
- Configuring Ingress and Egress SAP Parameters
- Using the Control Word
- Configuring PW Redundancy

Creating an Apipe Service

This section describes how to create an Apipe service. Additional topics in this section include:

- Configuring Apipe SAP Parameters
- Configuring Apipe SDP Bindings
- Configuring Apipe SAP Aggregation Groups
- Configuring Apipe SAPs as Aggregation Group Members

Use the following CLI syntax to create an Apipe service.

CLI Syntax:
```
config>service# apipe service-id [customer customer-id]
  [create] [vpn vpn-id] [vc-type {atm-vcc|atm-vpc|
  atm-cell}] [vc-switching]
      description description-string
      service-mtu octets
      no shutdown
```

PE router 1 (A:ALU-41):

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

PE router 2 (A:ALU-42):

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

The following example shows the Apipe service creation output.

PE Router 1 (ALU-41):

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

PE Router 2 (ALU-42):

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

Configuring Apipe SAP Parameters

Use the following CLI syntax to configure Apipe SAP parameters. For ingress and egress configuration information, see Configuring Ingress and Egress SAP Parameters.

CLI Syntax:

```
config>service# apipe service-id [customer customer-id]
  [create] [vpn vpn-id] [vc-type {atm-vcc|atm-vpc|
  atm-cell}] [vc-switching]
      sap sap-id [create]
           accounting-policy acct-policy-id
           atm
                egress
                     traffic-desc traffic-desc-profile-id
                ingress
                     traffic-desc traffic-desc-profile-id
                oam
                     alarm-cells
           collect-stats
           description description-string
           egress
                qos policy-id
           ingress
                qos policy-id
           no shutdown
```

Example:

```
A:ALU-41>config>service# apipe 5
A:ALU-41>config>service>apipe# sap 1/1/1.1:0/32 create
A:ALU-41>config>service>apipe>sap# ingress
A:ALU-41>config>service>apipe>sap>ingress# qos 102
A:ALU-41>config>service>apipe>sap>ingress# exit
A:ALU-41>config>service>apipe>sap# egress
A:ALU-41>config>service>apipe>sap>egress# qos 103
A:ALU-41>config>service>apipe>sap>egress# exit
A:ALU-41>config>service>apipe>sap# no shutdown
A:ALU-41>config>service>apipe>sap# exit
```

The following example shows Apipe SAP configuration output for PE Router 1 (ALU-41).

```
A:ALU-41>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 1/1/1.1:0/32 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            no shutdown
        exit
------------------------------------
```

To configure a basic local Apipe service (SAP-to-SAP), enter the sap *sap-id* command twice with different port IDs in the same service configuration.

The following example shows an ATM SAP-to-SAP configuration:

```
A:ALU-4>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "ATM sap2sap"
            service-mtu 1514
            sap 1/1/1.1:0/32
            sap 1/2/1.1:0/100
            no shutdown
        exit
------------------------------------
```

The following example shows an Apipe SAP configuration for a virtual trunk service:

```
A:ALU-4>config>service# info
------------------------------------
...
        apipe 5 customer 1 vc-type atm-cell create
            description "port VT apipe"
            service-mtu 1500
            sap 1/1/10 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
           exit
            no shutdown
        exit
------------------------------------
```

Configuring Apipe SDP Bindings

Use the following CLI syntax to create a spoke SDP binding with an Apipe service (for distributed service). For SDP configuration information, see Configuring SDPs.

CLI Syntax:

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

Example:

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

The following example shows the Apipe spoke SDP configuration output for PE Router 1 (ALU-41).

```
A:ALU-41>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 1/1/1.1:0/32 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:5 create
            exit
            no shutdown
        exit
...
------------------------------------
A:ALU-41>config>service#
```

Configuring Apipe SAP Aggregation Groups

Use the following CLI syntax to configure the parameters for an Apipe SAP aggregation group. The `vc-type` parameter must be set to `atm-vcc`. For ingress and egress configuration information, see Configuring Ingress and Egress SAP Parameters.

CLI Syntax:

```
config>service# apipe service-id [customer customer-id]
  [create] [vpn vpn-id] [vc-type {atm-vcc|atm-vpc|
  atm-cell}] [vc-switching]
      sap-aggregation-group group-id [create]
           accounting-policy acct-policy-id
           collect-stats
           description description-string
           egress
                qos policy-id
           ingress
                qos policy-id
           no shutdown
      no shutdown
```

Example:

```
A-ALU-1>config>service# apipe 3 customer 1 vc-type
  atm-vcc create
A-ALU-1>config>service>apipe# sap-aggregation-group
  GroupName1 create
A-ALU-1>config>service>apipe>sap-aggregation-group#
  ingress
A-ALU-1>config>service>apipe>sap-aggregation-
  group>ingress# qos 102
A-ALU-1>config>service>apipe>sap-aggregation-
  group>ingress# exit
A-ALU-1>config>service>apipe>sap-aggregation-group#
  egress
A-ALU-1>config>service>apipe>sap-aggregation-group#
  egress qos 103
A-ALU-1>config>service>apipe>sap-aggregation-
  group>egress# exit
A-ALU-1>config>service>apipe>sap-aggregation-group# no
  shutdown
A-ALU-1>config>service>apipe>sap-aggregation-group#
  exit
A-ALU-1>config>service>apipe#
```

The following example shows an Apipe SAP aggregation group configuration output.

```
A-ALU-1>config>service# info
------------------------------------
...
        apipe 3 customer 1 vc-type atm-vcc create
            description "SAP aggregation group 1"
            sap-aggregation-group GroupName1 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            no shutdown
        exit
...
------------------------------------
```

Configuring Apipe SAPs as Aggregation Group Members

Use the following CLI syntax to configure an Apipe SAP as a member of a SAP aggregation group. To add a SAP as a member of a SAP aggregation group, enter the sap-aggregation-group *group-id* command after the sap *sap-id* command. A SAP aggregation group can have up to 16 SAP members.

ATM attributes for aggregation group members can also be configured, including:

- vcid-translation
- traffic descriptor profiles for egress

Note: You cannot apply ATM QoS traffic descriptor profiles on ingress to a SAP in a SAP aggregation group; the profile is set to the default (1). Attempting to change the ingress traffic descriptor will cause an error message to be displayed.

- OAM alarm notifications

For ingress and egress configuration information, see Configuring Apipe SAP Parameters.

CLI Syntax:

```
config>service# apipe service-id [customer customer-id]
  [create] [vpn vpn-id] [vc-type {atm-vcc|atm-vpc|
  atm-cell}] [vc-switching]
      sap sap-id sap-aggregation-group group-id
           atm
                vcid-translation vpi/vci
                     egress
```

```
                                    traffic-desc traffic-desc-
                                      profile-id
                                ingress
                                    traffic-desc traffic-desc-
                                      profile-id
                                oam
                                    alarm-cells
                            description description-string
                            no shutdown
                        no shutdown
```

Example:

```
A-ALU-1>config>service# apipe 3 customer 1 create
A-ALU-1>config>service>apipe# sap 1/2/2.1:10/10
 sap-aggregation-group GroupName1
A-ALU-1>config>service>apipe>sap# exit
A-ALU-1>config>service>apipe# sap 1/2/2.1:20/20
 sap-aggregation-group GroupName1
A-ALU-1>config>service>apipe>sap# exit
A-ALU-1>config>service>apipe# sap 1/2/2.1:30/30
 sap-aggregation-group GroupName1
A-ALU-1>config>service>apipe>sap# exit
A-ALU-1>config>service>apipe#
```

The following example shows the output for an Apipe service with a SAP aggregation group containing three SAPs as aggregation group members.

```
A:ALU-41>config>service# info
------------------------------------
...
        apipe 3 customer 1 create
           sap-aggregation-group "GroupName1" create
           exit
           sap 1/2/2.1:10/10 sap-aggregation-group "GroupName1" create
           exit
           sap 1/2/2.1:20/20 sap-aggregation-group "GroupName1" create
           exit
           sap 1/2/2.1:30/30 sap-aggregation-group "GroupName1" create
           exit
        exit
...
------------------------------------
```

Creating a Cpipe Service

This section describes how to create a Cpipe service. Additional topics in this section include:

- Configuring Cpipe SAP Parameters
- Configuring Cpipe SDP Bindings

Use the following CLI syntax to create a Cpipe service.

CLI Syntax:
```
config>service# cpipe service-id [customer customer-id]
 [create] [vpn vpn-id] [vc-type {satop-e1 | satop-t1 |
 satop-e3 | satop-t3 | cesopsn | cesopsn-cas}] [vc-
 switching]
     description description-string
     service-mtu octets
     no shutdown
```

Example:
```
config>service# cpipe 234 customer 123 create vc-type
 cesopsn
config>service>cpipe# description "cpipe test"
config>service>cpipe# service-mtu 1400
config>service>cpipe# no shutdown
config>service>cpipe#
```

The following example shows the Cpipe service creation output for PE Router 1 (ALU-41).

```
A:ALU-41>config>service# info
------------------------------------
...
     cpipe 234 customer 123 create
            description "cpipe test"
            service-mtu 1400
            no shutdown
        exit
...
------------------------------------
A:ALU-41>config>service#
```

Configuring Cpipe SAP Parameters

Use the following CLI syntax to configure Cpipe SAP parameters. For ingress and egress configuration information, see Configuring Ingress and Egress SAP Parameters.

CLI Syntax:

```
config>service# cpipe service-id [customer customer-id]
 [create] [vpn vpn-id] [vc-type {satop-e1 | satop-t1 |
 satop-e3 | satop-t3 | cesopsn | cesopsn-cas}] [vc-
 switching]
    sap sap-id [create]
         cem
              [no] packet
              asym-delay-control [samples ksamples]
                 [min-repeat minutes]
                   jitter-buffer value | payload-size
                        size
                   payload-size size
              [no] report-alarm [stray] [malformed]
                 [pktloss]
                      [overrun] [underrun] [rpktloss]
                      [rfault] [rrdi]
              [no] rtp-header
         [no] collect-stats
         description description-string
         no description
         egress
              qos policy-id
              no qos
         ingress
              qos policy-id
              no qos
         [no] shutdown
```

Example:

```
A:ALU-41>config>service# cpipe 5 cesopsn
A:ALU-41>config>service>cpipe# sap 1/1/1.1 create
A:ALU-41>config>service>cpipe>sap# ingress
A:ALU-41>config>service>cpipe>sap>ingress# qos 102
A:ALU-41>config>service>cpipe>sap>ingress# exit
A:ALU-41>config>service>cpipe>sap# egress
A:ALU-41>config>service>cpipe>sap>egress# qos 103
A:ALU-41>config>service>cpipe>sap>egress# exit
A:ALU-41>config>service>cpipe>sap# no shutdown
A:ALU-41>config>service>cpipe>sap# exit
A:ALU-41>config>service>cpipe#
```

The following example shows the Cpipe SAP configuration output for PE Router 1 (ALU-41).

```
A:ALU-41>config>service# info
------------------------------------
...
        cpipe 5 customer 1 create
            description "cpipe test"
            service-mtu 1400
            sap 1/1/1.1 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            no shutdown
        exit
...
------------------------------------
```

To configure a basic local Cpipe service (SAP-to-SAP), enter the `sap` *`sap-id`* command twice with different port IDs in the same service configuration.

The following example shows a TDM SAP-to-SAP configuration:

```
A:ALU-41>config>service# info
------------------------------------
...
        cpipe 5 customer 1 create
            description "TDM sap2sap"
            service-mtu 1400
            sap 1/1/1.1
            sap 1/2/1.1
            no shutdown
        exit
...
------------------------------------
```

Use the following CLI syntax to configure a Cpipe SAP used by a Surveillance, Control, and Data Acquisition Support (SCADA) bridge. In this instance, the *sap-id* variable is in the format *slot/mda/bridge-id.branch-id*, where *bridge-id* represents an existing bridge number with a value of 1 to 16 and *branch-id* represents an existing branch number with a value of 1 to 32. Refer to the 7705 SAR OS Interface Configuration Guide, "Card, Adapter Card, and Port Command Reference", for information on how to configure bridges and branches for a SCADA application.

CLI Syntax:
```
config>service# cpipe service-id [customer customer-id]
  [create] vc-type cesopsn
      sap sap-id [create]
      description description-string
      no description
      [no] shutdown
```

Example:
```
A:ALU-41>config>service# cpipe 8 customer 1 create
 vc-type cesopsn
A:ALU-41>config>service>cpipe# sap 1/5/16.10 create
A:ALU-41>config>service>cpipe>sap# description "sap
 branch 10"
A:ALU-41>config>service>cpipe>sap# no shutdown
A:ALU-41>config>service>cpipe>sap# exit
A:ALU-41>config>service>cpipe#
```

The following example shows a SCADA bridge Cpipe SAP configuration:

```
A:ALU-41>config>service# info
------------------------------------
...
        cpipe 8 customer 1 vc-type cesopsn
            sap 1/5/16.10
            description "sap branch 10"
            no shutdown
        exit
```

Configuring Cpipe SDP Bindings

Use the following CLI syntax to create a spoke SDP binding with a Cpipe service. For SDP configuration information, see Configuring SDPs.

CLI Syntax:

```
config>service# cpipe service-id [customer customer-id]
  [create] [vpn vpn-id] [vc-type {satop-e1 | satop-t1 |
  satop-e3 | satop-t3 | cesopsn | cesopsn-cas}]
  [vc-switching]
      spoke-sdp sdp-id:vc-id [create]
           egress
                vc-label egress-vc-label
           ingress
                vc-label ingress-vc-label
           [no] shutdown
```

Example:

```
A:ALU-41>config>service# cpipe 5
A:ALU-41>config>service>cpipe# spoke-sdp 1:5 create
A:ALU-41>config>service>cpipe>spoke-sdp# no shutdown
A:ALU-41>config>service>cpipe>spoke-sdp# exit
```

The following example shows the Cpipe spoke SDP configuration output for PE Router 1 (ALU-41).

```
A:ALU-41>config>service# info
------------------------------------
...
        cpipe 5 customer 1 create
            description "cpipe test"
            service-mtu 1400
            sap 1/1/1.1 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:5 create
            exit
            no shutdown
        exit
...
------------------------------------
A:ALU-41>config>service#
```

Creating an Epipe Service

This section describes how to create an Epipe service. Additional topics in this section include:

- Configuring Epipe SAP Parameters
- Configuring Epipe SAP MEF 8 Parameters
- Configuring Epipe SAP Microwave Link Parameters for Interworking with TDM2Ethernet
- Configuring Epipe Spoke SDP Bindings

Use the following CLI syntax to create an Epipe service.

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

Example:
```
config>service# epipe 500 customer 5 create
config>service>epipe$ description "Local epipe service"
config>service>epipe# no shutdown
```

The following example shows the Epipe service creation output.

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

Configuring Epipe SAP Parameters

The 7705 SAR supports distributed Epipe service and local (SAP-to-SAP) Epipe service. A distributed Epipe consists of two SAPs on different nodes. A local Epipe consists of both SAPs on the same 7705 SAR. To configure a distributed Epipe service, you must configure service entities on the originating and far-end nodes.

Use the following CLI syntax to create distributed Epipe SAPs. For ingress and egress configuration information, see Configuring Ingress and Egress SAP Parameters. For SAP ETH-CFM configuration information, see Configuring ETH-CFM Parameters (802.1ag and Y.1731).

CLI Syntax:

```
config>service# epipe service-id [customer customer-id]
 [create] [vc-switching]
    sap sap-id [create]
         accounting-policy policy-id
         collect-stats
         description description-string
         no shutdown
         egress
              qos policy-id
         eth-cfm
         ingress
              filter [ip ip-filter-id]
              qos policy-id
```

Example:

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

ALU-2>config>service# epipe 5500 customer 5 create
config>service>epipe$ description "Distributed epipe
 service to west coast"
config>service>epipe# sap 1/1/4:550 create
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# qos 654
config>service>epipe>sap>ingress# exit
config>service>epipe>sap# egress
config>service>epipe>sap>egress# qos 432
config>service>epipe>sap>egress# exit
config>service>epipe>sap# no shutdown
config>service>epipe#
```

The following example shows the SAP configuration output for ALU-1 and ALU-2.

```
ALU-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
                    filter ip 1
                    qos 555
                exit
                egress
                    qos 627
                exit
            exit
        exit
...
----------------------------------------------
ALU-1>config>service#

ALU-2>config>service# info
----------------------------------------------
...
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to west coast"
            sap 1/1/4:550 create
                ingress
                    qos 654
               exit
               egress
                    qos 432
               exit
            exit
        exit
...
----------------------------------------------
ALU-2>config>service#
```

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

The following example shows an Ethernet SAP-to-SAP configuration:

```
A:ALU-4>config>service# info
------------------------------------
...
        epipe 2 customer 1 create
            description "Ethernet sap2sap"
               sap 1/1/1:1000
               sap 1/2/1:50
            no shutdown
        exit
...
------------------------------------
```

Configuring Epipe SAP MEF 8 Parameters

MEF 8 allows TDM services to be encapsulated across Epipes. To configure MEF 8, define one TDM SAP and one Ethernet SAP on the Epipe; see Configuring Epipe SAP Parameters for information on configuring an Ethernet SAP. The TDM SAP for the Epipe must include a local and remote ECID and a far-end destination MAC address. The TDM port's MAC address is used as the source MAC address for the circuit.

TDM can also be encapsulated into Ethernet which is then encapsulated in MPLS (or GRE) by configuring an Epipe with a TDM SAP and a spoke SDP. See Configuring SDPs for information on configuring a spoke SDP. The TDM SAP configured in the Epipe must include a local and remote ECID and a far-end destination MAC address. The TDM port's MAC address is used as the source MAC address for the circuit.

Use the following CLI syntax to configure MEF 8 parameters:

CLI Syntax:

```
config>service# epipe service-id [customer customer-id]
  [create] [vpn vpn-id] [vc-switching]
     description description-string
     sap sap-id [create]
          accounting-policy policy-id
          cem
               packet
                    jitter-buffer value | payload-size
                         size
                    payload-size size
               local-ecid value
               no local-ecid
               remote-ecid value
               no remote-ecid
               remote-mac ieee-mac-addr
               no remote-mac
               report-alarm [stray] [malformed]
                  [pktloss] [overrun] [underrun]
                  [rpktloss] [rfault] [rrdi]
          description description-string
```

The following example shows a TDM SAP-to-Ethernet SAP MEF 8 configuration and a sample output.

Example:

```
ALU-1>epipe 1
config>service>epipe# description "Test Epipe for
 service ID 1"
config>service>epipe# sap 1/2/4:1 create
config>service>epipe>sap# description "test SAP for
 service ID 1"
config>service>epipe>sap# exit
```

```
                config>service>epipe# sap 1/1/1.1 create
                config>service>epipe>sap# description "test SAP2 for
                 service ID 1"
                config>service>epipe>sap# cem
                config>service>epipe>sap>cem# report alarm rpktloss
                 rfault rrdi
                config>service>epipe>sap>cem# local-ecid 1
                config>service>epipe>sap>cem# remote-ecid 1
                config>service>epipe>sap>cem# remote-mac
                 10:00:50:00:00:02
                config>service>epipe>sap>cem# exit
                config>service>epipe>sap# exit
                config>service>epipe# no shutdown

*A:7705:Dut-A>config>service>epipe# info
----------------------------------------------
            description "Default epipe description for service id 1"
            sap 1/2/4:1 create
                description "Test SAP for service ID 1"
            exit
            sap 1/1/1.1 create
                description "Test SAP2 for service ID 1"
                cem
                    report-alarm rpktloss rfault rrdi
                    local-ecid 1
                    remote-ecid 1
                    remote-mac 10:00:50:00:00:02
                exit
            exit
            no shutdown
----------------------------------------------
*A:7705:Dut-A>config>service>epipe# back
```

The following example shows a TDM SAP-to-spoke SDP MEF 8 configuration and a sample output. See Configuring SDPs for more information about SDPs.

Example:

```
config>service>epipe 2
config>service>epipe# description "Test Epipe for
 service ID 2"
config>service>epipe# sap 1/1/1.2 create
config>service>epipe>sap# description "test SAP for
 service ID 2"
config>service>epipe>sap# exit
config>service>epipe# exit
config>service# sdp 2 mpls create
config>service>sdp# description "MPLS-10.10.10.104"
config>service>sdp# far-end "10.10.10.104"
config>service>sdp# no shutdown
config>service>sdp# exit
config service# epipe 2 sap
config>service>epipe>sap# cem
```

```
            config>service>epipe>sap>cem# report alarm rpktloss
             rfault rrdi
            config>service>epipe>sap>cem# local-ecid 2
            config>service>epipe>sap>cem# remote-ecid 2
            config>service>epipe>sap>cem# remote-mac
             10:00:50:00:00:02
            config>service>epipe>sap>cem# exit
            config>service>epipe>sap# exit
            config>service>epipe# no shutdown

*A:7705:Dut-A>config>service# epipe 2
*A:7705:Dut-A>config>service>epipe# info
----------------------------------------------
            description "Default epipe description for service id 2"
            sap 1/1/1.2 create
                description "Default sap description for service id 2"
                cem
                    report-alarm rpktloss rfault rrdi
                    local-ecid 2
                    remote-ecid 2
                    remote-mac 10:00:50:00:00:02
                exit
            exit
            spoke-sdp 2:2 create
            exit
            no shutdown
----------------------------------------------
*A:7705:Dut-A>config>service>epipe#
```

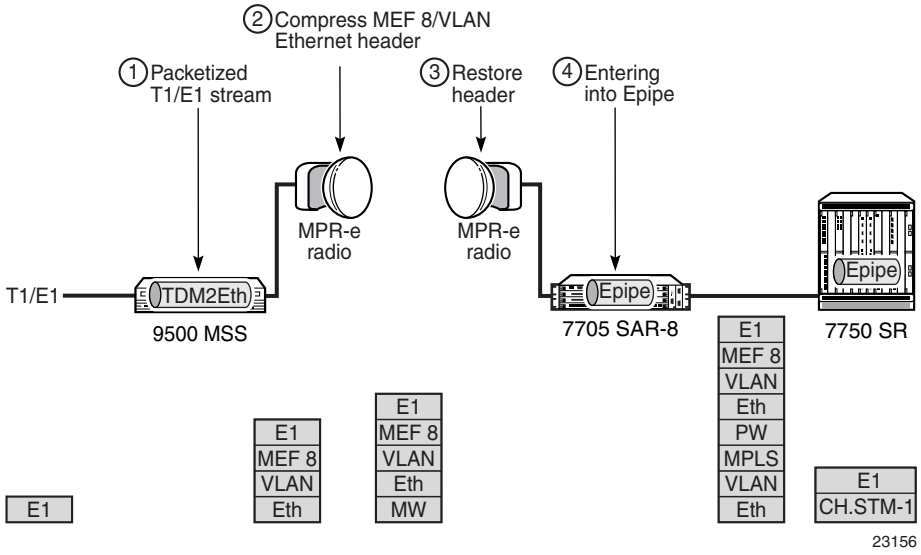
Configuring Epipe SAP Microwave Link Parameters for Interworking with TDM2Ethernet

In a Microwave Awareness environment, a mixed microwave link scenario may exist where an access link site has an MPR-e radio connected to a 7705 SAR-8 or 7705 SAR-18 as a standalone network element as well as an MPR-e radio connected to a 9500 MSS as an indoor unit, as shown in Figure 68.

When TDM2Ethernet transport is required, the 9500 MSS packetizes the TDM2Ethernet stream using a MEF 8 frame format; the 7705 SAR-8 or 7705 SAR-18 acts as the MEF 8 endpoint. From a packet flow perspective, the TDM2Ethernet MEF 8 stream is a VLAN flow that can be handled by a VLAN SAP into an Epipe service going into the 7705 SAR-8 or 7705 SAR-18. To optimize the microwave bandwidth for the TDM2Ethernet MEF 8 transport, the packet header is compressed at the ingress MPR-e radio (connected to the 9500 MSS) and restored at the egress MPR-e radio (connected to the 7705 SAR-8 or 7705 SAR-18). In order for the egress MPR-e radio to perform this function, the parameters listed below are configured for an Epipe SAP on a 7705 SAR-8 or 7705 SAR-18; see Epipe Service Configuration Commands for more information:

- TDM2Ethernet compression on the Epipe
- Real-time Transport Protocol (RTP) header, if required (see RTP Header for a description of how an RTP header is applied)
- source 9500 MSS MAC address and destination 7705 SAR-8 or 7705 SAR-18 MAC address

Figure 68: Mixed Microwave Link Scenario



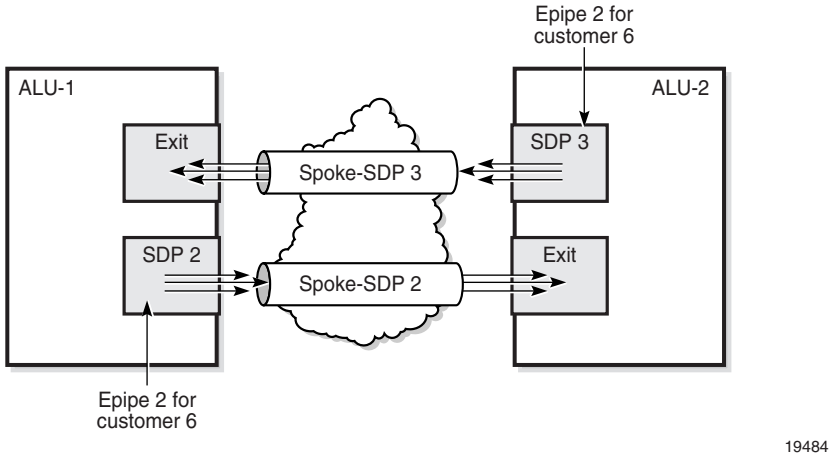
Apply Epipe SAP microwave link parameters for interworking with TDM2Ethernet as shown in the CLI syntax example below:

CLI Syntax:
```
config>service>epipe epipe-id customer customer-id
 create
     sap sap-id create
          mw
               compression source-mac destination-mac
                  [rtp]
```

Configuring Epipe Spoke SDP Bindings

Figure 69 shows 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. The `spoke-sdp sdp-id:vc-id` must match on both sides.

Figure 69: SDPs—Unidirectional Tunnels



An SDP must exist before it can be used with a spoke SDP. For SDP configuration information, see Configuring SDPs.

To configure a spoke SDP for ETH-CFM support, see Configuring ETH-CFM Parameters (802.1ag and Y.1731).

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]
  [create] [vc-switching]
      spoke-sdp sdp-id:vc-id [vc-type {ether|vlan}]
        [create] vlan-vc-tag 0..4094
          egress
               vc-label egress-vc-label
          eth-cfm
          ingress
               vc-label ingress-vc-label
          no shutdown
```

Example:

```
ALU-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

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

The following example shows the configuration output for binding an Epipe service between ALU-1 and ALU-2. This example assumes the SAPs have already been configured (see Configuring Epipe SAP Parameters).

```
ALU-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
                    filter ip 1
                    qos 555
                exit
                egress
                    qos 627
                exit
            exit
```

```
            spoke-sdp 2:123 create
                ingress
                    vc-label 6600
                exit
                egress
                    vc-label 5500
                exit
            exit
            no shutdown
        exit
...
-----------------------------------------------
ALU-1>config>service#



ALU-2>config>service# info
-----------------------------------------------
...
exit
        epipe 5500 customer 5 vpn 5500 create
            description "Distributed epipe service to west coast"
            sap 1/1/4:550 create
                ingress
                    qos 654
                exit
                egress
                    qos 432
                exit
            exit
            spoke-sdp 2:123 create
                ingress
                    vc-label 5500
                exit
                egress
                    vc-label 6600
                exit
            exit
            no shutdown
        exit
...
-----------------------------------------------
```

Creating an Fpipe Service

This section describes how to create an Fpipe service. Additional topics in this section include:

- Configuring Fpipe SAP Parameters
- Configuring Fpipe SDP Bindings

Use the following CLI syntax to create an Fpipe service.

CLI Syntax:

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

PE router 1 (A:ALU-41):

Example:

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

PE router 2 (A:ALU-42):

Example:

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

The following example shows the Fpipe service creation output.

PE Router 1 (ALU-41):

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

PE Router 2 (ALU-42):

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

Configuring Fpipe SAP Parameters

Use the following CLI syntax to configure Fpipe SAP parameters.

For ingress and egress configuration information, see Configuring Ingress and Egress SAP Parameters.

CLI Syntax:

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

Example:

```
A:ALU-41>config>service# fpipe 1
A:ALU-41>config>service>fpipe# sap 1/2/1:16 create
A:ALU-41>config>service>fpipe>sap# ingress
A:ALU-41>config>service>fpipe>sap>ingress# qos 101
A:ALU-41>config>service>fpipe>sap>ingress# exit
A:ALU-41>config>service>fpipe>sap# egress
A:ALU-41>config>service>fpipe>sap>egress# qos 1020
A:ALU-41>config>service>fpipe>sap>egress# exit
A:ALU-41>config>service>fpipe>sap# no shutdown
A:ALU-41>config>service>fpipe>sap# exit
A:ALU-41>config>service>fpipe#
```

The following example shows the Fpipe SAP configuration output for PE Router 1 (ALU-41).

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

To configure a basic local Fpipe service (SAP-to-SAP), enter the sap *sap-id* command twice with different port IDs in the same service configuration.

The following example shows an FR SAP-to-SAP configuration:

```
A:ALU-41>config>service# info
------------------------------------
...
        fpipe 1 customer 1 create
            description "FR sap2sap"
            service-mtu 1400
            sap 1/2/1:16
            sap 1/3/1:20
            no shutdown
        exit
...
------------------------------------
```

Configuring Fpipe SDP Bindings

Use the following CLI syntax to create a spoke SDP binding with an Fpipe service (for distributed service). For SDP configuration information, see Configuring SDPs.

CLI Syntax:

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

Example:

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

The following example shows the Fpipe spoke SDP configuration output for PE Router 1 (ALU-41).

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

Creating an Hpipe Service

This section describes how to create an Hpipe service. Additional topics in this section include:

- Configuring Hpipe SAP Parameters
- Configuring Hpipe SDP Bindings

Use the following CLI syntax to create an Hpipe service.

CLI Syntax:

```
config>service# hpipe service-id [customer customer-id]
  [create] [vpn vpn-id] [vc-type {hdlc}] [vc-switching]
      description description-string
      service-mtu octets
      no shutdown
```

PE router 1 (ALU-41):

Example:

```
A:ALU-41>config>service# hpipe 4 customer 1 create
A:ALU-41>config>service>hpipe# description "hpipe test"
A:ALU-41>config>service>hpipe# service-mtu 1400
A:ALU-41>config>service>hpipe# no shutdown
A:ALU-41>config>service>hpipe#
```

PE router 2 (ALU-42):

Example:

```
A:ALU-42>config>service# hpipe 4 customer 1 create
A:ALU-42>config>service>hpipe# description "hpipe test"
A:ALU-42>config>service>hpipe# service-mtu 1400
A:ALU-42>config>service>hpipe# no shutdown
A:ALU-42>config>service>hpipe#
```

The following example shows the Hpipe service creation output.

PE Router 1 (ALU-41):

```
A:ALU-41>config>service# info
----------------------------------
...
        hpipe 4 customer 1 create
            description "hpipe test"
            service-mtu 1400
            no shutdown
        exit
...
----------------------------------
A:ALU-41>config>service#
```

PE Router 2 (ALU-42):

```
A:ALU-42>config>service# info
------------------------------------
...
        hpipe 4 customer 1 create
            description "hpipe test"
            service-mtu 1400
            no shutdown
        exit
```

Configuring Hpipe SAP Parameters

Use the following CLI syntax to configure Hpipe SAP parameters. For ingress and egress configuration information, see Configuring Ingress and Egress SAP Parameters.

CLI Syntax:

```
config>service# hpipe service-id [customer customer-id]
  [create] [vpn vpn-id] [vc-type {hdlc}] [vc-switching]
      sap sap-id [create]
           accounting-policy acct-policy-id
           collect-stats
           description description-string
           egress
                qos policy-id
           ingress
                qos policy-id [shared queuing]
           no shutdown
```

Example:

```
A:ALU-41>config>service# hpipe 4
A:ALU-41>config>service>hpipe# sap 1/4/1.20 create
A:ALU-41>config>service>hpipe>sap# ingress
A:ALU-41>config>service>hpipe>sap>ingress# qos 102
A:ALU-41>config>service>hpipe>sap>ingress# exit
A:ALU-41>config>service>hpipe>sap# egress
A:ALU-41>config>service>hpipe>sap>egress# qos 103
A:ALU-41>config>service>hpipe>sap>egress# exit
A:ALU-41>config>service>hpipe>sap# no shutdown
A:ALU-41>config>service>hpipe>sap# exit
A:ALU-41>config>service>hpipe#
```

The following example shows the Hpipe SAP configuration output for PE Router 1 (ALU-41).

```
A:ALU-41>config>service# info
------------------------------------
...
        hpipe 4 customer 1 create
            description "hpipe test"
            service-mtu 1400
            sap 1/4/1.20 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            no shutdown
        exit
...
------------------------------------
```

To configure a basic local Hpipe service (SAP-to-SAP), enter the sap *sap-id* command twice with different port IDs in the same service configuration.

The following example shows an HDLC SAP-to-SAP configuration:

```
A:ALU-41>config>service# info
------------------------------------
...
        hpipe 4 customer 1 create
            description "HDLC sap2sap"
            service-mtu 1514
            sap 1/4/1.20
            sap 1/5/1.10
            no shutdown
        exit
...
------------------------------------
```

Configuring Hpipe SDP Bindings

Use the following CLI syntax to create a spoke SDP binding with an Hpipe service (for distributed service). For SDP configuration information, see Configuring SDPs.

CLI Syntax:

```
config>service# hpipe service-id [customer customer-id]
  [create] [vpn vpn-id] [vc-type {hdlc}] [vc-switching]
     spoke-sdp sdp-id:vc-id [create]
          egress
               vc-label egress-vc-label
          ingress
               vc-label ingress-vc-label
          no shutdown
```

Example:

```
A:ALU-41>config>service# hpipe 4
A:ALU-41>config>service>hpipe# spoke-sdp 1:4 create
A:ALU-41>config>service>hpipe>spoke-sdp# no shutdown
A:ALU-41>config>service>hpipe>spoke-sdp# exit
```

The following example shows the Hpipe spoke SDP configuration output for PE Router 1 (ALU-41).

```
A:ALU-41>config>service# info
------------------------------------
...
        hpipe 4 customer 1 create
            description "hpipe test"
            service-mtu 1400
            sap 1/4/1.20 create
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:4 create
            exit
            no shutdown
        exit
...
------------------------------------
A:ALU-41>config>service#
```

Creating an Ipipe Service

This section describes how to create an Ipipe service. Additional topics in this section include:

- Configuring Ipipe SAP Parameters
- Configuring Ipipe SDP Bindings

Use the following CLI syntax to create an Ipipe service.

CLI Syntax:

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

The following example shows an Ipipe configuration:

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

Configuring Ipipe SAP Parameters

The following example shows an Ipipe SAP configuration:

```
A:ALU-48>config>service# info
----------------------------------------------
...
        ipipe 202 customer 1 create
            sap 1/1/2:444 create
                description "eth_ipipe"
                ce-address 31.31.31.1
            exit
             spoke-sdp 16:516 create
                ce-address 31.31.31.2
            exit
            no shutdown
        exit
...
----------------------------------------------
```

The following shows a PPP to Ethernet local Ipipe example:

Example:

```
config>service# ipipe 206 customer 1 create
config>service>ipipe$ sap 1/1/2:447 create
config>service>ipipe>sap$ description "eth_ppp_ipipe"
config>service>ipipe>sap$ ce-address 33.33.33.1
config>service>ipipe>sap$ no shutdown
config>service>ipipe>sap$ exit
config>service>ipipe# spoke-sdp 15:516 create
config>service>ipipe>sap>spoke-sdp$ ce-address
 33.33.33.2
config>service>ipipe>sap-spoke-sdp$ exit
config>service>ipipe>$ exit
config>service>ipipe# no shutdown
config>service>ipipe# exit
config>service#
```

The following shows the output:

```
A:ALU-48>config>service# info
----------------------------------------------
ipipe 206 customer 1 create
            sap 1/1/2:447 create
                description "eth_ppp_ipipe"
                ce-address 33.33.33.1
            exit
            spoke-sdp 15:516 create
                ce-address 33.33.33.2
             exit
            exit
            no shutdown
        exit
exit
----------------------------------------------
```

The configuration of a Cisco HDLC SAP on an Ipipe is similar to the above example, except that the *sap-id* has the form *slot*/*mda*/*port*.*channel-group* (1/2/2.1).

Configuring Ipipe SDP Bindings

The following shows an Ipipe SDP configuration example:

```
A:ALU-48>config>service# info
----------------------------------------------
...
        sdp 16 mpls create
            far-end 4.4.4.4
            ldp
            path-mtu 1600
            keep-alive
                shutdown
            exit
            no shutdown
        exit
...
        ipipe 207 customer 1 create
            shutdown
            sap 1/1/2:449 create
                description "Remote_Ipipe"
                ce-address 34.34.34.1
            exit
            spoke-sdp 16:516 create
                ce-address 34.34.34.2
            exit
```

Configuring PW Switching

The `vc-switching` parameter defines a VLL service as a PW switching point, also called an S-PE. This is the point where a VLL service switches from one PW type to another. The `vc-switching` parameter must be specified when the VLL service is created.

Once a VLL service is configured as an S-PE, you cannot add a SAP to the configuration. The following example shows the error message generated by the CLI if you attempt to create a SAP on a VLL PW switching service.

Example:

```
*A:ALU>config>service>epipe 2 customer 1 create
vc-switching
*A:ALU>config>service>epipe$ sap 1/5/1 create
MINOR: SVCMGR #1311 SAP is not allowed under PW switching
 service
*A:ALU>config>service>epipe$
```

Use the following CLI syntax to configure pseudowire switching VLL services.

CLI Syntax:

```
config>service# apipe service-id [customer customer-id]
 [create] [vpn vpn-id] [vc-type {atm-vcc|atm-vpc|
 atm-cell}] [vc-switching]
```

```
    description description-string
    spoke-sdp sdp-id:vc-id
    no shutdown
```

CLI Syntax:
```
config>service# cpipe service-id [customer customer-id]
 [create] [vpn vpn-id] [vc-type {satop-e1 | satop-t1 |
 satop-e3 | satop-t3 | cesopsn | cesopsn-cas}]
 [vc-switching]
    description description-string
    service-mtu octets
    no shutdown
```

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

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

CLI Syntax:
```
config>service# hpipe service-id [customer customer-id]
 [create] [vpn vpn-id] [vc-type {hdlc}] [vc-switching]
    description description-string
    spoke-sdp sdp-id:vc-id
    no shutdown
```

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

The following shows an example of the command usage to configure a VLL service as a PW switching service:

Example:
```
ALU-1>config>service# apipe 5 customer 1 vpn 1
 vc-switching create
config>service>apipe$ description "Default apipe
 description for service id 100"
config>service>apipe# spoke-sdp 3:1 create
config>service>apipe>spoke-sdp# exit
config>service>apipe# spoke-sdp 6:200 create
config>service>apipe>spoke-sdp# exit
config>service>apipe# no shutdown
```

The following example shows configurations for each service:

```
*A:ALA-48>config>service# info
----------------------------------------------
apipe 100 customer 1 vpn 1 vc-switching create
description "Default apipe description for service id 100"
spoke-sdp 3:1 create
exit
spoke-sdp 6:200 create
exit
no shutdown
exit
...
cpipe 107 customer 1 vpn 107 vc-switching vc-type satop-e1 create
description "Default cpipe description for service id 107"
spoke-sdp 3:7 create
exit
spoke-sdp 6:207 create
exit
no shutdown
exit
...
epipe 108 customer 1 vpn 108 vc-switching create
description "Default epipe description for service id 108"
spoke-sdp 3:8 create
exit
spoke-sdp 6:208 create
exit
no shutdown
exit
...
fpipe 109 customer 1 vpn 75 vc-switching create
description "Default fpipe description for service id 109"
spoke-sdp 3:9 create
exit
spoke-sdp 6:209 create
exit
no shutdown
exit
...
hpipe 110 customer 1 vpn 76 vc-switching create
description "Default hpipe description for service id 110"
spoke-sdp 3:10 create
exit
spoke-sdp 6:210 create
exit
no shutdown
exit
...
ipipe 111 customer 1 vpn 77 vc-switching create
description "Default ipipe description for service id 111"
spoke-sdp 3:11 create
exit
spoke-sdp 6:211 create
exit
no shutdown
exit
...
----------------------------------------------
```

Configuring Ingress and Egress SAP Parameters

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

Ingress and egress QoS SAP parameters can be applied to distributed and local Epipe, Fpipe, Hpipe, and Ipipe service SAPs, and to local Apipe and Cpipe service SAPs.

Ingress and egress QoS parameters can also be applied to SAP aggregation groups for ATM VCC VLL services. QoS parameters cannot be applied to SAPs that are members of the aggregation group.

By default, there are no IP filters associated with interfaces or services. IP filter policies can be applied to ingress Epipe and Ipipe service SAPs.

Example:

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

The following example shows the Epipe SAP ingress and egress configuration output.

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

Using the Control Word

The control word is mandatory for Cpipe SAToP and CESoPSN configurations, and Fpipe one-to-one configurations. The control word is optional for Apipe, Epipe, Hpipe, and Ipipe configurations.

When the control word is enabled, the Admin Control Word is set to Preferred. Both sides of the VLL must be configured with a matching control word, either both enabled or both disabled, for the pipe to be up.

The control word state will be set to True or False depending on what is configured, either enabled (True) or disabled (False).

Example:
```
config>service# cpipe 2100 customer 1
config>service>cpipe$ description ”Default cpipe
 description for service id 2100”
config>service>cpipe$ sap 1/2/7.1:4 create
config>service>cpipe>sap$ description “Default sap
 description for service id 2100”
config>service>cpipe>sap$ exit
config>service>cpipe# spoke-sdp 1:2001 create
config>service>cpipe>spoke-sdp$ control-word
config>service>cpipe>spoke-sdp$ exit
config>service>cpipe# no shutdown
```

The following example shows the control word configuration output for a Cpipe service.

```
*A:ALU-Dut-B>config>service>cpipe# info
----------------------------------------------
            description "Default cpipe description for service id 2100"
            sap 1/2/7.1:4 create
                description "Default sap description for service id 2100"
            exit
            spoke-sdp 1:2001 create
                control-word
            exit
            no shutdown
----------------------------------------------
*A:ALU-Dut-B>config>service>cpipe#
```

The control word cannot be disabled on Cpipe services. To disable the control word option on Apipe, Epipe, Fpipe, Hpipe, or Ipipe services, use the `no control-word` command.

Example:
```
config>service>apipe# spoke-sdp 1:2001 no control-word
config>service>apipe>spoke-sdp$ exit
```

Configuring PW Redundancy

This section describes configuring PW redundancy. Additional topics in this section include:

- Configuring PW Redundancy—Standby Signaling
- Configuring PW Redundancy—ICB

For PW redundancy, create an explicit endpoint and then assign a primary spoke SDP and up to three secondary spoke SDPs, or up to four secondary spoke SDPs with no primary spoke SDP, to that endpoint.

CLI Syntax:

```
config>service# cpipe service-id [customer customer-id]
  [create]
      endpoint endpoint-name [create]
      spoke-sdp sdp-id:vc-id endpoint endpoint-name
         [create]
           precedence precedence-value
      no shutdown
```

Example:

```
config>service# cpipe 2100
config>service>cpipe$ endpoint "Endpoint_Y" create
config>service>cpipe$ spoke-sdp 1:100 endpoint
    "Endpoint_Y" create
config>service>cpipe>spoke-sdp$ precedence primary
config>service>cpipe$ spoke-sdp 2:200 endpoint
    "Endpoint_Y" create
config>service>cpipe>spoke-sdp$ precedence 1
no shutdown
```

The following example shows the PW redundancy configuration output for a Cpipe service.

```
*A:7705:Dut-C>config>service>cpipe# info
----------------------------------------------
            endpoint "Endpoint_Y" create
            exit
            spoke-sdp 1:100 endpoint "Endpoint_Y" create
                precedence primary
            exit
            spoke-sdp 2:200 endpoint "Endpoint_Y" create
                precedence 1
            exit
----------------------------------------------------
*A:7705:Dut-C>config>service>cpipe#
```

Configuring PW Redundancy—Standby Signaling

For standby signaling, create an endpoint as shown in the preceding section, but specify an SDP as the standby signaling master and the far-end endpoints as standby signaling slaves.

In the following example, precedence primary indicates that the spoke SDP is the active SDP and precedence 1 indicates that the spoke SDP is the standby SDP.

To create a standby signaling master:

CLI Syntax:
```
config>service# cpipe service-id [customer customer-id]
  [create]
      endpoint endpoint-name [create]
      standby-signaling-master
      exit
      spoke-sdp sdp-id:vc-id endpoint endpoint-name
         [create]
           precedence precedence-value
      no shutdown
```

Example:
```
config>service cpipe 555
config>service>cpipe# endpoint "stdbyMaster" create
config>service>cpipe>endpoint# standby-signaling-master
config>service>cpipe>endpoint# exit
config>service>cpipe# spoke-sdp 5:555 endpoint
    "stdbyMaster" create
config>service>cpipe>spoke-sdp$ precedence primary
exit
config>service>cpipe# spoke-sdp 131:555 endpoint
    "stdbyMaster" create
config>service>cpipe>spoke-sdp$ precedence 1
exit
```

To create the far-end endpoints as standby signaling slaves:

CLI Syntax:
```
config>service# cpipe service-id [customer customer-id]
  [create]
      endpoint endpoint-name [create]
      standby-signaling-slave
      exit
      spoke-sdp sdp-id:vc-id endpoint endpoint-name
         [create]
           precedence precedence-value
      no shutdown
```

Example:
```
config>service cpipe 555
config>service>cpipe# endpoint "FarendA" create
config>service>cpipe>endpoint# standby-signaling-slave
```

```
config>service>cpipe>endpoint# exit
config>service>cpipe# spoke-sdp 130:555 endpoint
    "FarendA" create
config>service>cpipe>spoke-sdp$ precedence primary
exit
```

Example:
```
config>service cpipe 555
config>service>cpipe# endpoint "FarendB" create
config>service>cpipe>endpoint# standby-signaling-slave
config>service>cpipe>endpoint# exit
config>service>cpipe# spoke-sdp 130:555 endpoint
    "FarendB" create
config>service>cpipe>spoke-sdp$ precedence primary
exit
```

Configuring PW Redundancy—ICB

ICB (Inter-Chassis Backup) spoke SDPs are supported for use with Cpipe services in an MC-APS configuration and Epipe services in an MC-LAG configuration. ICB improves switch times, provides additional protection in case of network failures, and reduces packet loss when an active endpoint is switched from a failed MC-APS or MC-LAG node to a protection node.

To configure ICB, create explicit endpoints on the access and network side on the active node and then assign a primary spoke SDP on the reverse side of the protection node.

CLI Syntax:
```
config>service# cpipe service-id [customer customer-id]
 [create]
    description description-string
    endpoint endpoint-name [create]
    service-mtu octets
    standby-signaling-master
    exit
    sap sap-id [create]
    spoke-sdp sdp-id:vc-id endpoint endpoint-name [icb]
       [create]
         precedence precedence-value
    no shutdown
```

To configure ICB on the active node:

Example:
```
config>service cpipe 555
config>service>cpipe# description "Cpipe for service id
 1"
config>service>cpipe# service-mtu 1000
config>service>cpipe# endpoint "X" create
```

```
config>service>cpipe>endpoint# exit
config>service>cpipe# endpoint "Y" create
config>service>cpipe>endpoint# exit
config>service>cpipe# sap aps-1.1.1.1.1 endpoint "X"
 create
config>service>cpipe>sap# description "SAP for service
 id 1"
config>service>cpipe>sap# exit
config>service>cpipe# spoke-sdp 2003:1 endpoint
    "Y" create
config>service>cpipe>spoke-sdp# exit
config>service>cpipe# spoke-sdp 2004:1001 endpoint
    "X" icb create
config>service>cpipe>spoke-sdp# exit
config>service>cpipe# spoke-sdp 2004:1001 endpoint
    "Y" icb create
config>service>cpipe>spoke-sdp# exit
config>service>cpipe# no shutdown
exit
```

To configure ICB on the protection node:

Example:

```
config>service cpipe 1 vc-type cesopsn create
config>service>cpipe# service-mtu 1000
config>service>cpipe# endpoint "X" create
config>service>cpipe>endpoint# exit
config>service>cpipe# endpoint "Y" create
config>service>cpipe>endpoint# exit
config>service>cpipe# sap aps-1.1.1.1.1 endpoint "X"
 create
config>service>cpipe>sap# description "SAP description
 for service id 1"
config>service>cpipe>sap# exit
config>service>cpipe# spoke-sdp 4003:1 endpoint
    "Y" create
config>service>cpipe>spoke-sdp# exit
config>service>cpipe# spoke-sdp 4002:1001 endpoint
    "Y" icb create
config>service>cpipe>spoke-sdp# exit
config>service>cpipe# spoke-sdp 4002:2001 endpoint
    "X" icb create
config>service>cpipe>spoke-sdp# exit
config>service>cpipe# no shutdown
exit
```

Service Management Tasks

The service management tasks are similar for Apipe, Cpipe, Epipe, and Ipipe services. This section discusses the following service management tasks:

- Modifying Service Parameters
- Disabling a Service
- Re-enabling a Service
- Deleting a Service

Modifying Service Parameters

Use the `show service service-using` command to display a list of configured VLL services.

To modify a VLL service:

1. Access the specific account by specifying the service ID.
2. Enter the service parameter to modify and then enter the new information.

PE router 1 (A:ALU-41):

Example:

```
A:ALU-41>config>service# apipe 5
A:ALU-41>config>service>apipe# sap 1/1/1.1:0/32 create
A:ALU-41>config>service>apipe>sap# accounting-policy 2
A:ALU-41>config>service>apipe>sap# exit
A:ALU-41>config>service>apipe# spoke-sdp 1:4
A:ALU-41>config>service>apipe>spoke-sdp# egress
A:ALU-41>config>service>apipe>spoke-sdp>egress# vc-
 label 2048
A:ALU-41>config>service>apipe>spoke-sdp>egress# exit
A:ALU-41>config>service>apipe>spoke-sdp# ingress
A:ALU-41>config>service>apipe>spoke-sdp>ingress# vc-
 label 18431
A:ALU-41>config>service>apipe>spoke-sdp>ingress# exit
A:ALU-41>config>service>apipe>spoke-sdp# exit
A:ALU-41>config>service>apipe#
```

PE router 2 (A:ALU-42):

Example:

```
A:ALU-42>config>service# apipe 5
A:ALU-42>config>service>apipe# sap 2/2/2.1:0/32 create
A:ALU-42>config>service>apipe>sap# accounting-policy 2
A:ALU-42>config>service>apipe>sap# exit
A:ALU-42>config>service>apipe# spoke-sdp 1:4
A:ALU-42>config>service>apipe>spoke-sdp# egress
A:ALU-42>config>service>apipe>spoke-sdp>egress# vc-
 label 18431
A:ALU-42>config>service>apipe>spoke-sdp>egress# exit
A:ALU-41>config>service>apipe>spoke-sdp# ingress
A:ALU-41>config>service>apipe>spoke-sdp>ingress# vc-
 label 2043
A:ALU-41>config>service>apipe>spoke-sdp>ingress# exit
A:ALU-42>config>service>apipe>spoke-sdp# exit
A:ALU-42>config>service>apipe#
```

The following example shows the configuration output when adding an accounting-policy to an existing SAP and modifying the spoke-sdp parameters on an existing Apipe service for PE Router 1 (ALU-41) and PE Router 2 (ALU-42).

Use a similar syntax to modify Cpipe, Epipe, and Ipipe services.

```
A:ALU-41>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 1/1/1.1:0/32 create
               accounting-policy 2
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:4 create
               egress
                   vc-label 2048
               ingress
                   vc-label 18431
          exit
            no shutdown
        exit
...
------------------------------------
A:ALU-41>config>service#
```

```
A:ALU-42>config>service# info
----------------------------------
...
        apipe 5 customer 1 create
            description "apipe test"
            service-mtu 1400
            sap 2/2/2.1:0/32 create
               accounting-policy 2
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:4 create
               egress
                   vc-label 18431
               ingress
                   vc-label 2048
          exit
            no shutdown
        exit
...
----------------------------------
A:ALU-42>config>service#
```

Disabling a Service

A service can be shut down without deleting the service parameters.

Use the `shutdown` command to shut down a VLL service. The following CLI syntax shows the command to shut down an Apipe service. Use a similar syntax to shut down Cpipe, Epipe, Fpipe, Hpipe, and Ipipe services.

CLI Syntax:

```
config>service#
apipe service-id
     shutdown
```

PE router 1 (A:ALU-41):

Example:

```
A:ALU-41>config>service# apipe 5
A:ALU-41>config>service>apipe# shutdown
A:ALU-41>config>service>apipe# exit
```

PE router 2 (A:ALU-42):

Example:

```
A:ALU-42>config>service# apipe 5
A:ALU-42>config>service>apipe# shutdown
A:ALU-42>config>service>apipe# exit
```

The following example shows the configuration output for deleting an Apipe service on PE Router 1 (ALU-41) and PE Router 2 (ALU-42).

```
A:ALU-41>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            shutdown
            description "apipe test"
            service-mtu 1400
            sap 1/1/1.1:0/32 create
               accounting-policy 2
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:4 create
               egress
                   vc-label 16
            exit
            no shutdown
        exit
...
------------------------------------
A:ALU-41>config>service#

A:ALU-42>config>service# info
------------------------------------
...
        apipe 5 customer 1 create
            shutdown
            description "apipe test"
            service-mtu 1400
            sap 2/2/2.1:0/32 create
               accounting-policy 2
               ingress
                   qos 102
               exit
               egress
                   qos 103
               exit
            exit
            spoke-sdp 1:4 create
               egress
                   vc-label 16
            exit
        exit
...
------------------------------------
A:ALU-42>config>service#
```

Re-enabling a Service

Use the `no shutdown` command to re-enable a previously disabled VLL service. The following CLI syntax shows the command to re-enable an Apipe service. Use a similar syntax to re-enable Cpipe, Epipe, Fpipe, Hpipe, and Ipipe services.

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

PE router 1 (A:ALU-41):

Example:
```
A:ALU-41>config>service# apipe 5
A:ALU-41>config>service>apipe# no shutdown
A:ALU-41>config>service>apipe# exit
```

PE router 2 (A:ALU-42):

Example:
```
A:ALU-42>config>service# apipe 5
A:ALU-42>config>service>apipe# no shutdown
A:ALU-42>config>service>apipe# exit
```

Deleting a Service

Use the `shutdown` command to delete a VLL service. The SAP, and any associated protocols and spoke SDPs, must be deleted from the VLL service before the VLL service can be deleted.

Perform the following steps to delete a service:

1. Shut down the SAP and SDP.
2. Delete the SAP and SDP.
3. Shut down the service.

Use the following syntax to delete Apipe services. Use a similar syntax to delete Cpipe, Epipe, Fpipe, Hpipe, and Ipipe services.

CLI Syntax:

```
config>service#
     apipe service-id
          sap sap-id
               shutdown
               exit
          no sap sap-id
          spoke-sdp [sdp-id:vc-id]
               shutdown
               exit
          no spoke-sdp [sdp-id:vc-id]
          shutdown
          exit
     no apipe service-id
```

Example:

```
A:ALU-41>config>service# apipe 5
A:ALU-41>config>service>apipe# sap 1/1/1.1:0/32
A:ALU-41>config>service>apipe>sap# shutdown
A:ALU-41>config>service>apipe>sap# exit
A:ALU-41>config>service>apipe# no sap 1/1/1.1:0/32
A:ALU-41>config>service>apipe# spoke-sdp 1:4
A:ALU-41>config>service>apipe>spoke-sdp# shutdown
A:ALU-41>config>service>apipe>spoke-sdp# exit
A:ALU-41>config>service>apipe# no spoke-sdp 1:4
A:ALU-41>config>service>apipe# shutdown
A:ALU-41>config>service>apipe# exit
A:ALU-41>config>service# no apipe 5
```

VLL Services Command Reference

Command Hierarchies

- VLL Services Configuration Commands
  - → Apipe Service Configuration Commands
  - → Apipe Service with SAP Aggregation Configuration Commands
  - → Cpipe Service Configuration Commands
  - → Epipe Service Configuration Commands
  - → Fpipe Service Configuration Commands
  - → Hpipe Service Configuration Commands
  - → Ipipe Service Configuration Commands
- Show Commands
- Clear Commands

VLL Services Configuration Commands

Apipe Service Configuration Commands

```
config
    — service
        — apipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-type {atm-vcc | atm-vpc |
            atm-cell}] [vc-switching]
        — no apipe service-id
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — description description-string
                — no description
                — revert-time [revert-time | infinite]
                — no revert-time
                — [no] standby-signaling-master
                — [no] standby-signaling-slave
            — sap sap-id [create]
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — atm
                    — egress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — ingress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — oam
                        — [no] alarm-cells
                — [no] collect-stats
                — description description-string
                — no description
                — egress
                    — qos policy-id
                    — no qos
                — ingress
                    — qos policy-id
                    — no qos
                — [no] shutdown
            — service-mtu octets
            — no service-mtu
            — service-name service-name
            — no service-name
            — [no] shutdown
            — spoke-sdp sdp-id:vc-id [create] [no-endpoint] (see Note:)
            — spoke-sdp sdp-id:vc-id [create] endpoint endpoint-name
            — no spoke-sdp sdp-id:vc-id
                — cell-concatenation
                    — [no] clp-change
```

```
                    — max-cells cell-count
                    — no max-cells [cell-count]
                    — max-delay delay-time
                    — no max-delay [delay-time]
                — [no] control-word
                — [no] egress
                    — vc-label egress-vc-label
                    — no vc-label [egress-vc-label]
                — [no] ingress
                    — vc-label ingress-vc-label
                    — no vc-label [ingress-vc-label]
                — precedence [precedence-value | primary]
                — no precedence
                — [no] shutdown
```

Note: The spoke-sdp configuration does not apply to ATM SAP-to-SAP configuration (local service). It only applies to SAP-to-SDP configuration (distributed service).

Apipe Service with SAP Aggregation Configuration Commands

```
config
    — service
        — apipe service-id [customer customer-id] [vpn vpn-id] [vc-type {atm-vcc | atm-vpc}]
            [vc-switching] [create]
        — no apipe service-id
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — description description-string
                — no description
                — revert-time [revert-time | infinite]
                — no revert-time
                — [no] standby-signaling-master
                — [no] standby-signaling-slave
            — sap-aggregation-group group-id [create]
            — no sap-aggregation-group group-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — [no] collect-stats
                — description description-string
                — no description
                — egress
                    — qos policy-id
                    — no qos
                — ingress
                    — qos policy-id
                    — no qos
                — [no] shutdown
            — [no] shutdown
```

```
— sap sap-id [sap-aggregation-group group-id] [create]
— no sap sap-id
    — atm
        — vcid-translation vpi-vci
        — no vcid-translation
        — egress
            — traffic-desc traffic-desc-profile-id
            — no traffic-desc
        — ingress
            — traffic-desc traffic-desc-profile-id
            — no traffic-desc
        — oam
            — [no] alarm-cells
    — description description-string
    — no description
— [no] shutdown
— sap sap-id  [sap-aggregation-group group-id] [create] (additional SAP group
    members up to a total of 16 for each Apipe)
— service-mtu octets
— no service-mtu
— service-name service-name
— no service-name
— spoke-sdp sdp-id:vc-id [create] [no-endpoint] (see  Note:)
— spoke-sdp sdp-id:vc-id [create] endpoint endpoint-name
— no spoke-sdp sdp-id:vc-id
    — cell-concatenation
        — [no] clp-change
        — max-cells cell-count
        — no max-cells [cell-count]
        — max-delay delay-time
        — no max-delay [delay-time]
    — [no] control-word
    — [no] egress
        — vc-label egress-vc-label
        — no vc-label [egress-vc-label]
    — [no] ingress
        — vc-label ingress-vc-label
        — no vc-label [ingress-vc-label]
    — precedence [precedence-value | primary]
    — no precedence
    — [no] shutdown
```

Note: The spoke-sdp configuration does not apply to ATM SAP-to-SAP configuration (local service). It only applies to SAP-to-SDP configuration (distributed service).

Cpipe Service Configuration Commands

```
config
    — service
        — [no] cpipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-type {satop-e1 |
          satop-t1 | satop-e3 | satop-t3 | cesopsn | cesopsn-cas}] [vc-switching]
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — description description-string
                — no description
                — revert-time revert-time
                — no revert-time
                — [no] standby-signaling-master
                — [no] standby-signaling-slave
            — sap sap-id [create]
            — [no] sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — cem
                    — asym-delay-control [samples ksamples] [min-repeat minutes]
                    — no asym-delay-control
                    — [no] packet
                        — jitter-buffer value [payload-size size]
                        — payload-size size
                    — [no] report-alarm [stray] [malformed] [pktloss] [overrun]
                          [underrun] [rpktloss] [rfault] [rrdi]
                    — [no] rtp-header
                — [no] collect-stats
                — description description-string
                — no description
                — egress
                    — qos policy-id
                    — no qos
                — ingress
                    — qos policy-id
                    — no qos
                — [no] shutdown
            — service-mtu octets
            — no service-mtu
            — service-name service-name
            — no service-name
            — [no] shutdown
            — spoke-sdp sdp-id:vc-id [create] [no-endpoint] (see Note:)
            — spoke-sdp sdp-id:vc-id [create] endpoint endpoint-name [icb]
            — no spoke-sdp sdp-id:vc-id
                — [no] control-word
                — [no] egress
                    — [no] vc-label egress-vc-label
                — [no] ingress
                    — [no] vc-label ingress-vc-label
                — precedence [precedence-value | primary]
                — no precedence
```

— [no] shutdown

Note: The spoke-sdp configuration does not apply to TDM SAP-to-SAP configuration (local service). It only applies to SAP-to-SDP configuration (distributed service).

Epipe Service Configuration Commands

```
config
    — service
        — [no] epipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-switching]
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — description description-string
                — no description
                — revert-time [revert-time | infinite]
                — no revert-time
                — [no] standby-signaling-master
                — [no] standby-signaling-slave
            — sap sap-id [create]
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — cem
                    — [no] packet
                        — jitter-buffer value [payload-size size]
                        — payload-size size
                    — local-ecid value
                    — no local-ecid
                    — remote-ecid value
                    — no remote-ecid
                    — remote-mac ieee-mac-addr
                    — no remote-mac
                    — [no] report-alarm [stray] [malformed] [pktloss] [overrun]
                            [underrun] [rpktloss] [rfault] [rrdi]
                    — [no] rtp-header
                — [no] collect-stats
                — description description-string
                — no description
                — egress
                    — agg-rate-limit agg-rate [cir cir-rate]
                    — no agg-rate-limit
                    — [no] qinq-mark-top-only
                    — qos policy-id
                    — no qos
                    — scheduler-mode {4-priority | 16-priority}
                    — [no] shaper-group shaper-group-name
                — eth-cfm
                    — [no] hold-mep-up-on-failure
```

- **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**
  - **description** *description-string*
  - **no description**
  - [**no**] **dual-ended-loss-test-enable**
    - **alarm-threshold** *percentage*
    - **no alarm-threshold**
    - **alarm-clear-threshold** *percentage*
    - **no alarm-clear-threshold**
  - [**no**] **eth-test-enable**
    - **bit-error-threshold** *bit-errors*
    - [**no**] **test-pattern** {**all-zeros** | **all-ones**} [**crc-enable**]
  - **low-priority-defect** {**allDef** | **macRemErrXcon** | **remErrXcon** | **errXcon** | **xcon** | **noXcon**}
  - **one-way-delay-threshold** *seconds*
  - [**no**] **shutdown**
- **ethernet**
  - [**no**] **llf**
- **ingress**
  - **agg-rate-limit** *agg-rate* [**cir** *cir-rate*]
  - **no agg-rate-limit**
  - **filter ip** *ip-filter-id*
  - **no filter** [**ip** *ip-filter-id*]
  - **match-qinq-dot1p** {**top** | **bottom**}
  - **no match-qinq-dot1p**
  - **qos** *policy-id*
  - **no qos**
  - **scheduler-mode** {**4-priority** | **16-priority**}
  - [**no**] **shaper-group** *shaper-group-name*
- **mw**
  - **compression** *source-mac destination-mac* [**rtp**]
  - **no compression**

- **service-mtu** *octets*
- **no service-mtu**
- **service-name** *service-name*
- **no service-name**
- [**no**] **shutdown**
- **spoke-sdp** *sdp-id*:*vc-id* [**vc-type** {**ether** | **vlan**}] [**create**] [**no-endpoint**]
- **spoke-sdp** *sdp-id*:*vc-id* [**vc-type** {**ether** | **vlan**}] [**create**] **endpoint** *endpoint-name* [**icb**]
- **no spoke-sdp** *sdp-id*:*vc-id*
  - [**no**] **control-word**
  - **egress**

```
                    — vc-label egress-vc-label
                    — 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] ccm-enable
                        — ccm-ltm-priority priority
                        — no ccm-ltm-priority
                        — description description-string
                        — no description
                        — low-priority-defect {allDef | macRemErrXcon |
                              remErrXcon | errXcon | xcon | noXcon}
                        — [no] shutdown
                — ingress
                    — vc-label ingress-vc-label
                    — no vc-label [ingress-vc-label]
                — [no] shutdown
                — precedence [precedence-value | primary]
                — no precedence
                — [no] pw-status-signaling
                — vlan-vc-tag 0..4094
                — no vlan-vc-tag [0..4094]
```

Note:

- The spoke-sdp configuration does not apply to Ethernet SAP-to-SAP configuration (local service). It only applies to SAP-to-SDP configuration (distributed service).
- For information on configuring ETH-CFM on Epipes, refer to the 7705 SAR OS OAM and Diagnostics Guide.

Fpipe Service Configuration Commands

```
config
    — service
        — fpipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-type {fr-dlci}]
              [vc-switching]
        — no fpipe service-id
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — description description-string
                — no description
                — revert-time [revert-time | infinite]
                — no revert-time
                — [no] standby-signaling-master
                — [no] standby-signaling-slave
            — sap sap-id [create]
            — no sap sap-id
                — accounting-policy acct-policy-id
```

```
                    — no accounting-policy
                    — [no] collect-stats
                    — description description-string
                    — no description
                    — egress
                        — qos policy-id
                        — no qos
                    — frame-relay
                    — ingress
                        — qos policy-id
                        — no qos
                    — [no] shutdown
                — service-mtu octets
                — no service-mtu
                — service-name service-name
                — no service-name
                — [no] shutdown
                — spoke-sdp sdp-id:vc-id [create] [no-endpoint]
                — spoke-sdp sdp-id:vc-id [create] endpoint endpoint-name
                — no spoke-sdp sdp-id:vc-id
                    — [no] control-word
                    — egress
                        — vc-label egress-vc-label
                        — no vc-label [egress-vc-label]
                    — ingress
                        — vc-label ingress-vc-label
                        — no vc-label [ingress-vc-label]
                    — precedence [precedence-value | primary]
                    — no precedence
                    — [no] shutdown
```

Hpipe Service Configuration Commands

```
config
    — service
        — hpipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-type {hdlc}]
          [vc-switching]
        — no hpipe service-id
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — description description-string
                — no description
                — revert-time [revert-time | infinite]
                — no revert-time
                — [no] standby-signaling-master
                — [no] standby-signaling-slave
            — sap sap-id [create]
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
```

```
                    — [no] collect-stats
                    — description description-string
                    — no description
                    — egress
                        — qos policy-id
                        — no qos
                    — ingress
                        — qos policy-id
                        — no qos
                    — [no] shutdown
                — service-mtu octets
                — no service-mtu
                — service-name service-name
                — no service-name
                — [no] shutdown
                — spoke-sdp sdp-id:vc-id [create] [no-endpoint]
                — spoke-sdp sdp-id:vc-id [create] endpoint endpoint-name
                — no spoke-sdp sdp-id:vc-id
                    — [no] control-word
                    — egress
                        — vc-label egress-vc-label
                        — no vc-label [egress-vc-label]
                    — ingress
                        — vc-label ingress-vc-label
                        — no vc-label [ingress-vc-label]
                    — precedence [precedence-value | primary]
                    — no precedence
                    — [no] shutdown
```

Ipipe Service Configuration Commands

```
config
    — service
        — ipipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-switching]
        — [no] ipipe service-id
            — ce-address-discovery
            — no ce-address-discovery
            — description description-string
            — no description
            — [no] endpoint endpoint-name
                — description description-string
                — no description
                — revert-time [revert-time | infinite]
                — no revert-time
                — [no] standby-signaling-master
                — [no] standby-signaling-slave
            — sap sap-id [create]
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — ce-address ip-address
```

```
— no ce-address
— no collect-stats
— description description-string
— no description
— egress
    — agg-rate-limit agg-rate [cir cir-rate]
    — no agg-rate-limit
    — [no] qinq-mark-top-only
    — qos policy-id
    — no qos
    — scheduler-mode {4-priority | 16-priority}
    — [no] shaper-group shaper-group-name
— frame-relay
— ingress
    — agg-rate-limit agg-rate [cir cir-rate]
    — no agg-rate-limit
    — filter ip ip-filter-id
    — no filter [ip ip-filter-id]
    — match-qinq-dot1p {top | bottom}
    — no match-qinq-dot1p
    — qos policy-id
    — no qos
    — scheduler-mode {4-priority | 16-priority}
    — [no] shaper-group shaper-group-name
— [no] ipcp
    — [no] assign-peer-ce-addr
    — [no] dns ip-address-1 [secondary ip-address-2]
— [no] mac ieee-address
— mac-refresh refresh-interval
— no mac-refresh
— [no] shutdown
— service-mtu octets
— no service-mtu
— service-name service-name
— no service-name
— [no] shutdown
— spoke-sdp sdp-id:vc-id [create] [no-endpoint]
— spoke-sdp sdp-id:vc-id [create] endpoint endpoint-name
— no spoke-sdp sdp-id:vc-id
    — ce-address ip-address
    — no ce-address
    — [no] control-word
    — egress
        — vc-label egress-vc-label
        — no vc-label [egress-vc-label]
    — ingress
        — vc-label ingress-vc-label
        — no vc-label [ingress-vc-label]
    — [no] shutdown
    — precedence [precedence-value | primary]
    — no precedence
```

Show Commands

```
show
    — eth-cfm (refer to the 7705 SAR OS OAM and Diagnostics Guide for show>eth-cfm commands)
    — service
        — egress-label start-label [end-label]
        — id service-id
            — all
            — base
            — endpoint endpoint-name
            — labels
            — sap [sap-id] [atm | base | detail | qos | sap-stats | stats]
            — sap-aggregation-group group-id [base | detail | group-stats | stats]
            — sdp [sdp-id | far-end ip-address] [detail | keep-alive-history]
        — ingress-label start-label [end-label]
        — sap-using [sap sap-id]
        — sap-using [sap-aggregation-group group-id]
        — sap-using description
        — sap-using [ingress | egress] atm-td-profile td-profile-id
        — sap-using [ingress | egress] qos-policy qos-policy-id
        — sap-using [ingress | egress] scheduler-mode {4-priority | 16-priority}
        — scada bridge-id
        — sdp-using [sdp-id[:vc-id] | far-end ip-address]
        — service-using [epipe] [ies] [vpls] [vprn] [apipe] [fpipe] [ipipe] [cpipe] [hpipe] [sdp-id
            sdp-id ] customer-id customer-id
```

Clear Commands

```
clear
    — service
        — id service-id
            — arp
            — spoke-sdp sdp-id:vc-id ingress-vc-label
        — statistics
            — id service-id
                — counters
                — spoke-sdp sdp-id:vc-id {all | counters}
            — sap sap-id {all | cem | counters}
            — sap-aggregation-group group-id {all | counters}
            — sdp sdp-id keep-alive
```

Command Descriptions

- VLL Service Configuration Commands
- Show Commands
- Clear Commands

VLL Service Configuration Commands

- Generic Commands
- VLL Global Commands
- VLL SAP Commands
- SAP Aggregation Group Commands
- SAP cem Commands
- SAP QoS and IP Filter Policy Commands
- SAP Microwave Link Commands
- Service Billing Commands
- VLL SDP Commands
- SDP Cell Concatenation Commands
- ATM Commands
- ATM OAM Commands

Generic Commands

description

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

Context config>service>apipe
config>service>apipe>endpoint
config>service>apipe>sap
config>service>apipe>sap-aggregation-group
config>service>cpipe
config>service>cpipe>endpoint
config>service>cpipe>sap
config>service>epipe
config>service>epipe>endpoint
config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>fpipe
config>service>fpipe>endpoint
config>service>fpipe>sap
config>service>fpipe>spoke-sdp
config>service>hpipe
config>service>hpipe>endpoint
config>service>hpipe>sap
config>service>hpipe>spoke-sdp
config>service>ipipe
config>service>ipipe>endpoint
config>service>ipipe>sap
config>service>ipipe>spoke-sdp

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

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

Default No description is associated with the configuration context.

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

shutdown

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

Context config>service>apipe
config>service>apipe>sap
config>service>apipe>spoke-sdp
config>service>apipe>sap-aggregation-group
config>service>cpipe
config>service>cpipe>sap
config>service>cpipe>spoke-sdp
config>service>epipe
config>service>epipe>sap
config>service>epipe>spoke-sdp
config>service>fpipe
config>service>fpipe>sap
config>service>fpipe>spoke-sdp
config>service>hpipe
config>service>hpipe>sap
config>service>hpipe>spoke-sdp
config>service>ipipe
config>service>ipipe>sap
config>service>ipipe>spoke-sdp

Description The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many objects must be shut down before they can be deleted. Many entities must be explicitly enabled using the **no shutdown** command.

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

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 are described in the following Special Cases.

Special Cases **Service Admin State —** bindings to an SDP within the service will be put into the out-of-service state when the service is shut down. While the service is shut down, all customer packets are dropped and counted as discards for billing and debugging purposes.

Service Operational State — a service is considered operational if at least one SAP and one SDP are operational

SDP (global) — when an SDP is shut down at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would normally be transmitted using this SDP binding will be discarded and counted as dropped packets.

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

VLL Global Commands

apipe

Syntax **apipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-type** {**atm-vcc** | **atm-vpc** | **atm-cell**}] [**vc-switching**]
no apipe *service-id*

Context config>service

Description This command configures a point-to-point ATM service. The Apipe service provides a point-to-point Layer 2 VPN connection to a local or remote SAP. An Apipe can connect an ATM endpoint locally (in the same 7705 SAR) or over a PSN to a remote endpoint of the same type.

Apipes support SAP aggregation groups in which multiple VCC SAPs can be bound in a single service.

Parameters *service-id* — uniquely identifies a 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 or name used for every 7705 SAR on which this service is defined.

Values 1 to 2147483647 or *service-name*

create — keyword used to create an Apipe. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

customer-id **—** specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn-id — specifies the VPN ID number that allows you to identify virtual private networks (VPNs) by a VPN identification number. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 to 2147483647

vc-type — specifies a 15-bit value that defines the type of the VC signaled to the peer. Its values are defined in *draft-ietf-pwe3-iana-allocation* and it defines both the signaled VC type as well as the resulting datapath encapsulation over the Apipe.

Values atm-vcc, atm-vpc, atm-cell

Default atm-vcc

The **vc-type** must be set to **atm-vcc** in order to configure an Apipe for N-to-one cell mode where N > 1.

The **vc-type** must be set to **atm-cell** in order to configure an Apipe for ATM virtual trunking.

vc-switching — specifies that pseudowire switching signaling is used for the spoke SDPs configured for this service

cpipe

Syntax [**no**] **cpipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-type** {**satop-e1** | **satop-t1** | **satop-e3** | **satop-t3** | **cesopsn** | **cesopsn-cas**}] [**vc-switching**]

Context config>service

Description This command configures a circuit emulation service using MPLS or GRE encapsulation. The **vc-type** defines the type of unstructured or structured circuit emulation service to be configured. All other parameters (**service-id**, **customer**) have common usage with other service types.

Default no cpipe

Parameters *service-id* — uniquely identifies a 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 or name used for every 7705 SAR on which this service is defined.

Values 1 to 2147483647 or *service-name*

customer-id — specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

create — keyword used to create a Cpipe. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

vpn-id — specifies the VPN ID number that allows you to identify virtual private networks (VPNs) by a VPN identification number. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 to 2147483647

vc-type — specifies a value that defines the type of the VC signaled to the peer. This optional parameter is included when the Cpipe service is created.

Values satop-e1: unstructured E1 circuit emulation service
satop-t1: unstructured DS1 circuit emulation service
satop-e3: unstructured E3 circuit emulation service
satop-t3: unstructured DS3 circuit emulation service
cesopsn: basic structured n × 64 kb/s circuit emulation service
cesopsn-cas: structured n × 64 kb/s circuit emulation service with signaling

Default cesopsn

vc-switching — specifies that pseudowire switching signaling is used for the spoke SDPs configured for this service

epipe

Syntax [**no**] **epipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-switching**]

Context config>service

Description This command configures a point-to-point Ethernet service. An Epipe connects two endpoints defined as SAPs. Both SAPs are defined on separate routers (7705 SAR routers or other Alcatel-Lucent service routers) connected over the service provider network. When the endpoint SAPs are separated by the service provider network, the far-end SAP is generalized into an SDP. This SDP describes a destination 7705 SAR and the encapsulation method used to reach it.

No MAC learning or filtering is provided (or needed) on an Epipe.

When a service is created, the **customer** keyword and *customer-id* must be specified, which 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, Epipe services do not exist until they are explicitly created with this command.

The **no** form of this command deletes the Epipe service instance with the specified *service-id*. The service cannot be deleted until the service has been shut down.

Parameters *service-id* — uniquely identifies a 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 or name used for every 7705 SAR on which this service is defined.

Values 1 to 2147483647 or *service-name*

customer-id — specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

create — keyword used to create an Epipe. The create keyword requirement can be enabled/disabled in the **environment>create** context

vpn-id — specifies the VPN ID number that allows you to identify virtual private networks (VPNs) by a VPN ID. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 to 2147483647

vc-switching — specifies that pseudowire switching signaling is used for the spoke SDPs configured for this service

fpipe

Syntax **fpipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-type** {**fr-dlci**}] [**vc-switching**]
no fpipe *service-id*

Context config>service

Description This command configures a point-to-point FR service. The Fpipe service provides a point-to-point Layer 2 VPN connection to a local or remote SAP. An Fpipe connects only FR endpoints. Endpoints can be connected locally (in the same 7705 SAR) or remotely over a PSN.

Parameters *service-id* — uniquely identifies a 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 or name used for every 7705 SAR on which this service is defined.

Values 1 to 2147483647 or *service-name*

create — keyword used to create an Fpipe. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

customer-id **—** specifies the customer ID number to be associated with the service. This parameter is required on service creation and is optional for service editing or deleting.

Values 1 to 2147483647

vpn-id — specifies the VPN ID number that allows you to identify virtual private networks by a VPN identification number. If this parameter is not specified, the VPN ID uses the service ID number.

Values 1 to 2147483647

vc-type — specifies a 15-bit value that defines the type of the VC signaled to the peer. The value is defined in *draft-ietf-pwe3-iana-allocation* and is used for the signaled VC type and the resulting datapath encapsulation over the Fpipe.

Values fr-dlci

vc-switching — specifies that pseudowire switching signaling is used for the spoke SDPs configured for this service

hpipe

Syntax **hpipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-type** {**hdlc**}] [**vc-switching**]
no hpipe *service-id*

Context config>service

Description This command configures a point-to-point HDLC service.

Parameters *service-id* — uniquely identifies a 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 or name used for every 7705 SAR on which this service is defined.

Values 1 to 2147483647 or *service-name*

create — keyword used to create an Hpipe. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

customer-id **—** specifies the customer ID number to be associated with the service. This parameter is required on service creation and is optional for service editing or deleting.

Values 1 to 2147483647

vpn-id — specifies the VPN ID number that allows you to identify virtual private networks by a VPN identification number. If this parameter is not specified, the VPN ID uses the service ID number.

Values 1 to 2147483647

vc-type — specifies a 15-bit value that defines the type of the VC signaled to the peer. The value is defined in *draft-ietf-pwe3-iana-allocation* and is used for the signaled VC type and the resulting datapath encapsulation over the Hpipe.

Values hdlc

vc-switching — specifies that pseudowire switching signaling is used for the spoke SDPs configured for this service

ipipe

Syntax **ipipe** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-switching**]
[**no**] **ipipe** *service-id*

Context config>service

Description This command configures an IP interworking service. An Ipipe can be configured as a SAP-to-SAP or SAP-to-SDP service. An Ipipe can connect the following types of SAPs over an MPLS or IP network:

- Ethernet SAP
- LAG SAP
- PPP/MLPPP SAP
- Frame Relay SAP
- cHDLC SAP

Parameters *service-id* — uniquely identifies a 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 or name used for every 7705 SAR on which this service is defined.

Values 1 to 2147483647 or *service-name*

customer-id — specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

create — keyword used to create an Ipipe. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

vpn-id — specifies the VPN ID number that allows you to identify virtual private networks (VPNs) by a VPN ID. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 to 2147483647

vc-switching — specifies that pseudowire switching signaling is used for the spoke SDPs configured for this service

endpoint

Syntax [**no**] **endpoint** *endpoint-name*

Context config>service>apipe
config>service>cpipe
config>service>epipe
config>service>fpipe
config>service>hpipe
config>service>ipipe

Description This command provides access to the service endpoint context.

Parameters *endpoint-name* — specifies an endpoint name (up to 32 alphanumeric characters)

ce-address-discovery

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

Context config>service>ipipe

Description This command enables CE address discovery for an Ipipe service.

Default no ce-address-discovery

revert-time

Syntax **revert-time** [*revert-time* | **infinite**]
no revert-time

Context config>service>apipe>endpoint
config>service>cpipe>endpoint
config>service>epipe>endpoint
config>service>fpipe>endpoint
config>service>hpipe>endpoint
config>service>ipipe>endpoint

Description This command configures the time to wait before reverting back to the primary spoke SDP defined on this service endpoint, after having switched over to a backup spoke SDP after a failure of the primary spoke SDP.



Note: The **infinite** option (for non-revertive behavior) does not apply to Cpipes.

Parameters *revert-time* — specifies the time, in seconds, to wait before reverting to the primary SDP

Values 0 to 600

infinite — causes the endpoint to be non-revertive

standby-signaling-master

Syntax [**no**] **standby-signaling-master**

Context config>service>apipe>endpoint
config>service>cpipe>endpoint
config>service>epipe>endpoint
config>service>fpipe>endpoint
config>service>hpipe>endpoint
config>service>ipipe>endpoint

Description This command sends the pseudowire standby bit (value 0x00000020) to the targeted LDP (T-LDP) peer whenever a spoke SDP of the endpoint is selected as a standby (see precedence). This bit informs the far end that the pseudowire is not currently active.

Default no standby-signaling-master

standby-signaling-slave

Syntax [**no**] **standby-signaling-slave**

Context config>service>apipe>endpoint
config>service>cpipe>endpoint
config>service>epipe>endpoint
config>service>fpipe>endpoint
config>service>hpipe>endpoint
config>service>ipipe>endpoint

Description This command blocks the spoke SDP from transmitting data when the pseudowire standby bit (value 0x00000020) is received from a targeted LDP (T-LDP) peer. In order to have standby-signaling-slave working properly end-to-end, standby-signaling-master must be enabled on the ingress LER (see standby-signaling-master).

If this command is enabled, the s**how service id** *id* **all** command output shows the Flags field with a value of "StandbySigSlaveTxDown" and the Peer Pw Bits field with a value of "PwFwdingStandby", indicating that transmission is blocked but the spoke SDP is still up.

The **no** form of the command disables the blocking of traffic in the transmit direction, and data received via the associated SAP or service continues to be transmitted.

Default no standby-signaling-slave

service-mtu

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

Context config>service>apipe
config>service>cpipe
config>service>epipe
config>service>fpipe
config>service>hpipe
config>service>ipipe

Description This command configures the service payload (MTU), in octets, for the service. This MTU value overrides the service-type default MTU.

The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP binding's operational state within the service.

The service MTU and a SAP's service delineation encapsulation overhead (4 bytes for a dot1q tag or 8 bytes for qinq tags) 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, the SAP will be placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP will be able to transition to the operative state.

When binding an SDP to a service, the service MTU is compared to the path MTU associated with the SDP. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced due to the MTU capabilities discovered by the tunneling mechanism of the SDP or the egress interface MTU capabilities based on the next hop in the tunnel path. If the service MTU is larger than the path MTU, the SDP binding for the service will be placed in an inoperative state. If the service MTU is equal to or less than the path MTU, then the SDP binding will be placed in an operational state.

In the event that a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, then all associated SAP and SDP binding operational states are automatically re-evaluated.

The **no** form of this command returns the default service-mtu for the indicated service type to the default value.

Parameters *octets* — specifies the size of the MTU, expressed as a decimal integer

Values 1 to 1514 (apipe, cpipe, fpipe, and hpipe)
1 to 9670 (epipe and ipipe)

Default The default values depend on the PW type:

- apipe: 1508
- cpipe: 1514
- epipe: 1514
- fpipe: 1508
- hpipe: 1514
- ipipe: 1500

Table 38 displays MTU values for specific VC types.

Table 38: Service MTU Values

VC-Type	Example of Service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
Ethernet (with preserved qinq)	1522	1508
VLAN (dot1p transparent to MTU value)	1514	1500

service-name

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

Context config>service>apipe
config>service>cpipe
config>service>epipe
config>service>fpipe
config>service>hpipe
config>service>ipipe

Description This command configures a service name that can be used for reference in configuration and show commands.

Parameters *service-name* — up to 64 characters

VLL SAP Commands

sap

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

Context config>service>apipe
config>service>cpipe
config>service>epipe
config>service>fpipe
config>service>hpipe
config>service>ipipe

Description This command creates a SAP within a service. Each SAP must be unique.

All SAPs must be explicitly created with the **create** keyword. If no SAPs are created within a service or an IP interface, a SAP will not exist on that object.

To edit SAP parameters, enter an existing SAP without the **create** keyword.

A SAP can only be associated with a single service. The SAP is owned by the service in which it was created. A SAP can only be defined on a port that has been configured as an access port in the **config>port** *port-id* context using the **mode access** command, or on a Surveillance, Control, and Data Acquisition Support (SCADA) bridge that has been configured on an Integrated Services card using the **config>scada** command. The Integrated Services card is a resource card that does not have any ports; as well, it supports an access functionality only.

Fractional TDM ports are always access ports. Refer to the 7705 SAR OS Interface Configuration Guide.

If a port or SCADA bridge is shut down, all SAPs on that port or SCADA bridge become operationally down. When a service is shut down, SAPs for the service are displayed as operationally down and all traffic traversing the service is discarded. The operational state of a SAP is relative to the operational state of the port or SCADA bridge on which the SAP is defined.

The following SAP types are supported:

- ATM VPI/VCI on an ATM port for **vc-type atm-vcc**
- ATM VPI on an ATM port for **vc-type atm-vpc**
- ATM virtual trunk on an ATM port for **vc-type atm-cell**
- Ethernet-Ethernet
- SAToP
- CESoPSN (with and without CAS)
- PPP IPCP encapsulation of an IPv4 packet for Ipipe service (RFC 1332)
- MLPPP bundle

- LAG
- Ethernet SAPs supporting null, dot1q, and qinq for Ipipe service
- FR DLCI
- HDLC
- cHDLC (supported only for Ipipe services)

The **no** form of this command deletes the SAP with the specified port or SCADA bridge. When a SAP is deleted, all configuration parameters for the SAP will also be deleted.

SAPs that are configured with this command cannot be bound to a SAP aggregation group. See SAP Aggregation Group Commands for more information about the sap *sap-id* sap-aggregation-group *group-id* command, which is used to support N-to-one cell mode where N > 1.

Default no sap

Parameters *sap-id* — specifies the physical port or SCADA bridge identifier portion of the SAP definition

The *sap-id* can be configured in one of the formats described in Table 39. The range of values for the parameters follow the table.

Table 39: SAP ID Configurations

Type	Syntax	Example
port-id	*slot*/*mda*/*port*[.*channel*]	1/1/5
SCADA bridge	*slot/mda/bridge-id.branch-id*	1/5/16.10
null	[*port-id* | *bundle-id* | lag-*id* | aps-*id* | mw-link-*id*]	*port-id*: 1/1/3 *bundle-id*: bundle-ppp-1/1.1 lag-*id*: lag-1 aps-*id*: aps-1 mw-link-*id*: mw-link-1
dot1q	[*port-id* | lag-*id* | aps-*id* | mw-link-*id*]:qtag1	*port-id*:qtag1: 1/1/3:100 lag-*id*: lag-1:10 aps-*id*: aps-1 mw-link-*id*: mw-link-1
qinq	[*port-id* | lag-*id*]:qtag1.qtag2	*port-id*:qtag1.qtag2: 1/1/3:100.30 lag-*id*: lag-1:10.10

Table 39: SAP ID Configurations (Continued)

Type	Syntax	Example
atm	[*port-id* | aps-*id*][:vpi/vci | vpi | vpi1.vpi2] [1]	port-id: 1/1/1 or 1/1/1.1 (for T1/E1 channelized ports) aps-*id*: aps-1 vpi/vci: 16/26 vpi: 16 vpi1.vpi2: 16.22
lag	lag-*id*	lag-2
frame	[*port-id*| aps-*id*]:*dlci*	1/1/1 aps-*id*: aps-1 *dlci*: 16
frame relay	[*port-id*]:*dlci*	1/1/1 *dlci*: 16
cisco-hdlc	*slot/mda/port.channel*	1/1/1.3
cem	*slot/mda/port.channel*	1/1/1.3
ima-grp	*bundle-id*[:vpi/vci | vpi | vpi1.vpi2]	1/1/3.1
ipcp	*slot/mda/port.channel*	1/2/2.4
hdlc	*slot/mda/port.channel*	1/1/3.1
lag-id	lag-*id*	lag-1
mw-link-id	mw-link-*id*	mw-link-1
aps-id	aps-*group-id*[.channel]	aps-1
bundle-id	bundle-[ima | ppp]-*slot/mda.bundle-num*	bundle-ima-1/1.1
tunnel-id	tunnel-<*id*>.[private | public]:<*tag*>	tunnel-1.private:1

Note:

1. For Apipes in virtual trunking mode, vpi/vci, vpi, and vpi1.vpi2 are omitted.

Values *sap-id*:

null	[*port-id* | *bundle-id* | *lag-id* | *aps-id* | *mw-link-id*]
dot1q	[*port-id* | *lag-id* | *aps-id* | *mw-link-id*]:*qtag1*
qinq	[*port-id* | *lag-id*]:*qtag1.qtag2*
atm	[*port-id* | *aps-id*][:vpi/vci |vpi | vpi1.vpi2]
frame	[*port-id* | *aps-id*]:*dlci*

cisco-hdlc	*slot/mda/port.channel*
cem	*slot/mda/port.channel*
ipcp	*slot/mda/port.channel*
ima-grp	*bundle-id*[:vpi/vci | vpi | vpi1.vpi2]
hdlc	*slot/mda/port.channel*
port-id	*slot/mda/port*[*.channel*]
SCADA bridge	*slot/mda/bridge-id.branch-id* *bridge-id* 1 to 16 *branch-id* 1 to 32
bundle-id	bundle-*type-slot/mda.bundle-num* bundle keyword *type* ima, ppp *bundle-num* 1 to 32
aps-id	aps-*group-id*[*.channel*] aps keyword *group-id* 1 to 24
mw-link-id	mw-link-*id* *id* 1 to 24
lag-id	lag-*id* lag keyword *id* 1 to 32
qtag1	*, 0 to 4094
qtag2	*, 0 to 4094
vpi	NNI 0 to 4095 UNI 0 to 255
vci	1, 2, 5 to 65535
dlci	16 to 1022
tunnel-id	tunnel-*id*.[private | public]:*tag* tunnel keyword *id* 1 to 16 ("1" is the only valid value) *tag* 0 to 4094

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

If the card in the slot has an adapter card installed, the *port-id* must be in the slot-number/MDA-number/port-number format. For example, 1/2/3 specifies port 3 on MDA 2 in slot 1.

The *port-id* must reference a valid port type. When the *port-id* parameter represents 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.

bundle-id — specifies the multilink bundle identifier. The **bundle** keyword must be entered at the beginning of the parameter. The command syntax must be configured as follows:

bundle-id: **bundle**-*type-slot*/*mda*.*bundle-num*

type: ima, ppp

bundle-num: 1 to 32

For example:

```
*A:ALU-12>config# port bundle-ppp-5/1.1
*A:ALU-12>config>port# multilink-bundle
```

qtag1, qtag2 — specifies the encapsulation value used to identify the SAP on the port or sub-port. For dot1q encapsulation, only *qtag1* is used; for qinq encapsulation, both *qtag1* and *qtag2* are used. If *qtag1* or *qtag2* is not specifically defined, the value 0 is used. The "*" value represents all *qtag* values between 0 and 4094 that are not specifically defined within another SAP context under the same port. In addition, the following *qtag1.qtag2* values are invalid options:

- *.*qtag2*
- *.0
- 0.*qtag2*

Values *qtag1*: *, 0 to 4094

qtag2: *, 0 to 4094

The values depend on the encapsulation type configured for the interface. Table 40 describes the allowed values for the port and encapsulation types.

Table 40: Port and Encapsulation Values

Port Type	Encap-Type	Allowed Values	Comments
Ethernet	Null	—	The SAP is identified by the port.
Ethernet	Dot1q	*, 0 to 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, and that a * qtag1 value accepts any VLAN ID that is not specifically configured on the port. [1]
Ethernet	QinQ	*, 0 to 4094	The SAP is identified by the two 802.1Q tags on the port. Note that a 0 qtag1 or qtag 2 value also accepts untagged packets on the qinq port, and that a * qtag1 or qtag2 value accepts any VLAN ID that is not specifically configured on the port. [1]

Note:

1. Traffic matching the * qtag value uses VLAN 4095 internally.

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

mac

Syntax **[no] mac** *ieee-address*

Context config>service>ipipe>sap

Description This command assigns a specific MAC address to an Ipipe Ethernet SAP.

The **no** form of this command returns the MAC address of the SAP to the default value.

Default The default is the physical MAC address associated with the Ethernet interface where the SAP is configured.

Parameters *ieee-address* — specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers

mac-refresh

Syntax **mac-refresh** *refresh-interval*
no mac-refresh

Context config>service>ipipe>sap

Description This command specifies the interval between ARP requests sent on an Ipipe Ethernet SAP. When the SAP is first enabled, an ARP request will be sent to the attached CE device and the received MAC address will be used in addressing unicast traffic to the CE. Although this MAC address will not expire while the Ipipe SAP is enabled and operational, it is verified by sending periodic ARP requests at the specified interval.

The **no** form of this command restores **mac-refresh** to the default value.

Default 14400

Parameters *refresh-interval* — specifies the interval, in seconds, between ARP requests sent on an Ipipe Ethernet SAP

Values 0 to 65535

ipcp

Syntax **[no] ipcp**

Context config>service>ipipe>sap

Description This command enables the context to configure IPCP. Within this context, IPCP extensions can be configured to define the remote IP address and DNS IP address to be signaled via IPCP on the associated PPP interface.

This command is only applicable if the associated SAP is a PPP/MLPPP interface.

assign-peer-ce-addr

Syntax [**no**] **assign-peer-ce-addr**

Context config>service>ipipe>sap>ipcp

Description This command assigns the IP address, defined by the **config>service>ipipe>sap>ce-address** command, to the far end of the associated PPP/MLPPP link via IPCP extensions. This command is only applicable if the associated SAP or port is a PPP/MLPPP interface with an IPCP encapsulation.

Default no assign-peer-ce-addr

dns

Syntax [**no**] **dns** *ip-address-1* [**secondary** *ip-address-2*]

Context config>service>ipipe>sap>ipcp

Description This command defines the dns address(es) to be assigned to the far end of the associated PPP/MLPPP link via IPCP extensions. This command is only applicable if the associated SAP or port is a PPP/MLPPP interface with an IPCP encapsulation.

Default no dns

Parameters *ip-address-1* — specifies a unicast IPv4 address for the primary DNS server to be signaled to the far end of the associated PPP/MLPPP link via IPCP extensions

ip-address-2 — specifies a unicast IPv4 address for the secondary DNS server to be signaled to the far end of the associated PPP/MLPPP link via IPCP extensions

frame-relay

Syntax **frame-relay**

Context config>service>fpipe>sap
config>service>ipipe>sap

Description Use this command to configure frame relay properties for the SAP.

ethernet

Syntax **ethernet**

Context config>service>epipe>sap

Description Use this command to configure Ethernet properties for the SAP.

llf

Syntax [**no**] **llf**

Context config>service>epipe>sap>ethernet

Description This command enables Link Loss Forwarding (LLF) on an Ethernet port. LLF can only be enabled on Ethernet ports configured for null encapsulation.

LLF provides an end-to-end OAM fault notification for Ethernet VLL service. When LLF is enabled and there is a local fault on the pseudowire or service, or a remote fault on the SAP or pseudowire, the Ethernet port is brought down. Using label withdrawal or T-LDP status bits, LLF signals to connected equipment that the VLL is down. LLF stops signaling when the fault disappears.

The **no** form of the command disables LLF.

SAP Aggregation Group Commands

sap-aggregation-group

Syntax **sap-aggregation-group** *group-id* [**create**]
no sap-aggregation-group *group-id*

Context config>service>apipe

Description This command configures a SAP aggregation group on an Apipe. A SAP aggregation group *group-id* is a unique identifier on each node.

To configure an Apipe with a SAP aggregation group, the SAP aggregation group must be configured before any SAPs are bound to the Apipe. SAPs that are to be bound to the same ATM PW payload must be configured with the same SAP aggregation group *group-id*.

All common access configuration parameters, such as accounting, statistics, and packet layer QoS profile can be configured under the sap-aggregation-group.

Each sap-aggregation-group SAP has its own ingress and egress Layer 2 traffic descriptors. These descriptors are used for policing at ingress and shaping and scheduling priority at egress.

Default n/a

Parameters *group-id* — the name identifying the aggregation group

Values 1 to 32 alphanumeric characters

sap

Syntax **sap** *sap-id* [**sap-aggregation-group** *group-id*] [**create**]
no sap *sap-id* [**sap-aggregation-group** *group-id*]

Context config>service>apipe

Description This command configures a SAP and associates the SAP as a member of a specified SAP aggregation group. The aggregation group must be configured before the SAP. The SAP must be associated with the aggregation group at the time of configuration. See sap-aggregation-group for more information.

Up to 16 SAPs can be bound to a SAP aggregation group.

The following restrictions apply to the SAP.

- An Apipe can either have SAPs that are associated with a SAP aggregation group or regular SAPs (that is, SAPs that are not associated with an aggregation group). An Apipe cannot have both types of SAPs.

- The *sap-id* specifies the physical port identifier of the SAP that is to be bound to the aggregation group. The port must be preconfigured for ATM VPI/VCI service with the syntax described in Table 41.

Table 41: SAP ID Preconfiguration for SAP Aggregation Groups

Type	Syntax	Example
atm	*slot/mda/port*[.*channel*]:*vpi/vci*	slot/mda/port[.channel]: 1/1/1.1 vpi/vci: 16/26

- The SAP can only be defined on a port that has been configured as an access port in the **config>port** *port-id* context using the **mode access** command.
- The Apipe for the SAP must be configured to support **vc-type atm-vcc**.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP are also deleted.

Default no sap

Parameters *sap-id* — see Table 41. See Table 39 for a full list of SAP IDs.

group-id — the name identifying the aggregation group to which the SAP is to be bound. See sap-aggregation-group.

Values 1 to 32 alphanumeric characters

SAP cem Commands

cem

Syntax **cem**

Context config>service>cpipe>sap
config>service>epipe>sap

Description This command configures the circuit emulation service parameters. In the **epipe** context, circuit emulation parameters are configured to encapsulate TDM services using Epipes according to the MEF 8 standard.

This command is blocked for all SAPs except for E1, DS1, E3, DS3,and *n* × 64 kb/s channels configured for encap-type cem.

asym-delay-control

Syntax **asym-delay-control** [**samples** *ksamples*] [**min-repeat** *minutes*]
no asym-delay-control

Context config>service>cpipe>sap>cem

Description This command enables or disables asymmetric delay control on the Cpipe. When enabled, the TDM PW is analyzed and adjusted for asymmetric delay when the service is initially started or restarted. The service is considered to be down during this period.

The optional **min-repeat** keyword sets the analysis to repeat at configured time intervals after the service is up. The service remains up while the repeated analysis runs.

Asymmetric delay control must be enabled on both ends of the Cpipe; otherwise, a service parameter mismatch state occurs and the service is brought down.

Default no asym-delay-control

Parameters *ksamples* — specifies the number of packets that are analyzed (k = 1024)

Values 1, 2, 4, 8, 16, 32 (corresponding to 1024, 2048, 4096, 8192, 16 384, 32 768)

Default 16 (16 384)

minutes — the interval at which the asymmetric delay control function will repeat after the service is up, in minutes

Values 1 to 1440; 0 disables the periodic analysis

Default 60

packet

Syntax **[no] packet**

Context config>service>cpipe>sap>cem
config>service>epipe>sap>cem

Description This command enables the context to configure packet parameters on the SAP.

jitter-buffer

Syntax **jitter-buffer** *value* [**payload-size** *size*]

Context config>service>cpipe>sap>cem>packet
config>service>epipe>sap>cem>packet

Description This command defines the size of the receive jitter buffer for the circuit emulation service SAP.

If asym-delay-control is enabled, the jitter buffer size must match on both ends of the Cpipe; otherwise, a service parameter mismatch state occurs and the service is brought down.

Default The default value varies depending on the SAP bandwidth, as follows:

- 5 ms, where SAP bandwidth ≥ 16 DS0s (1024 kb/s)
- 8 ms, where SAP bandwidth is between 5 and 15 DS0s (between 320 and 960 kb/s)
- 16 ms, where SAP bandwidth is between 2 and 4 DS0s (between 128 and 256 kb/s)
- 32 ms, where SAP bandwidth = 1 DS0 (64 kb/s)

Parameters *value* — specifies the size of the receive jitter buffer, expressed in milliseconds. The range of supported values is 1 to 250 ms. Setting the value to 0 sets the default (depends on SAP bandwidth). The buffer size must be set to at least 3 times the value of the packetization delay and no greater than 32 times the value of the packetization delay.

To calculate the size of the buffer (in bytes), multiply the value of the buffer size (in ms) by the SAP TDM bandwidth (in bits per second) and divide by 8. After the initialization of the circuit emulation service, transmission of TDM data begins when the buffer is half full (50%).

size — for convenience, the payload size can be configured at the same time as the jitter buffer. This avoids any configuration errors due to interactions between the jitter buffer and payload size settings. See payload-size.

payload-size

Syntax **payload-size** *size*

Context config>service>cpipe>sap>cem>packet
config>service>epipe>sap>cem>packet

Description This command defines the payload size, in bytes, for one circuit emulation service packet.

Default For SAToP, see Table 20. For CESoPSN without CAS, see Table 21. For CESoPSN with CAS, see Table 22.

Parameters *size* — specifies the payload size (in octets) to be encapsulated in one circuit emulation service packet. The range of supported values is 2 to 1514 bytes. If only one timeslot is being transported, the payload size must be an even number. The packetization delay for the circuit emulation service can be calculated by multiplying the payload size (in octets) by 8 (bits/octet), then dividing by the SAP TDM bandwidth (in bits/second).

For CESoPSN or CESoETH with CAS, the 7705 SAR supports 1 to 8 multiframes. For T1, 1 multiframe equals 24 frames. For E1, 1 multiframe equals 16 frames. The payload size must correspond to *n* x 16 E1 frames or *n* x 24 T1 frames of TDM data, where *n* is an integer between 1 and 8. The restriction for payload sizes of 1 to 8 multiframes also applies to CESoPSN without CAS if the payload is within a T1/E1 channel that has a mix of CESoPSN with and without CAS.

The configured value of the payload size does not need to include the extra bytes for the transport of CAS bits. Note that the configured value of the service MTU size takes the extra CAS bytes into account. See Structured T1/E1 CES with CAS for details.

For CESoPSN or CESoETH, if a port on a 16-port T1/E1 ASAP Adapter card, version 2, or 32-port T1/E1 ASAP Adapter card is configured for DCR, the port timing is associated with the service clock of the Cpipe of channel group 1. For a framed T1 port, there is a restriction on the Cpipe payload size of channel group 1:

- for DCR with a timestamp frequency of 77.76 MHz, the payload size = 2 x I x (number of timeslots), where I = 1 to 20
- for DCR with a timestamp frequency of 19.44 MHz, the payload size = 8 x I x (number of timeslots), where I = 1 to 5

This restriction does not apply to framed E1 ports or unframed T1/E1 ports.

For CESoPSN or CESoETH, the payload size may be specified as the number of bytes to be included in the packet.

For SAToP T1, SAToP E1, SAToP T3 and SAToP E3 circuit emulation services, the size must be specified in bytes (with minimum 64 bytes) and the value must be a multiple of 32. The minimum for SAToP T3 and SAToP E3 is 1024 bytes.

Interactions — The jitter-buffer value must be greater than or equal to twice the payload size to ensure that a frame arrives prior to the start of play-out. Therefore, the payload size may have to be decreased prior to setting the jitter-buffer value. Alternatively, the jitter-buffer value may have to be increased prior to setting the payload-size.

local-ecid

Syntax **local-ecid** *value*
no local-ecid

Context config>service>epipe>sap>cem

Description This command defines the Emulated Circuit Identifier to be used for the (local) source end of the MEF 8 Epipe.

The **no** form of the command removes the ECID value from the service configuration.

Parameters *value* — specifies the source ECID value for the CES. Upon CES packet reception, the ECID in the packet is compared to the local ECID value. If the local and remote ECID values match, the packet payload is used for the TDM circuit. The remote ECID value is inserted into all MEF 8 CES packets that are transmitted.

Values 0 to 1048575

remote-ecid

Syntax **remote-ecid** *value*
no remote-ecid

Context config>service>epipe>sap>cem

Description This command defines the Emulated Circuit Identifiers to be used for the remote (destination) end of the MEF 8 Epipe.

The **no** form of the command removes the ECID value from the service configuration.

Parameters *value* — specifies the destination ECID value for the CES. Upon CES packet reception, the ECID in the packet is compared to the local ECID value specified by the local-ecid command. If the local and remote ECID values match, the packet payload is used for the TDM circuit. The remote ECID value is inserted into all MEF 8 CES packets that are transmitted.

Values 0 to 1048575

remote-mac

Syntax **remote-mac** *ieee-mac-addr*
no remote-mac

Context config>service>epipe>sap>cem

Description This command defines the destination IEEE MAC address to be used to reach the remote end of the MEF 8 Epipe.

Default 00:00:00:00:00:00

Parameters *ieee-mac-addr* — specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff, where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any unicast MAC addresses and non-IEEE-reserved group multicast MAC addresses. Users should exercise caution when using multicast MAC addresses, as Ethernet frames with a multicast destination address could be flooded when traversing an Ethernet broadcast domain.

report-alarm

Syntax [**no**] **report-alarm** [**stray**] [**malformed**] [**pktloss**] [**overrun**] [**underrun**] [**rpktloss**] [**rfault**] [**rrdi**]

Context config>service>cpipe>sap>cem
config>service>epipe>sap>cem

Description This command enables or disables alarm reporting for CES circuit alarm conditions.

Default On: stray, malformed, pktloss, overrun and underline

Off: rpktloss, rfault, rrdi

Parameters **stray —** reports the reception of packets not destined for this CES circuit

malformed — reports the reception of packets not properly formatted as CES packets

pktloss — reports the lack of reception of CES packets

overrun — reports the reception of too many CES packets resulting in an overrun of the receive jitter buffer. Overrun is supported on Cpipes and MEF 8 Epipes.

underrun — reports the reception of too few CES packets resulting in an underrun of the receive jitter buffer. Underrun is supported on Cpipes and MEF 8 Epipes.

rpktloss — reports that the remote peer is currently in packet loss status

rfault — reports that the remote TDM interface is currently not in service

rrdi — reports that the remote TDM interface is currently in RDI status

rtp-header

Syntax [**no**] **rtp-header**

Context config>service>cpipe>sap>cem
config>service>epipe>sap>cem

Description This optional command inserts RTP headers operating in absolute mode into the CES packets. If asym-delay-control is enabled, this command cannot be enabled for Cpipes.

The **no** form of this command will not insert RTP headers into CES packets.

Default no rtp-header

SAP QoS and IP Filter Policy Commands

egress

Syntax **egress**

Context config>service>apipe>sap
config>service>apipe>sap-aggregation-group
config>service>cpipe>sap
config>service>epipe>sap
config>service>fpipe>sap
config>service>hpipe>sap
config>service>ipipe>sap

Description This command enables the context to configure egress QoS policies for a SAP or a SAP aggregation group.

If no SAP egress QoS policy is defined, the system default SAP egress QoS policy is used for egress processing.

ingress

Syntax **ingress**

Context config>service>apipe>sap
config>service>apipe>sap-aggregation-group
config>service>cpipe>sap
config>service>epipe>sap
config>service>fpipe>sap
config>service>hpipe>sap
config>service>ipipe>sap

Description This command enables the context to configure ingress QoS policies for a SAP or a SAP aggregation group.

If no SAP ingress QoS policy is defined, the system default SAP ingress QoS policy is used for ingress processing.

agg-rate-limit

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

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

Description This command sets the aggregate rate limits (PIR and CIR) for the SAP. The *agg-rate* sets the PIR value. The *cir-rate* sets the CIR value. When aggregate rate limits are configured on a second-generation (Gen-2) adapter card, the scheduler mode must be set to 16-priority. On the third-generation (Gen-3) Ethernet adapter card, only the PIR component (*agg-rate*) of the **agg-rate-limit** can be set and the scheduler mode is always 4-priority-hqos. For information on adapter card generations, refer to the "Evolution of Ethernet Adapter Cards, Modules, and Platforms" section in the 7705 SAR OS Interface Configuration Guide.

Configuring the *cir-rate* is optional. If a *cir-rate* is not entered, then the *cir-rate* is set to its default value (0 kb/s). If a *cir-rate* has been set and the *agg-rate* is changed without re-entering the *cir-rate*, the *cir-rate* automatically resets to 0 kb/s. For example, to change the *agg-rate* from 2000 to 1500 while maintaining a *cir-rate* of 500, use the command **agg-rate-limit** 1500 **cir** 500.

The *cir-rate* for this command is blocked on Gen-3 Ethernet adapter cards and Gen-3 platforms with Ethernet ports, such as the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X. Gen-3 Ethernet adapter cards and Gen-3 platforms with Ethernet ports support 4-priority-hqos scheduling mode. Gen-3 platforms with a TDM MDA support 4-priority scheduling mode.

If the specified SAP is a LAG SAP, then *agg-rate* can be configured regardless of the scheduler mode setting, and *cir-rate* can be configured but is not applied to Gen-3 Ethernet adapter cards or platforms. If the active port is on a Gen-3 card or platform, then only agg-rate is applicable. If the active port is on a Gen-2 card, then agg-rate and cir-rate apply when the scheduler mode is set to 16-priority. If the active port is on a Gen-1 card, then *agg-rate* and *cir-rate* are not applicable.

The **no** form of the command sets the *agg-rate* to the maximum and the *cir-rate* to 0 kb/s.

Default no agg-rate-limit

Parameters *agg-rate* — sets the PIR for the aggregate of all the queues on the SAP. The **max** keyword applies the maximum physical port rate possible.

Values 1 to 10000000 kb/s, or **max**

Default max

cir-rate — sets the CIR for the aggregate of all the queues on the SAP

Values 0 to 10000000 kb/s, or **max**

Default 0 kb/s

filter

Syntax **filter ip** *ip-filter-id*
no filter [**ip** *ip-filter-id*]

Context config>service>epipe>sap>ingress
config>service>ipipe>sap>ingress

Description This command associates an IP filter policy with an ingress Epipe or Ipipe SAP.

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 *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 will be displayed.

The **no** form of the command removes any configured filter ID association with the SAP. The filter policy cannot be deleted until it is removed from all SAPs where it is applied.

Default no filter

Parameters *ip-filter-id* — specifies the IP filter policy. The filter ID or filter name must already exist within the created IP filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

match-qinq-dot1p

Syntax **match-qinq-dot1p** {**top** | **bottom**}
no match-qinq-dot1p

Context config>service>epipe>sap>ingress
config>service>ipipe>sap>ingress

Description This command specifies which dot1q tag position (top or bottom) in a qinq-encapsulated packet should be used when QoS evaluates dot1p classification.

The **no** form of the command restores the default dot1p evaluation behavior for the SAP, which means that the inner (bottom) tag (second tag) dot1p bits are used for classification.

By default, the dot1p bits from the inner tag service-delineating dot1q tag are used.

Table 42 shows which set of dot1p bits are used for QoS purposes when **match-qinq-dot1p** is configured. To use the table, find the row that represents the settings for **Port/SAP Type** and **Match-QinQ-Dot1q Setting**. Use the **Existing Packet Tags** column to identify which dot1q tags are available in the packet. Then use the **P-bits Used for Match** column to identify which dot1q tag contains the dot1p bits that are used for QoS dot1p classification.

Default no match-qinq-dot1p

Parameters **top —** the **top** parameter and **bottom** parameter are mutually exclusive. When the **top** parameter is specified, the outer tag's dot1p bits (topmost P-bits) are used (if existing) to match any **dot1p** *dot1p-value* entries.

bottom — the **bottom** parameter and **top** parameter are mutually exclusive. When the **bottom** parameter is specified, the bottommost P-bits (second tag's P-bits) are used (if existing) to match any **dot1p** *dot1p-value* entries.

Table 42: Match-QinQ-Dot1p Matching Behavior

Port/ SAP Type	Match-QinQ-Dot1p Setting [1]	Existing Packet Tags	P-bits Used for Match
Null	n/a	None	None
Null	n/a	Dot1p (VLAN ID 0)	None [2]
Null	n/a	Dot1q	None [2]
Null	n/a	TopQ BottomQ	None [2]
Dot1Q	n/a	None	None
Dot1Q	n/a	Dot1p (default SAP VLAN ID 0)	Dot1p P-bits
Dot1Q	n/a	Dot1q	Dot1q P-bits
QinQ/ X.Y	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.Y	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ X.0	Top	TopQ	TopQ P-bits
QinQ/ X.0	Default or Bottom	TopQ	TopQ P-bits
QinQ/ X.0	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.0	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ X.*	Top	TopQ	TopQ P-bits
QinQ/ X.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ X.*	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ 0.*	Top	None	None
QinQ/ 0.*	Default or Bottom	None	None
QinQ/ 0.*	Top	TopQ	TopQ P-bits
QinQ/ 0.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ 0.*	Top	TopQ BottomQ	TopQ P-bits

Table 42: Match-QinQ-Dot1p Matching Behavior (Continued)

Port/ SAP Type	Match-QinQ-Dot1p Setting [1]	Existing Packet Tags	P-bits Used for Match
QinQ/ 0.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ *.*	Top	None	None
QinQ/ *.*	Default or Bottom	None	None
QinQ/ *.*	Top	TopQ	TopQ P-bits
QinQ/ *.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ *.*	Top	TopQ BottomQ	TopQ P-bits
QinQ/ *.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits

Notes:

1. "Default" in this column refers to the **no** form of the **match-qinq-dot1p** command.
2. For null encapsulation, the 7705 SAR does not process dot1p bits.

qinq-mark-top-only

Syntax **[no] qinq-mark-top-only**

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

Description When enabled, the **qinq-mark-top-only** command specifies which P-bits to mark during packet egress. When disabled, both sets of P-bits are marked. When enabled, only the P-bits in the top Q-tag are marked. The **no** form of the command is the default state (disabled).

Table 43 shows the dot1p re-marking behavior for different egress port type/SAP type combinations and **qinq-mark-top-only** state, where "False" represents the default (disabled) state.

If a new tag is pushed, the dot1p bits of the new tag will be zero (unless the new tag is re-marked by the egress policy. The dot1p bits are configured using the **dot1p** parameter under the **config**>**qos** context.

Table 43: Dot1P Re-marking Behavior for the QinQ-mark-top-only Command

Egress Port Type/SAP Type	QinQ-mark-top-only State	Egress P-Bits Marked or Re-marked
Null [1]	n/a	None
Dot1q/ X [1]	n/a	Outer tag
Dot1q/ * [2]	n/a	None
QinQ/ X.Y [1]	False	Two outer tags [3]
	True	Outer tag [3]
QinQ/ X.* [4]	True or False	Outer tag
QinQ/ X.0 [1]	True or False	Outer tag
QinQ/ 0.* [1]	True or False	None
QinQ/ *.* [2]	True or False	None

Notes:

1. This port type/SAP type is supported by the following services: Epipe, Ipipe, VPLS, IES, and VPRN.
2. This port type/SAP type is supported by the following services: Epipe and VPLS.
3. Normally, when a new tag is pushed, the dot1p bits of the new tag will be zero, unless the P-bits are re-marked by the egress policy. However, an exception to this occurs when the egress SAP type is X.Y and only one new outer tag must be pushed. In this case, the new outer tag will have its dot1p bits set to the inner tag's dot1p bits.
4. This port type/SAP type is supported by the following services: Epipe, Ipipe, and VPLS.

Default no qinq-mark-top-only (disabled)

qos

Syntax **qos** *policy-id*
no qos

Context config>service>apipe>sap>egress
config>service>apipe>sap>ingress
config>service>apipe>sap-aggregation-group>egress
config>service>apipe>sap-aggregation-group>ingress
config>service>cpipe>sap>egress
config>service>cpipe>sap>ingress
config>service>epipe>sap>egress
config>service>epipe>sap>ingress
config>service>fpipe>sap>egress
config>service>fpipe>sap>ingress
config>service>hpipe>sap>egress
config>service>hpipe>sap>ingress
config>service>ipipe>sap>egress
config>service>ipipe>sap>ingress

Description This command associates a QoS policy with an ingress or egress for a SAP or a SAP aggregation group.

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the *policy-id* does not exist, an error will be returned.

For SAPs, the **qos** command is used to associate both ingress and egress QoS policies on a per-SAP basis. The **qos** command only allows ingress policies to be associated on the SAP ingress and egress policies on the SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

For SAP aggregation groups, the **qos** command is used to apply an ingress and egress QoS policy to each N > 1 service. The configuration of a QoS policy on a per-SAP basis is not permitted. Any existing QoS profile can be applied to the N > 1 service. The QoS policy governs the whole PW service at the packet layer, irrespective of the number of SAPs that are bound to the N > 1 service.

For SAPs that are bound to an aggregation group, VC-based QoS using Layer 2 traffic descriptor profiles can be applied at ingress and egress. See the traffic-desc command for more information.

Only one ingress and one egress QoS policy can be associated with a SAP or a SAP aggregation group at one time.

By default, no specific QoS policy is associated with the SAP or SAP aggregation group 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 SAP aggregation group, and the QoS policy reverts to the default.

Parameters *policy-id* — associates the ingress or egress policy ID with the SAP or SAP aggregation group on ingress or egress. The policy ID or name must already exist. The *policy-name* does not apply to a SAP aggregation group.

Values 1 to 65535, or *policy-name* (up to 64 characters)

scheduler-mode

Syntax **scheduler-mode** {**4-priority** | **16-priority**}

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

Description This command sets the scheduler mode for the SAP and is part of the hierarchical QoS (H-QoS) feature on the 7705 SAR.

If the mode is 4-priority, then the SAP is considered an unshaped 4-priority SAP and the agg-rate-limit cannot be changed from its default values.

If the mode is 16-priority and the **agg-rate limit** parameters are configured to be non-default values, then the SAP is considered a shaped SAP. If the **agg-rate limit** parameters are left in their default settings, the SAP is considered an unshaped, 16-priority SAP.

This command is blocked on third-generation (Gen-3) Ethernet adapter cards and platforms, such as the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X with Ethernet ports, which use only 4-priority-hqos scheduling mode. It is also blocked on the 7705 SAR-X with a TDM MDA, which uses only 4-priority scheduling mode.

If the specified SAP is a LAG SAP, **scheduler-mode** can be configured but is not applied to Gen-3 adapter cards and platforms. If one of the ports in the LAG is on a Gen-1 adapter card, then **scheduler-mode** cannot be configured.

Default 4-priority

Parameters **4-priority —** sets the scheduler mode for the SAP to be 4-priority mode

16-priority — sets the scheduler mode for the SAP to be 16-priority mode

shaper-group

Syntax [**no**] **shaper-group** *shaper-group-name*

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

Description This command applies a shaper group to a SAP. The shaper group must already be created and must be within the shaper policy assigned to the Ethernet MDA (for ingress) or port (for egress). A shaper group is a dual-rate aggregate shaper used to shape aggregate access ingress or egress SAPs at a shaper group rate. Multiple aggregate shaper groups ensure fair sharing of available bandwidth among different aggregate shapers.

The default shaper group cannot be deleted.

The **no** form of this command removes the configured **shaper-group**.

Default shaper-group “default”

Parameters *shaper-group-name* — the name of the shaper group. To access the default shaper group, enter “default”.

create — keyword used to create a shaper group

SAP Microwave Link Commands

mw

Syntax **mw**

Context config>service>epipe>sap

Description This command enables the context to configure microwave link parameters for an Epipe SAP in a Microwave Awareness mixed link scenario. See Configuring Epipe SAP Microwave Link Parameters for Interworking with TDM2Ethernet for more information.

compression

Syntax **compression** *source-mac destination-mac* [**rtp**]
no compression

Context config>service>epipe>sap>mw

Description This command configures the TDM2 Ethernet compression on an Epipe SAP that is using a microwave link.

The **no** form of the command removes any configured TDM2 Ethernet compression associated with the Epipe SAP.

Default n/a

Parameters *source-mac* — the source 9500 MSS MAC address

Values xx:xx:xx:xx:xx:xx

destination-mac — the destination 7705 SAR-8 or 7705 SAR-18 MAC address

Values xx:xx:xx:xx:xx:xx

rtp — enables the RTP header on the Epipe SAP

Service Billing Commands

accounting-policy

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

Context config>service>apipe>sap
config>service>apipe>sap-aggregation-group
config>service>cpipe>sap
config>service>epipe>sap
config>service>fpipe>sap
config>service>hpipe>sap
config>service>ipipe>sap

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

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

The **no** form of this command removes the accounting policy association from the SAP or SAP aggregation group, and the accounting policy reverts to the default.

Default no accounting-policy

Parameters *acct-policy-id* — the accounting *policy-id* as configured in the **config>log>accounting-policy** context

Values 1 to 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>apipe>sap
config>service>apipe>sap-aggregation-group
config>service>cpipe>sap
config>service>epipe>sap
config>service>fpipe>sap
config>service>hpipe>sap
config>service>ipipe>sap

Description This command enables accounting and statistical data collection for the SAP or SAP aggregation group. 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 CSM. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued, the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default collect-stats

VLL SDP Commands

spoke-sdp

Syntax **spoke-sdp** *sdp-id*:*vc-id* [**create**] [**no-endpoint**]
spoke-sdp *sdp-id*:*vc-id* [**create**] **endpoint** *endpoint-name* [**icb**]
no spoke-sdp *sdp-id*:*vc-id*

Context config>service>apipe
config>service>cpipe
config>service>fpipe
config>service>hpipe
config>service>ipipe

Description This command binds a service to an existing Service Destination Point (SDP). The syntax for an Epipe spoke SDP has additional parameters. See spoke-sdp for the Epipe syntax.

A spoke SDP is treated as 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 on which it was received.

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 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 7705 SAR devices can participate in the service.

The **endpoint** command allows multiple spoke SDPs to be associated with the endpoint, providing PW redundancy capability. The endpoint must be defined using the **create** command before multiple spoke SDPs can be associated with the endpoint. The **no-endpoint** command removes the endpoint and the spoke SDP associations.

On Cpipe spoke SDPs, you can configure **icb** to provide resiliency by reducing packet loss when an active endpoint is switched from a failed node of an MC-APS group to a standby node. For example, if a port on an active MC-APS node fails, the port on the peer becomes active, but traffic continues to route to the previously active MC-APS node until it detects the failure. ICB spoke SDPs ensure that in-flight packets are delivered to the newly active MC-APS node. Two ICB spoke SDPs must be created. The ICB associated with the MC-APS on the first node must be associated with the pseudowire on the second node. Likewise, the ICB associated with the MC-APS on the second node must be associated with the pseudowire on the first node.

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

Default no *sdp-id* is bound to a service

Parameters *sdp-id* — uniquely identifies the SDP

Values 1 to 17407

vc-id — identifies the virtual circuit

Values 1 to 4294967295

endpoint-name **—** specifies the name of the service endpoint

no-endpoint — removes a spoke SDP association

icb — enables ICB (inter-chassis backup)

spoke-sdp

Syntax **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** *endpoint-name* [**icb**]
no spoke-sdp *sdp-id*:*vc-id*

Context config>service>epipe

Description This command binds an Epipe service to an existing Service Destination Point (SDP). The syntax for an Apipe, Cpipe, or Ipipe spoke SDP has additional parameters. See spoke-sdp for the Apipe, Cpipe, or Ipipe syntax.

A spoke SDP is treated as 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 on which it was received.

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 an Epipe 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 7705 SAR devices can participate in the service.

The **endpoint** command allows multiple spoke SDPs to be associated with the endpoint, providing PW redundancy capability. The endpoint must already be defined in the **config>service>epipe** context in order to associate multiple spoke SDPs with the endpoint.

The **icb** (inter-chassis backup) spoke SDP provides resiliency by reducing packet loss when an active endpoint is switched from a failed node of an MC-LAG group to a standby node. For example, if a port on an active MC-LAG node fails, the port on the peer becomes active, but traffic continues to route to the previously active MC-LAG node until it detects the failure. ICB spoke SDPs ensure that in-flight packets are delivered to the newly active MC-LAG node. Two ICB spoke SDPs must be created. The ICB associated with the MC-LAG on the first node must be associated with the pseudowire on the second node. Likewise, the ICB associated with the MC-LAG on the second node must be associated with the pseudowire on the first node.

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

Default no *sdp-id* is bound to a service

Parameters *sdp-id* — uniquely identifies the SDP

Values 1 to 17407

vc-id — identifies the virtual circuit

Values 1 to 4294967295

vc-type — overrides the default VC type signaled for the spoke binding to the far end of the SDP. The VC type is a 15-bit quantity containing a value that 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 binding's VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

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

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

Values ether | vlan

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

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.

endpoint-name **—** specifies the name of the service endpoint

no-endpoint — removes a spoke SDP association

icb — enables **icb** (inter-chassis backup) spoke SDP to provide resiliency by reducing packet loss when an active endpoint is switched from a failed node of an MC-LAG group to a standby node.

egress

Syntax **egress**

Context config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>hpipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description This command configures the egress SDP context.

ingress

Syntax **ingress**

Context config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>hpipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description This command configures the ingress SDP context.

precedence

Syntax **precedence** [*precedence-value* | **primary**]
no precedence

Context config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>hpipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description This command specifies the precedence of the spoke SDP when there are multiple spoke SDPs associated with one service endpoint. One SDP binding can be assigned to be the primary SDP binding, leaving three bindings for secondary bindings, or, if no primary spoke SDP is defined, up to four secondary spoke SDPs can be configured. When an SDP binding goes down, the next highest precedence SDP binding will begin to forward traffic.

The **no** form of the command returns the precedence value to the default.

Default 4

Parameters *precedence-value* — specifies the spoke SDP precedence

Values 1 to 4 (where 1 is the highest precedence)

primary — makes the specified spoke SDP the primary spoke SDP (primary is indicated on the CLI display as the value 0)

pw-status-signaling

Syntax [**no**] **pw-status-signaling**

Context config>service>epipe>spoke-sdp

Description This command enables pseudowire status signaling for spoke SDPs. The **no** form of this command disables pseudowire status signaling. When pseudowire status signaling is disabled, a 7705 SAR does not include the PW status TLV in the initial label mapping message of the pseudowire that is used for a spoke SDP. This forces both 7705 SAR PEs to use the pseudowire label withdrawal method for signaling pseudowire status.

If pseudowire status signaling is enabled, the 7705 SAR includes the pseudowire status TLV in the initial label mapping message for the pseudowire.

Default pw-status-signaling

vc-label

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

Context config>service>apipe>spoke-sdp>egress
config>service>cpipe>spoke-sdp>egress
config>service>epipe>spoke-sdp>egress
config>service>fpipe>spoke-sdp>egress
config>service>hpipe>spoke-sdp>egress
config>service>ipipe>spoke-sdp>egress

Description This command configures the egress VC label.

Parameters *egress-vc-label* — indicates a specific connection

Values 16 to 1048575

vc-label

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

Context config>service>apipe>spoke-sdp>ingress
config>service>cpipe>spoke-sdp>ingress
config>service>epipe>spoke-sdp>ingress
config>service>fpipe>spoke-sdp>ingress
config>service>hpipe>spoke-sdp>ingress
config>service>ipipe>spoke-sdp>ingress

Description This command configures the ingress VC label.

Parameters *ingress-vc-label* — indicates a specific connection

Values 2048 to 18431

vlan-vc-tag

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

Context config>service>epipe>spoke-sdp

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

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

The **no** form of this command disables the command

Default no vlan-vc-tag

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

ce-address

Syntax **ce-address** *ip-address*
no ce-address

Context config>service>ipipe>sap
config>service>ipipe>spoke-sdp

Description This command specifies the IP address of the CE device associated with an Ipipe SAP or spoke SDP. In the case of a SAP, it is the address of the CE device directly attached to the SAP. For a spoke SDP, it is the address of the CE device reachable through that spoke SDP (for example, attached to the SAP on the remote node). The address must be a host address (no subnet addresses are accepted) as there must be only one CE device attached to an Ipipe SAP. The CE address specified at one end of an Ipipe will be used in processing ARP messages at the other endpoint, as the router acts as a proxy for ARP messages.

Default n/a

Parameters *ip-address* — specifies the IP address of the CE device associated with an Ipipe SAP

control-word

Syntax **control-word**
no control-word

Context config>service>apipe>spoke-sdp
config>service>cpipe>spoke-sdp
config>service>epipe>spoke-sdp
config>service>fpipe>spoke-sdp
config>service>hpipe>spoke-sdp
config>service>ipipe>spoke-sdp

Description This command indicates whether the control word is used or not. The value of the control word is negotiated with the peer.

This command is mandatory for SAToP and CESoPSN encapsulation.

The control word must be used for any pseudowire service that transports packets that are less than 64 bytes. When the control word is enabled and the packet length is less than 64 bytes (that is, the length of the Layer 2 payload plus the length of the control word), the length field in the control is set to the length of the packet. The CE-bound PE uses the length field in the control word to determine the size of the padding that was added by the PSN so that the PE can extract the PW payload from the PW packet.

Note: If the control word is not set, the PE cannot determine the original length of the packet and will forward the payload, including the padding bits. On reception of the padded packet, the CE will drop the packet. See RFC 4385 for more information.

Default The control word is not enabled.

SDP Cell Concatenation Commands

cell-concatenation

Syntax **cell-concatenation**

Context config>service>apipe>spoke-sdp

Description This command enables the context to provide access to the various options that control the termination of ATM cell concatenation into an MPLS frame. Several options can be configured simultaneously. The concatenation process for a given MPLS packet ends when the first concatenation termination condition is met. The concatenation parameters apply only to ATM N-to-1 cell mode VLL.

Frame boundaries are not configurable.

clp-change

Syntax [**no**] **clp-change**

Context config>service>apipe>spoke-sdp>cell-concatenation

Description This command enables the configuration of CLP change to be an indication to complete the cell concatenation operation.

The **no** form of the command resets the configuration to ignore the CLP change as an indication to complete the cell concatenation.

max-cells

Syntax **max-cells** *cell-count*
no max-cells [*cell-count*]

Context config>service>apipe>spoke-sdp>cell-concatenation

Description This command enables the configuration of the maximum number of ATM cells to accumulate in an MPLS packet. The remote peer will also signal the maximum number of concatenated cells it is willing to accept in an MPLS packet. When the lesser of the configured value and the signaled value is reached, the MPLS packet is queued for transmission onto the pseudowire. It is ensured that the MPLS packet MTU conforms to the configured service MTU.

If the max-delay and jitter buffer options are not configured, then the maximum number of cells allowed in a single VLL frame must be less than the configured service-mtu size.

The **no** form of this command sets max-cells to the value "1", indicating that no concatenation will be performed.

Parameters *cell-count* — specifies the maximum number of ATM cells to be accumulated in an MPLS packet before queuing the packet for transmission onto the pseudowire

Values 1 to 29

Default 1

max-delay

Syntax **max-delay** *delay-time*
no max-delay [*delay-time*]

Context config>service>apipe>spoke-sdp>cell-concatenation

Description This command enables the configuration of the maximum amount of time to wait while performing ATM cell concatenation into an MPLS packet before transmitting the MPLS packet. This places an upper bound on the amount of delay introduced by the concatenation process. When this amount of time is reached from when the first ATM cell for this MPLS packet was received, the MPLS packet is queued for transmission onto the pseudowire.

The **no** form of this command resets **max-delay** to its default value.

Parameters *delay-time* — specifies the maximum amount of time, in hundreds of microseconds, to wait before transmitting the MPLS packet with whatever ATM cells have been received. For example, a value of 1 equals 100 µs, and a value of 400 equals 40000 µs, or 40 ms.

Values 1 to 400

Default 400

ATM Commands

atm

Syntax **atm**

Context config>service>apipe>sap

Description This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:

- configuring ATM port or ATM port-related functionality on T1/E1 ASAP Adapter cards or OC3/STM1 Adapter cards
- configuring ATM-related configuration for ATM-based SAPs that exist on T1/E1 ASAP Adapter cards or on OC3/STM1 Adapter cards

If ATM functionality is not supported for a given context, the command returns an error.

egress

Syntax **egress**

Context config>service>apipe>sap>atm

Description This command provides access to the context to configure egress ATM traffic policies for the SAP.

ingress

Syntax **ingress**

Context config>service>apipe>sap>atm

Description This command provides access to the context to configure ingress ATM traffic policies for the SAP.

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
no traffic-desc

Context config>service>apipe>sap>atm>egress
config>service>apipe>sap>atm>ingress

Description This command assigns an ATM traffic descriptor profile to a SAP.

When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction. When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.

For SAPs that belong to a SAP aggregation group, the egress traffic descriptor can be changed. The ingress traffic descriptor cannot be changed from the default (1). Attempting to change the ingress traffic descriptor will cause an error message to be displayed.

The **no** form of the command reverts the traffic descriptor to the default traffic descriptor profile.

Default The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.

Parameters *traffic-desc-profile-id* — specifies a defined traffic descriptor profile (see the QoS **atm-td-profile** command)

vcid-translation

Syntax **vcid-translation** *vpi-vci*
no vcid-translation

Context config>service>apipe>sap>atm

Description This optional command is used only with SAPs that are configured as members of a SAP aggregation group.

The **vcid-translation** command is used when traffic arrives on multiple SAPs within the same SAP aggregation group with the same VPI/VCI value. In this case, the VPI/VCI values for incoming ATM cells that are to be aggregated in a single ATM PW must be translated to preserve their individual identification before the cells are mapped to the ATM PW payload.

When vcid-translation is configured for a SAP, ingress cells have their VPI/VCI values translated (rewritten) to the vcid-translation value. On the same node, egressing cells have their VPI/VCI values translated back to the VPI/VCI of the SAP.

If the vcid-translation for any ATM sap-aggregation-group SAP is not configured, the ingress VPI/VCI values are retained.

The 7705 SAR performs a check to ensure the uniqueness of the translated VPI/VCI values for all of the SAPs of the same ATM PW service, that is, within the same SAP aggregation group. If there are duplicate VC identifiers, the status of the VCs are set to operationally down and flagged as `ApipeSapVcIdNotUnique`.

It is the responsibility of the user to ensure uniqueness of the VPI and VCI values after translation.

Default no vcid-translation

Parameters *vpi-vci* — specifies the VPI and VCI translation values

ATM OAM Commands

oam

Syntax **oam**

Context config>service>apipe>sap>atm

Description This command enables the context to configure OAM functionality for a PVCC delimiting a SAP.

The T1/E1 ASAP Adapter cards and 4-port OC3/STM1 Clear Channel Adapter card support the generation of F4 (VP) and F5 (VC) AIS cells when the Apipe service is operationally down. When the Apipe service is operationally up, OAM cells are transported over the Apipe and are transparent to the 7705 SAR. This capability is in accordance with ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance.

alarm-cells

Syntax [**no**] **alarm-cells**

Context config>service>apipe>sap>atm>oam

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC terminations to monitor and report the status of their connection by propagating fault information through the network and by driving the PVCC operational status.

The 7705 SAR Apipe does not support PVCC terminations. Instead, it allows OAM cells to be transported transparently from end-to-end. When this command is enabled, AIS cells are generated when an Apipe or corresponding SAP is operationally down.

The **no** command disables alarm-cells functionality for the Apipe. When alarm-cells functionality is disabled, AIS cells are not generated as result of the Apipe or corresponding SAP going into the operationally down state.

Default enabled

Show Commands

Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

all

Syntax **all**

Context show>service>id

Description This command displays detailed information for all aspects of the service.

Output The Sample Output (Apipe ATMVcc Service) is an example of **id** *service-id* **all** information, and Table 44 describes the fields. Following the table are output examples for:

- Sample Output (Apipe ATMVpc Service)
- Sample Output (Cpipe Service)
- Sample Output (Epipe Service)
- Sample Output (TDM SAP-to-Ethernet SAP MEF 8 Epipe Service)
- Sample Output (Fpipe Service)
- Sample Output (Hpipe Service)
- Sample Output (Ipipe Service)
- Sample Output (Ipipe Frame Relay to Ethernet Service)
- Sample Output (Ipipe cHDLC to Ethernet Service)

Sample Output (Apipe ATMVcc Service)

```
===============================
*A:ALU-A>show>service# id 2 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 2
Service Type      : Apipe               VLL Type          : ATMVCC
Name              : apipe2
Description       : (Not Specified)
Customer Id       : 2
Last Status Change: 03/11/2008 19:58:19
Last Mgmt Change  : 03/28/2008 19:49:51
Admin State       : Down                Oper State        : Down
MTU               : 1508
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
-------------------------------------------------------------------------------
```

```
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 2:2  -(138.120.38.1)
-------------------------------------------------------------------------------
SDP Id             : 2:2                       Type              : Spoke
VC Type            : ATMVCC                    VC Tag            : 0
Admin Path MTU     : 0                         Oper Path MTU     : 0
Far End            : 138.120.38.1              Delivery          : MPLS

Admin State        : Up                        Oper State        : Down
Acct. Pol          : None                      Collect Stats     : Disabled
Ingress Label      : 0                         Egress Label      : 0
Ing mac Fltr       : n/a                       Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                       Egr ip Fltr       : n/a
Admin ControlWord  : Not Preferred             Oper ControlWord  : False
Admin BW(Kbps)     : 0                         Oper BW(Kbps)     : 0
Last Status Change : 03/11/2008 19:58:19       Signaling         : TLDP
Last Mgmt Change   : 03/28/2008 19:49:51
Endpoint           : N/A                       Precedence        : 4
PW Status Sig      : Enabled
Class Fwding State : Down
Flags              : StandbySigSlaveTxDown
Mac Move           : Ukwn                      Blockable Level   : Unknown
Local Pw Bits      : None
Peer Pw Bits       : pwFwdingStandby
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

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

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

Associated LSP LIST :
No LSPs Associated


-------------------------------------------------------------------------------
APIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Admin Concat Limit : 1                         Oper Concat Limit : 1
Peer Concat Limit  : n/a                       Max Concat Delay  : 400
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
```

```
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/4/1.1:0/32
-------------------------------------------------------------------------------
Service Id         : 2
SAP                : 1/4/1.1:0/32             Encap             : atm
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown L2OperDown
Multi Svc Site     : None
Last Status Change : 03/11/2008 19:58:19
Last Mgmt Change   : 03/28/2008 19:35:51
Sub Type           : regular

Admin MTU          : 1572                     Oper MTU          : 1572
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
tod-suite          : None                     qinq-pbit-marking : both
Egr Agg Rate Limit : max
Endpoint           : N/A

Acct. Pol          : None                     Collect Stats     : Disabled

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        n/a
Off. HiPrio           : 39192                    n/a
Off. LowPrio          : n/a                      n/a

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        n/a
Dro. LowPrio          : n/a                      n/a
For. InProf           : 19596                    19596
For. OutProf          : 19596                    19596

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        n/a
Dro. OutProf          : n/a                      n/a
For. InProf           : 39192                    39192
For. OutProf          : n/a                      n/a
-------------------------------------------------------------------------------
```

```
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 39192                   n/a
Off. LoPrio           : n/a                     n/a
Dro. HiPrio           : 0                       n/a
Dro. LoPrio           : n/a                     n/a
For. InProf           : 19596                   19596
For. OutProf          : 19596                   19596

Egress Queue 1
For. InProf           : 39192                   39192
For. OutProf          : n/a                     n/a
Dro. InProf           : 0                       n/a
Dro. OutProf          : n/a                     n/a
-------------------------------------------------------------------------------
ATM SAP Configuration Information
-------------------------------------------------------------------------------
Ingress TD Profile : 1                        Egress TD Profile : 1
Alarm Cell Handling: Enabled                  AAL-5 Encap       : n/a
OAM Termination    : Disabled                 Periodic Loopback : Disabled
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
```

Table 44: Show Service-ID All Command Output Fields

Label	Description
Service Detailed Information	
Service Id	Identifies the service by its ID number
Service Type	Specifies the type of service
VLL Type	Specifies the VLL type
Name	Specifies the optional configured service name
Description	Displays generic information about the service
Customer Id	Identifies the customer by its ID number
Creation Origin	Specifies how the service was created
Last Status Change	Displays the date and time of the most recent status change to this service
Last Mgmt Change	Displays the date and time of the most recent management-initiated change to this service
Admin State	Specifies the desired state of the service
Oper State	Specifies the operating state of the service

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
MTU	Specifies the service MTU
Vc Switching	Specifies whether the service is configured as a PW switching point
SAP Count	Displays the number of SAPs specified for this service
SDP Bind Count	Displays the number of SDPs bound to this service
Service Destination Points (SDPs)	
Description	Displays generic information about the SDP
SDP Id	Identifies the SDP
Type	Identifies the service SDP binding type (for example, spoke)
Split Horiz Grp	Not applicable
VC Type	Displays the VC type for the SDP (for example, CESoPSN)
VC Tag	The explicit dot1q value used when encapsulating to the SDP far end
Admin Path MTU	Specifies the desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented
Oper Path 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
Delivery	Specifies the type of delivery used by the SDP (MPLS or GRE)
Far End	Displays the IP address of the remote end of the MPLS or GRE tunnel defined by this SDP
Admin State	Specifies the administrative state of this SDP
Oper State	Specifies the operational state of this SDP
Acct. Pol	The accounting policy ID assigned to the SAP
Collect Stats	Specifies whether collect stats is enabled
Ingress Label	Displays the label used by the far-end device to send packets to this device in this service by this SDP
Egress Label	Displays the label used by this device to send packets to the far-end device in this service by this SDP

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
Admin ControlWord	Specifies the administrative state of the control word: Preferred (control word enabled) or Not Preferred (control word disabled)
Oper ControlWord	Specifies the operational state of the control word: True (control word enabled) or False (control word disabled)
Last Status Change	Specifies the time of the most recent operating status change to this spoke SDP
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this spoke SDP
PW Status Sig	Specifies whether pseudowire status signaling for spoke SDPs is enabled or disabled
Flags	Displays the conditions that affect the operating status of this spoke SDP. Display output includes PathMTUtooSmall, SdpOperDown, NoIngVCLabel, NoEgrVCLabel, StandbySigSlaveTxDown, and so on.
Mac Move	Indicates the administrative state of the MAC movement feature associated with the service
Local Pw Bits	Displays the setting of the local pseudowire bits
Peer Pw Bits	Displays the setting of the peer pseudowire bits. Display output includes pwNotforwarding, psnIngressFault, psnEgressFault, IacIngressFault, lacEgressFault, PwFwdingStandby, and so on.
Peer Fault Ip	N/A
Peer Vccv CV Bits	Displays the setting of the pseudowire peer VCCV control verification bits (lspPing)
Peer Vccv CC Bits	Displays the setting of the pseudowire peer VCCV control channel bits (pwe3ControlWord and/or mplsRouterAlertLabel)
Keepalive Information	
Admin State	Specifies the administrative state of the keepalive protocol
Oper State	Specifies the operational state of the keepalive protocol
Hello Time	Specifies how often the SDP Echo Request messages are transmitted on this SDP

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
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
Statistics	
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets
I. Dro. Pkts.	Specifies the number of dropped ingress packets
I. Fwd. Octs.	Specifies the number of forwarded ingress octets
I. Dro. Octs.	Specifies the number of dropped ingress octets
E. Fwd. Pkts.	Specifies the number of forwarded egress packets
E. Fwd. Octets	Specifies the number of forwarded egress octets
Eth-Cfm Configuration Information	
Md-index	Displays the value of the MD index
Direction	Displays the direction of the MEP
Ma-index	Displays the value of the MA index
Admin	Displays the administrative state of the MEP (enabled or disabled)
MepId	Displays the MEP-ID
CCM-Enable	Displays the status of the Continuity Check Message (CCM)
LowestDefectPri	Displays a configured value that defects are evaluated against
HighestDefect	Displays the highest defect
Defect Flags	Indicates the defect flags
Mac Address	Displays the MAC address (the MAC address for a spoke SDP is the system MAC address; for a SAP, it is the port MAC address)
CcmLtmPriority	Displays the priority of the CCM Linktrace Message (LTM)
CcmTx	Displays the number of CCM transmissions

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
CcmSequenceErr	Displays the number of CCM sequence errors
DmrRepliesTx	Displays the number of delay measurement replies transmitted
LmrRepliesTx	Displays the number of loss measurement replies transmitted
Dual-Loss Test	Displays the status of the dual-ended loss measurement test (enabled or disabled)
Dual-Loss Thresh	Displays the frame error threshold beyond which an alarm will be raised. The threshold is expressed as a percentage.
Eth-Ais	Displays the status of the ETH-AIS test (enabled or disabled)
Eth-Ais Rx Ais	Indicates whether any ETH-AIS messages have been received
Eth-Ais Tx Priorit*	Displays the priority value of a transmitted ETH-AIS frame
Eth-Ais Rx Interv*	Indicates the interval of a received ETH-AIS frame
Eth-Ais Tx Interva*	Displays the interval of a transmitted ETH-AIS frame
Eth-Ais Tx Counte*	Displays the number of ETH-AIS frames that have been sent
Eth-Ais Tx Levels	Indicates the MD level of transmitted ETH-AIS frames
Eth-Tst	Indicates the status of the ETH-Test (enabled or disabled)
LbRxReply	Displays the number of received loopback (LB) replies
LbRxBadOrder	Displays the number of LB replies that have been received in the wrong order
LbRxBadMsdu	Displays the number of LB replies that have been received with the wrong destination MAC address (MSDU = MAC Service Data Unit)
LbTxReply (Total)	Displays the total number of LBRs (loopback replies) transmitted from this MEP
LbTxReplyNoTLV	Displays the number of LBRs (loopback replies) transmitted from this MEP with no TLV Note that because only LBMs with no TLVs are used for throughput testing, the LbTxReply (Total), LbTxReplyNoTLV, and LbTxReplyWithTLV counters can help debug problems if throughput testing is not working
LbTxReplyWithTLV	Displays the number of LBRs (loopback replies) transmitted from this MEP with TLV

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
LbNextSequence	Displays the sequence number of the next LB transmission
LtNextSequence	Displays the sequence number of the next Linktrace (LT) message transmitted
LtRxUnexplained	Displays the number of the unexplained Linktrace (LT) messages
Associated LSP LIST	
Lsp Name	Specifies the name of the static LSP
Admin State	Specifies the administrative state of the associated LSP
Oper State	Specifies the operational state of the associated LSP
Time Since Last Tr*	Specifies the time that the associated static LSP has been in service
APIPE Service Destination Point specifics	
Admin Concat Limit	Specifies the administrative (configured) value for the maximum number of cells for cell concatenation, as defined via the `max-cells` command
Oper Concat Limit	Specifies the operational value for the maximum number of cells for cell concatenation
Peer Concat Limit	Specifies the far-end value for the maximum number of cells for cell concatenation
Max Concat Delay	Specifies the amount of time to wait while cell concatenation is occurring, as defined via the `max-delay` command
CPIPE Service Destination Point specifics	
Local Bit-rate	Specifies the number of DS0s used by the local SDP
Peer Bit-rate	Specifies the number of DS0s used by the far-end SDP
Local Payload Size	Specifies the local payload size, in bytes, used by the local SDP
Peer Payload Size	Specifies the peer payload size, in bytes, used by the far-end SDP
Local Jitter Buffer	Specifies the jitter buffer size, in milliseconds, used by the local SDP
Peer Jitter Buffer	Specifies the jitter buffer size, in milliseconds, used by the far-end SDP

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
Local Asym Delay	Specifies whether asymmetric delay control is enabled on the local SDP
Peer Asym Delay	Specifies whether asymmetric delay control is enabled on the far-end SDP
Local Sig Pkts	Specifies the type of signaling packets used by the local SDP
Peer Sig Pkts	Specifies the type of signaling packets used by the far-end SDP
Local CAS Framing	Specifies the type of CAS framing used by the local SDP
Peer CAS Framing	Specifies the type of CAS framing used by the far-end SDP
Local RTP Header	Specifies whether the local router inserts the RTP header
Peer RTP Header	Specifies whether the peer router inserts the RTP header
Number of SDPs	Specifies the number of SDPs bound to the service
FPIPE Service Destination Point specifics	
Split Horiz Grp	Not applicable
Endpoint	Not applicable
Precedence	Specifies the precedence level of the SDP binding
Class Fwding State	Specifies the admin state of class-based forwarding on this SDP
IPIPE Service Destination Point specifics	
Precedence	Specifies the precedence level of the SDP binding
IpipeSdpBindCeIpAd*	Specifies the IP address of the Ipipe spoke-sdp
Service Access Points	
Service Id	Identifies the service
SAP	Specifies the ID of the access port where this SAP is defined
Encap	Specifies the encapsulation type for this SAP on the access port
Admin State	Specifies the desired state of the SAP
Oper State	Specifies the operating state of the SAP
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes ServiceAdminDown, PortOperDown, and so on.

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
Last Status Change	Specifies the date and time of the most recent status change to this SAP
Last Mgmt Change	Specifies the date and time of the most recent management-initiated change to this SAP
Dot1Q Ethertype	Identifies the value of the dot1q Ethertype
QinQ Ethertype	Identifies the value of the qinq Ethertype
qinq-pbit-marking	Indicates the qinq P-bit marking for the service: both or top
Hold Meps Up	Specifies if hold-mep-up-on-failure is enabled or disabled for a SAP Epipe
LLF Admin State	Specifies the Link Loss Forwarding administrative state
LLF Oper State	Specifies the Link Loss Forwarding operational state
Admin MTU	Specifies the desired largest service frame size (in octets) that can be transmitted through this SAP to the far-end router, without requiring the packet to be fragmented
Oper MTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SAP to the far-end router, without requiring the packet to be fragmented
Ingr IP Fltr-ID	Specifies the ingress IP filter policy ID assigned to the SAP
Egr IP Fltr-Id	Specifies the egress IP filter policy ID assigned to the SAP
Ingr Mac Fltr-ID	Specifies the ingress MAC filter policy ID assigned to the SAP
Egr Mac Fltr-Id	Specifies the egress MAC filter policy ID assigned to the SAP
Acct. Pol	Specifies the accounting policy applied to the SAP
Collect Stats	Specifies whether accounting statistics are collected on the SAP
FPIPE Service Access Points specifics	
Multi Svc Site	Indicates the multiservice site that the SAP is a member of. Not applicable.
Sub Type	The supported subtype: regular
Split Horiz Grp	Not applicable
Ingr IP Fltr-Id	Not applicable
Egr IP Fltr-Id	Not applicable

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
tod-suite	Indicates whether a time-based policy is applied to a multiservice site. Not applicable.
Ing Agg Rate Limit	Indicates the maximum total rate for all ingress queues on a service SAP in kb/s
Egr Agg Rate Limit	Indicates the maximum total rate for all egress queues on a service SAP in kb/s
Endpoint	Not applicable
FRF-12	Not applicable
IPIPE Service Access Points specifics	
Ipipe SAP ARP Entry Info	Displays the MAC address of the connected CE address after being resolved through the ARP mechanism
qinq-pbit-marking	Indicates the qinq P-bit marking for the service: both or top
QOS	
Ingress qos-policy	Displays the SAP ingress QoS policy ID
FPIPE QOS	
Shared Q plcy	Not applicable
Multipoint shared	
I. Sched Pol	
E. Sched Pol	
SAP Statistics	
Last Cleared Time	Displays the date and time that a clear command was issued on statistics
Forwarding Engine Stats	
Dropped	Indicates the number of packets or octets dropped by the forwarding engine
Off. HiPrio	Indicates the number of high-priority packets or octets offered to the forwarding engine
Off. LowPrio	Indicates the number of low-priority packets offered to the forwarding engine
Queueing Stats (Ingress QoS Policy)	

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
Dro. HiPrio	Indicates the number of high-priority packets or octets discarded, as determined by the SAP ingress QoS policy
Dro. LowPrio	Indicates the number of low-priority packets discarded, as determined by the SAP ingress QoS policy
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP ingress QoS policy
For. OutProf	Indicates the number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP ingress QoS policy
Queueing Stats (Egress QoS Policy)	
Dro. InProf	Indicates the number of in-profile packets or octets discarded, as determined by the SAP egress QoS policy
Dro. OutProf	Indicates the number of out-of-profile packets or octets discarded, as determined by the SAP egress QoS policy
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP egress QoS policy
For. OutProf	Indicates the number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP egress QoS policy
Sap per Queue stats	
Ingress Queue *n*	Specifies the index of the ingress QoS queue of this SAP, where *n* is the index number
Off. HiPrio	Indicates the packets or octets count of the high-priority traffic for the SAP (offered)
Off. LoPrio	Indicates the packets or octets count of the low-priority traffic for the SAP (offered)
Dro. HiPrio	Indicates the number of high-priority traffic packets/octets dropped
Dro. LoPrio	Indicates the number of low-priority traffic packets/octets dropped
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded
For. OutPro	Indicates the number of out-of-profile octets (rate above CIR) forwarded

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
Egress Queue *n*	Specifies the index of the egress QoS queue of the SAP, where *n* is the index number
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	Indicates the number of out-of-profile packets or octets (rate above CIR) forwarded
Dro. InProf	Indicates the number of in-profile packets or octets dropped for the SAP
Dro. OutProf	Indicates the number of out-of-profile packets or octets discarded
ATM SAP Configuration Information	
Ingress TD Profile	The profile ID of the traffic descriptor applied to the ingress SAP
Egress TD Profile	The profile ID of the traffic descriptor applied to the egress SAP
Alarm Cell Handling	Indicates that OAM cells are being processed
OAM Termination	Indicates whether this SAP is an OAM termination point
CEM SAP Configuration Information	
Endpoint Type	Specifies the type of endpoint
Bit-rate	Specifies the number of DS0s or timeslots in the channel group
Payload Size	Specifies the number of octets contained in the payload of a TDM PW packet when the packet is transmitted
Jitter Buffer (ms)	Specifies the size of the receive jitter buffer, expressed in milliseconds
Jitter Buffer (packets)	Specifies the number of packets in the receive jitter buffer
Playout Threshold (packets)	Indicates the number of packet buffers for the playout buffer packets threshold
Use RTP Header	Specifies whether RTP headers are used in CES packets (Yes or No)
Differential	Indicates whether DCR is active
Timestamp Freq	Specifies the optional timestamp frequency
CAS Framing	Specifies the type of CAS framing

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
Effective PDVT	Displays the peak-to-peak packet delay variation (PDV) used by the circuit emulation service. Since the operating system may adjust the jitter buffer setting in order to ensure no packet loss, the configured jitter buffer value may not be the value used by the system. The effective PDVT provides an indication that the PDV has been adjusted by the operating system (see Jitter Buffer)
AsymDelayControl	Specifies whether asymmetric delay control is enabled
RepeatPeriod	Specifies the number of times that the asymmetric delay control analysis is repeated
Samples	Specifies the number of packets that are analyzed (k = 1024)
Cfg Alarm	Specifies the alarms that have alarm reporting enabled
Alarm Status	Indicates the current alarm state (stray, malformed, packet loss, overrun, underrun, remote packet loss, remote fault, or remote RDI)
CEM SAP Statistics	
Packets	(Column heading) Displays the number of packets counted for the statistic since the last counter reset
Seconds	(Column heading) Displays the number of seconds elapsed for the statistic since the last counter reset
Events	(Column heading) Displays the number of events counted for the statistic since the last counter reset
Egress Stats	Indicates that the following statistics are egress statistics
Forwarded	Displays the number of forwarded packets
Dropped	Displays the number of dropped packets
Missing	Displays the number of missing packets
Reordered Forwarded	Displays the number of packets that have been reordered and forwarded
Underrun	Displays the accumulated number of underrun packets for the number of underrun events
Overrun	Displays the accumulated number of overrun packets for the number of overrun events
Misordered Dropped	Displays the number of misordered packets that have been dropped

Table 44: Show Service-ID All Command Output Fields (Continued)

Label	Description
Malformed Dropped	Displays the number of malformed packets that have been dropped
LBit Dropped	Displays the number of L bit marked packets that have been dropped
Error	Displays the accumulated number of seconds that have passed while any error has occurred
Severely Error	Displays the accumulated number of seconds that have passed while severe errors have occurred
Unavailable	Displays the accumulated number of seconds that have passed while the Cpipe or MEF 8 Epipe is unavailable
Failure Count	Displays the accumulated number of failed events
Jitter Buffer Depth	Displays the number of packets sitting in the jitter buffer at that instant for a given Cpipe or MEF 8 Epipe
Jitter Buffer Sampling Complete	Displays the number of asymmetric delay control analysis periods completed
Jitter Buffer Adjust	Displays the number of jitter buffer adjustments made for asymmetric delay control
Ingress Stats	Indicates that the following statistics are ingress statistics
Forwarded	Displays the number of forwarded packets
Dropped	Displays the number of dropped packets

Sample Output (Apipe ATMVpc Service)

```
===============================
*A:ALU-A>show>service# id 5 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 5
Service Type      : Apipe               VLL Type          : ATMVPC
Name              : apipe5
Description       : (Not Specified)
Customer Id       : 2
Last Status Change: 03/11/2008 19:58:19
Last Mgmt Change  : 04/01/2008 16:51:59
Admin State       : Down                Oper State        : Down
MTU               : 1508
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
-------------------------------------------------------------------------------
```

```
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 5:5  -(138.120.20.1)
-------------------------------------------------------------------------------
SDP Id             : 5:5                       Type              : Spoke
VC Type            : ATMVPC                    VC Tag            : 0
Admin Path MTU     : 0                         Oper Path MTU     : 0
Far End            : 138.120.20.1              Delivery          : MPLS

Admin State        : Up                        Oper State        : Down
Acct. Pol          : None                      Collect Stats     : Disabled
Ingress Label      : 0                         Egress Label      : 0
Ing mac Fltr       : n/a                       Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                       Egr ip Fltr       : n/a
Admin ControlWord  : Not Preferred             Oper ControlWord  : False
Admin BW(Kbps)     : 0                         Oper BW(Kbps)     : 0
Last Status Change : 03/11/2008 19:58:19       Signaling         : TLDP
Last Mgmt Change   : 04/01/2008 16:51:59
Endpoint           : N/A                       Precedence        : 4
Class Fwding State : Down
Flags              : SdpOperDown SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
Mac Move           : Ukwn                      Blockable Level   : Unknown
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

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

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

Associated LSP LIST :
No LSPs Associated

-------------------------------------------------------------------------------
APIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Admin Concat Limit : 1                         Oper Concat Limit : 1
Peer Concat Limit  : n/a                       Max Concat Delay  : 400
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
```

```
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/4/14.1:55
-------------------------------------------------------------------------------
Service Id         : 5
SAP                : 1/4/14.1:55              Encap             : atm
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown L2OperDown
Multi Svc Site     : None
Last Status Change : 03/11/2008 19:58:19
Last Mgmt Change   : 04/01/2008 17:03:42
Sub Type           : regular

Admin MTU          : 1572                     Oper MTU          : 1572
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
tod-suite          : None                     qinq-pbit-marking : both
Egr Agg Rate Limit : max
Endpoint           : N/A

Acct. Pol          : None                     Collect Stats     : Disabled

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        n/a
Off. HiPrio           : 30                       n/a
Off. LowPrio          : n/a                      n/a

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        n/a
Dro. LowPrio          : n/a                      n/a
For. InProf           : 15                       15
For. OutProf          : 15                       15

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        n/a
Dro. OutProf          : n/a                      n/a
For. InProf           : 30                       30
For. OutProf          : n/a                      n/a
-------------------------------------------------------------------------------
```

```
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 30                      n/a
Off. LoPrio           : n/a                     n/a
Dro. HiPrio           : 0                       n/a
Dro. LoPrio           : n/a                     n/a
For. InProf           : 15                      15
For. OutProf          : 15                      15

Egress Queue 1
For. InProf           : 30                      30
For. OutProf          : n/a                     n/a
Dro. InProf           : 0                       n/a
Dro. OutProf          : n/a                     n/a
-------------------------------------------------------------------------------
ATM SAP Configuration Information
-------------------------------------------------------------------------------
Ingress TD Profile : 1                        Egress TD Profile : 1
Alarm Cell Handling: Enabled
OAM Termination    : Disabled                 Periodic Loopback : Disabled

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:ALU-A>show>service#
```

Sample Output (Cpipe Service)

```
================================
*A:ALU-A>show>service# id 51 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 51
Service Type      : Cpipe               VLL Type          : CESoPSN
Name              : (Not Specified)
Description       : Cpipe for service 51
Customer Id       : 2                   Creation Origin   : manual
Last Status Change: 03/11/2015 19:58:19
Last Mgmt Change  : 03/31/2015 20:41:13
Admin State       : Down                Oper State        : Down
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 51:51  -(138.120.38.1)
-------------------------------------------------------------------------------
Description        : (Not Specified)
SDP Id             : 51:51                   Type              : Spoke
```

```
Split Horiz Grp    : (Not Specified)
VC Type            : CESoPSN                    VC Tag            : 0
Admin Path MTU     : 0                          Oper Path MTU     : 0
Delivery           : MPLS
Far End            : 138.120.38.1
Tunnel Far End     : n/a                        LSP Types         : RSVP
Admin State        : Up                         Oper State        : Down
Acct. Pol          : None                       Collect Stats     : Disabled
Ingress Label      : 0                          Egress Label      : 0
Ing Mac Fltr-Id    : n/a                        Egr Mac Fltr-Id   : n/a
Ing IP Fltr-Id     : n/a                        Egr IP Fltr-Id    : n/a
Admin ControlWord  : Preferred                  Oper ControlWord  : True
Admin BW(Kbps)     : 0                          Oper BW(Kbps)     : 0
Last Status Change : 03/11/2015 19:58:19        Signaling         : TLDP
Last Mgmt Change   : 03/31/2015 20:41:13
Endpoint           : N/A                        Precedence        : 4
PW Status Sig      : Enabled
Class Fwding State : Down
Flags              : SdpOperDown SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
Mac Move           : Ukwn                       Blockable Level   : Unknown
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

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

KeepAlive Information :
Admin State        : Disabled                   Oper State        : Disabled
Hello Time         : 100                        Hello Msg Len     : 0
Max Drop Count     : 3                          Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 1579587                    I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 142162830                  I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 1607658                    E. Fwd. Octets    : 144689220
E. Dro. Pkts.      : 0

-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
Lsp Name           : static_33
Admin State        : Up                         Oper State        : Up
Time Since Last Tr*: 01h53m21s


-------------------------------------------------------------------------------
CPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Local Bit-rate     : 10                         Peer Bit-rate     : n/a
Local Payload Size : 160                        Peer Payload Size : n/a
Local Jitter Buffer: 32                         Peer Jitter Buffer: n/a
Local Asym Delay   : enabled                    Peer Asym Delay   : n/a
```

```
Local Sig Pkts     : No Sig.                  Peer Sig Pkts     : No Sig.
Local CAS Framing  : No CAS                   Peer CAS Framing  : No CAS
Local RTP Header   : Yes                      Peer RTP Header   : No
Local Differential : No                       Peer Differential : No
Local Timestamp    : 0                        Peer Timestamp    : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
* indicates that the corresponding row element may have been truncated.
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
SAP 1/4/5.1
-------------------------------------------------------------------------------
Service Id         : 51
SAP                : 1/4/5.1                  Encap             : cem
Description        : Default sap description for service id 1
Admin State        : Up                       Oper State        : Down
Flags              : AsymDelaySampling
Multi Svc Site     : None
Last Status Change : 03/11/2015 19:58:19
Last Mgmt Change   : 03/31/2015 21:38:50
Sub Type           : regular
Split Horizon Group: (Not Specified)

Admin MTU          : 1572                     Oper MTU          : 1572
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                     qinq-pbit-marking : both
Ing Scheduler Mode : 4-priority               Egr Scheduler Mode: 4-priority
Ing Agg Rate Limit : n/a                      Egr Agg Rate Limit: n/a
Ing Agg cir        : n/a                      Egr Agg cir       : n/a
Endpoint           : N/A

Acct. Pol          : Default                  Collect Stats     : Enabled

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1
Ingress FP QGrp    : (none)                   Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                   Egr Port QGrp Inst: (none)
Shared Q plcy      : n/a                      Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                          Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 1607924                  102907136
Off. LowPrio          : n/a                      n/a
Off. Managed          : 0                        0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        0
```

```
Dro. LowPrio          : n/a                    n/a
For. InProf           : 1607924                102907136
For. OutProf          : 0                      0

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

Ingress Queue 1 (Priority)
Off. Combined         : 0                      0
Dro. HiPrio           : 0                      0
Dro. LowPrio          : n/a                    n/a
For. InProf           : 1607924                102907136
For. OutProf          : 0                      0

Egress Queue 1
For. InProf           : n/a                    n/a
For. OutProf          : n/a                    n/a
Dro. InProf           : n/a                    n/a
Dro. OutProf          : n/a                    n/a
-------------------------------------------------------------------------------
CEM SAP Configuration Information
-------------------------------------------------------------------------------
Endpoint Type         : NxDS0           Bit-rate                  : 1
Payload Size          : 64              Jitter Buffer (ms)        : 32
Jitter Buffer (packets): 4              Playout Threshold (packets): 3
Use RTP Header        : No              Differential              : No
Timestamp Freq        : 0               CAS Framing               : No CAS
Effective PDVT        : +/-16.0 ms

AsymDelayControl      : enabled
RepeatPeriod          : 1               Samples                   : 4K

Cfg Alarm      : stray malformed pktloss overrun underrun
Alarm Status   : pktloss underrun rpktloss
-------------------------------------------------------------------------------
CEM SAP Statistics
-------------------------------------------------------------------------------
                                  Packets          Seconds          Events
Egress Stats
Forwarded                    :   0
Dropped                      :   0
Missing                      :   1152
Reordered Forwarded          :   0
Underrun                     :   723                                   0
Overrun                      :   0                                     0
Misordered Dropped           :   0
Malformed Dropped            :   0
LBit Dropped                 :   0
Error                        :                    0
Severely Error               :                    0
Unavailable                  :                    15
Failure Count                :                                         0
```

```
Jitter Buffer Depth            :   0
Jitter Buffer Sampling Complete:                                        0
Jitter Buffer Adjust           :                                        0

Ingress Stats
Forwarded                      :  1157
Dropped                        :  0

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
```

Sample Output (Epipe Service)

```
A:ALU-1>show>service>id# all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 2
Service Type      : Epipe
Name              : epipe1
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 07/13/2009 18:50:40
Last Mgmt Change  : 07/13/2009 18:50:40
Admin State       : Down                Oper State        : Down
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 1:11  -(10.10.10.10)
-------------------------------------------------------------------------------
SDP Id             : 1:11                     Type              : Spoke
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 0
Far End            : 10.10.10.10              Delivery          : MPLS

Admin State        : Up                       Oper State        : Down
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 0                        Egress Label      : 0
Ing mac Fltr       : n/a                      Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                      Egr ip Fltr       : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Admin BW(Kbps)     : 0                        Oper BW(Kbps)     : 0
Last Status Change : 07/13/2009 18:50:40      Signaling         : TLDP
Last Mgmt Change   : 07/13/2009 18:50:40      Force Vlan-Vc     : Disabled
Endpoint           : N/A                      Precedence        : 4
PW Status Sig      : Enabled
Class Fwding State : Down
Flags              : SvcAdminDown SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
```

```
Time to RetryReset : 476014240 seconds        Retries Left      : -1
Mac Move           : Ukwn                     Blockable Level   : Unknown
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

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

Statistics           :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0
-------------------------------------------------------------------------------
Eth-Cfm Configuration Information
-------------------------------------------------------------------------------
Md-index           : 1                        Direction         : Down
Ma-index           : 1                        Admin             : Disabled
MepId              : 2                        CCM-Enable        : Disabled
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : None
Mac Address        : a4:58:ff:00:00:00        CcmLtmPriority    : 7
CcmTx              : 0                        CcmSequenceErr    : 0
LbRxReply          : 0                        LbRxBadOrder      : 0
LbRxBadMsdu        : 0                        LbTxReply         : 0
LbNextSequence     : 1                        LtNextSequence    : 1
LtRxUnexplained    : 0

Associated LSP LIST :
No LSPs Associated

-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/5/1
-------------------------------------------------------------------------------
Service Id         : 2
SAP                : 1/5/1                    Encap             : null
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Multi Svc Site     : None
Last Status Change : 07/13/2009 18:50:40
Last Mgmt Change   : 07/13/2009 18:50:40
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Hold Meps Up       : Disabled
LLF Admin State    : Down                     LLF Oper State    : Clear
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                      qinq-pbit-marking : both
Egr Agg Rate Limit : max
Endpoint           : N/A
Q Frame-Based Acct : Disabled
Vlan-translation   : None

Acct. Pol          : None                      Collect Stats     : Disabled

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                         Egress qos-policy : 1
Shared Q plcy      : n/a                       Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A
                        Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 0                        0
Off. LowPrio          : 0                        0

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

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

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

Egress Queue 1
For. InProf           : 0                        0
For. OutProf          : 0                        0
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
-------------------------------------------------------------------------------
Eth-Cfm Configuration Information
-------------------------------------------------------------------------------
Md-index           : 1                         Direction         : Down
Ma-index           : 1                         Admin             : Disabled
```

```
MepId              : 1                          CCM-Enable        : Disabled
LowestDefectPri    : macRemErrXcon              HighestDefect     : none
Defect Flags       : None
Mac Address        : 00:00:00:00:00:00
CcmTx              : 0                          CcmSequenceErr    : 0
LbRxReply          : 0                          LbRxBadOrder      : 0
LbRxBadMsdu        : 0                          LbTxReply         : 0
LbNextSequence     : 1                          LtNextSequence    : 1
LtRxUnexplained    : 0
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
A:ALU-1>show>service>id#
```

Sample Output (TDM SAP-to-Ethernet SAP MEF 8 Epipe Service)

```
*A:7705:Dut-A>config>service>epipe# show service id 1 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1
Service Type      : Epipe
Name              : (Not Specified)
Description       : Default epipe description for service id 1
Customer Id       : 1
Last Status Change: 10/12/2012 13:42:38
Last Mgmt Change  : 10/12/2012 13:42:34
Admin State       : Up                 Oper State        : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 2                  SDP Bind Count    : 0
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/2/4:1
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/2/4:1                Encap             : q-tag
Description        : Default sap description for service id 1
Admin State        : Up                     Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 10/12/2012 13:32:42
Last Mgmt Change   : 10/12/2012 13:42:34
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                 QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Hold Meps Up       : Disabled
```

```
Admin MTU          : 1518                    Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                    qinq-pbit-marking : n/a
Ing Scheduler Mode : 4-priority              Egr Scheduler Mode: 4-priority


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

Acct. Pol          : None                    Collect Stats     : Disabled
Application Profile: None


-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 5891091                 907228014

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 5891091                 907228014

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 5891109                 930795222
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 5891091                 907228014
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 5891091                 907228014

Egress Queue 1
For. InProf           : 5891109                 930795222
For. OutProf          : 0                       0
```

```
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0

-------------------------------------------------------------------------------
SAP 1/1/1.1
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/1/1.1                  Encap             : cem
Description        : Default sap description for service id 1
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 10/12/2012 13:42:38
Last Mgmt Change   : 10/12/2012 13:42:34
Sub Type           : regular
Split Horizon Group: (Not Specified)

Admin MTU          : 1514                     Oper MTU          : 1514
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                     qinq-pbit-marking : n/a
Ing Scheduler Mode : 4-priority               Egr Scheduler Mode: 4-priority


Endpoint           : N/A
Vlan-translation   : None

Acct. Pol          : None                     Collect Stats     : Disabled
Application Profile: None


-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 5886512                  753473536
Off. LowPrio          : n/a                      n/a

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        0
Dro. LowPrio          : n/a                      n/a
For. InProf           : 5886512                  753473536
For. OutProf          : 0                        0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : n/a                      n/a
Dro. OutProf          : n/a                      n/a
For. InProf           : n/a                      n/a
For. OutProf          : n/a                      n/a
```

```
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Priority)
Off. HiPrio           : 5886512                 753473536
Off. LoPrio           : n/a                     n/a
Dro. HiPrio           : 0                       0
Dro. LoPrio           : n/a                     n/a
For. InProf           : 5886512                 753473536
For. OutProf          : 0                       0

Egress Queue 1
For. InProf           : n/a                     n/a
For. OutProf          : n/a                     n/a
Dro. InProf           : n/a                     n/a
Dro. OutProf          : n/a                     n/a
-------------------------------------------------------------------------------
CEM SAP Configuration Information
-------------------------------------------------------------------------------
Endpoint Type          : NxDS0            Bit-rate                   : 16
Payload Size           : 128              Jitter Buffer (ms)         : 5
Jitter Buffer (packets): 5                Playout Threshold (packets): 3
Use RTP Header         : No               Differential               : No
Timestamp Freq         : 0                CAS Framing                : No CAS
Effective PDVT         : +/-2.0 ms
Local Ecid             : 2                Remote Ecid                : 2
Remote MacAddr         : 10:00:50:00:00:02

Cfg Alarm      : stray malformed pktloss overrun underrun rpktloss rfault rrdi
Alarm Status   :
-------------------------------------------------------------------------------
CEM SAP Statistics
-------------------------------------------------------------------------------
                          Packets           Seconds           Events
Egress Stats
Forwarded             :  5907374
Dropped               :  0
Missing               :  0
Reordered Forwarded   :  0
Underrun              :  118                                  4
Overrun               :  0                                    0
Misordered Dropped    :  0
Malformed Dropped     :  0
LBit Dropped          :  0
Error                 :                    1
Severely Error        :                    0
Unavailable           :                    0
Failure Count         :                                       0
Jitter Buffer Depth   :  2

Ingress Stats
Forwarded             :  5907402
Dropped               :  0

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
```

```
No Endpoints found.
===============================================================================
*A:7705:Dut-A>config>service>epipe#
```

Sample Output (Fpipe Service)

```
================================
*A:ALU-A>show>service# id 10021 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 10021
Service Type      : Fpipe               VLL Type          : FRDLCI
Name              : fpipe10021
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 07/07/2011 17:04:33
Last Mgmt Change  : 07/07/2011 17:03:06
Admin State       : up                  Oper State        : up
MTU               : 1508
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 1:1021  -(100.100.100.148)
-------------------------------------------------------------------------------
Description        : (Not Specified)
SDP Id             : 1:1021                  Type              : Spoke
Split Horiz Grp     (Not Specified)
VC Type            : FRDLCI                  VC Tag            : 0
Admin Path MTU     : 0                       Oper Path MTU     : 1550
Far End            : 100.100.100.148         Delivery          : LDP
Admin State        : Up                      Oper State        : Up
Acct. Pol          : None                    Collect Stats     : Disabled
Ingress Label      : 131056                  Egress Label      : 131056
Ing mac Fltr       : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                     Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                     Egr ipv6 Fltr       : n/a
Admin ControlWord  : Preferred               Oper ControlWord  : True
Admin BW(Kbps)     : 0                       Oper BW(Kbps)     : 0
Last Status Change : 07/07/2011 17:04:33     Signaling         : TLDP
Last Mgmt Change   : 07/07/2011 17:03:06
Endpoint           : N/A                     Precedence        : 4
Class Fwding State : Down
Flags              : None
Mac Move           : Ukwn                    Blockable Level   : Unknown
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : pwe3ControlWord mplsRouterAlertLabel
KeepAlive Information :
Admin State        : Disabled                Oper State        : Disabled
Hello Time         : 10                      Hello Msg Len     : 0
Max Drop Count     : 3                       Hold Down Time    : 10
Statistics           :
```

```
I. Fwd. Pkts.      : 1413                    I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 737586                  I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 551723                  E. Fwd. Octets    : 287999406
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/4/2:21
-------------------------------------------------------------------------------
Service Id         : 10021
SAP                : 1/4/2:21                Encap             : frRel
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 07/06/2011 18:06:20
Last Mgmt Change   : 07/06/2011 18:00:35
Sub Type           : regular
Split Horiz Grp     (Not Specified)
Admin MTU          : 1600                    Oper MTU          : 1600
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
tod-suite          : None                    qinq-pbit-marking : n/a
Ing Agg Rate Limit : max                     Egr Agg Rate Limit : max
Endpoint           : N/A
FRF-12             : Disabled
Acct. Pol          : None                    Collect Stats     : Disabled

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A
                          Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 0                        0
Off. LowPrio          : 574548                   286124904
Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        0
Dro. LowPrio          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 574548                   286124904
Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 1540                     766920
-------------------------------------------------------------------------------
```

```
Sap per Queue stats
-------------------------------------------------------------------------------
                      Packets                 Octets
Ingress Queue 1 (Priority)
Off. HiPrio           : 0                     0
Off. LoPrio           : 574548                286124904
Dro. HiPrio           : 0                     0
Dro. LoPrio           : 0                     0
For. InProf           : 0                     0
For. OutProf          : 574548                286124904
Egress Queue 1
For. InProf           : 0                     0
For. OutProf          : 1540                  766920
Dro. InProf           : 0                     0
Dro. OutProf          : 0                     0
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
```

Sample Output (Hpipe Service)

```
A:7705:Dut-C>config>service# show service id 501 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id      : 501
Service Type    : Hpipe                         VLL Type       : Hpipe
Name            : hpipe501
Description     : Default Hpipe description for service id 501
Customer Id       : 500
Last Status Change: 07/11/2011 14:16:53
Last Mgmt Change  : 07/11/2011 14:16:14
Admin State       : Up                          Oper State     : Up
MTU               : 1514
Vc Switching      : False
SAP Count         : 1                           SDP Bind Count : 1
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Sdp Id 999:501 -(10.10.10.2)
-------------------------------------------------------------------------------
Description       : RSVP_SdpToDut-B
SDP Id            : 999:501                     Type             : Spoke
Split Horiz Grp   : (Not Specified)
VC Type           : Hpipe                       VC Tag           : 0
Admin Path MTU    : 0                           Oper Path MTU    : 1550
Far End           : 10.10.10.2                  Delivery         : MPLS
Admin State       : Up                          Oper State       : Up
Acct. Pol         : None                        Collect Stats    : Disabled
Ingress Label     : 131069                      Egress Label     : 131069
Ing mac Fltr      : n/a                         Egr mac Fltr     : n/a
Ing ip Fltr       : n/a                         Egr ip Fltr      : n/a
Admin ControlWord : Not Preferred               Oper ControlWord : False
Admin BW(Kbps)    : 0                           Oper BW(Kbps)    : 0
Last Status Change: 07/11/2011 14:16:35         Signaling        : TLDP
```

```
Last Mgmt Change  : 07/11/2011 14:16:14
Endpoint          : N/A                          Precedence       : 4
Class Fwding State: Down
Flags             : None
Local Pw Bits     : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : mplsRouterAlertLabel
KeepAlive Information :
Admin State       : Disabled                     Oper State      : Disabled
Hello Time        : 10                           Hello Msg Len   : 0 Max
Drop Count        : 3                            Hold Down Time  : 10
Statistics :
I. Fwd. Pkts.     : 0                            I. Dro. Pkts.   : 0
I. Fwd. Octs.     : 0                            I. Dro. Octs.   : 0
E. Fwd. Pkts.     : 0                            E. Fwd. Octets  : 0
Associated LSP LIST :
Lsp Name          : LSPToDut-B
Admin State       : Up                           Oper State        : Up
Time Since Last Tr*: 00h01m53s
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
SAP 1/10/9.1
-------------------------------------------------------------------------------
Service Id : 501
SAP : 1/10/9.1                                   Encap             : hdlc
Description : Default sap description for service id 501
Admin State : Up                                 Oper State        : Up
Flags : None
Multi Svc Site : None
Last Status Change : 07/11/2011 14:16:53
Last Mgmt Change : 07/11/2011 14:16:14
Sub Type : regular
Split Horizon Group: (Not Specified)
Admin MTU : 1514                                 Oper MTU          : 1514
Ingr IP Fltr-Id : n/a                            Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id : n/a                           Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id : n/a                          Egr IPv6 Fltr-Id  : n/a
tod-suite : None                                 qinq-pbit-marking : n/a
Ing Agg Rate Limit : max                         Egr Agg Rate Limit: max
Endpoint : N/A
Acct. Pol : None                                 Collect Stats     : Disabled
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                           Egress qos-policy : 1
Shared Q plcy     : n/a                          Multipoint shared : Disabled
I. Sched Pol      : (Not Specified)
E. Sched Pol      : (Not Specified)
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Sap Statistics
```

```
-------------------------------------------------------------------------------
Last Cleared Time  : N/A
Packets Octets
Forwarding Engine Stats
Dropped            : 0 0
Off. HiPrio        : 0 0
Off. LowPrio       : 0 0
Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio        : 0 0
Dro. LowPrio       : 0 0
For. InProf        : 0 0
For. OutProf       : 0 0
Queueing Stats(Egress QoS Policy 1)
Dro. InProf        : 0 0
Dro. OutProf       : 0 0
For. InProf        : 0 0
For. OutProf       : 0 0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
Packets Octets
Ingress Queue 1 (Priority)
Off. HiPrio        : 0 0
Off. LoPrio        : 0 0
Dro. HiPrio        : 0 0
Dro. LoPrio        : 0 0
For. InProf        : 0 0
For. OutProf       : 0 0

Egress Queue 1
For. InProf        : 0 0
For. OutProf       : 0 0
Dro. InProf        : 0 0
Dro. OutProf       : 0 0
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.

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

Sample Output (Ipipe Service)

```
*A:ALU-A# show service id 1301 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1301
Service Type      : Ipipe
Name              : epipe1
Description       : Default ipipe description for service id 1301
Customer Id       : 1
Last Status Change: 01/20/2009 16:44:14
Last Mgmt Change  : 01/20/2009 16:02:02
Admin State       : Up                  Oper State        : Up
MTU               : 1514
Vc Switching      : False
```

```
SAP Count         : 1                    SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 123:1301  -(10.20.1.3)
-------------------------------------------------------------------------------
Description     : Default sdp description
SDP Id            : 123:1301                Type              : Spoke
VC Type           : Ipipe                   VC Tag            : 0
Admin Path MTU    : 0                       Oper Path MTU     : 1516
Far End           : 10.20.1.3               Delivery          : LDP

Admin State       : Up                      Oper State        : Up
Acct. Pol         : None                    Collect Stats     : Disabled
Ingress Label     : 131069                  Egress Label      : 131069
Ing mac Fltr      : n/a                     Egr mac Fltr      : n/a
Ing ip Fltr       : n/a                     Egr ip Fltr       : n/a
Admin ControlWord : Not Preferred           Oper ControlWord  : False
Admin BW(Kbps)    : 0                       Oper BW(Kbps)     : 0
Last Status Change : 01/20/2009 16:05:49    Signaling         : TLDP
Last Mgmt Change  : 01/20/2009 16:02:02
Endpoint          : N/A                     Precedence        : 4
Class Fwding State : Down
Flags             : None
Time to RetryReset : 1 seconds              Retries Left      : 213236003
Mac Move          : Ukwn                    Blockable Level   : Unknown
Local Pw Bits     : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits  : lspPing
Peer Vccv CC Bits  : mplsRouterAlertLabel

Ipipe Sdp Bind Info :
IpipeSdpBindCeIpAd*: 88.1.10.4

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.     : 600                     I. Dro. Pkts.     : 0
I. Fwd. Octs.     : 60000                   I. Dro. Octs.     : 0
E. Fwd. Pkts.     : 21817053                E. Fwd. Octets    : 1919900664
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/2/8:11
-------------------------------------------------------------------------------
Service Id        : 1301
SAP               : 1/2/8:11                Encap             : q-tag
Description       : Default sap description for service id 1301
Admin State       : Up                      Oper State        : Up
```

```
Flags              : None
Multi Svc Site     : None
Last Status Change : 01/20/2009 16:44:14
Last Mgmt Change   : 01/21/2009 16:31:04
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                    QinQ Ethertype    : 0x8100

Admin MTU          : 1572                      Oper MTU          : 1572
Ingr IP Fltr-Id    : n/a                       Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                       Egr Mac Fltr-Id   : n/a
tod-suite          : None                      qinq-pbit-marking : both
Egr Agg Rate Limit : max
Endpoint           : N/A
Q Frame-Based Acct : Disabled

Acct. Pol          : Default                   Collect Stats     : Enabled
Ce IP Address      : 88.1.10.3
SAP MAC Address    : 00:1a:f0:bd:ab:b0         Mac Refresh Inter*: 14400

-------------------------------------------------------------------------------
Ipipe SAP ARP Entry Info
-------------------------------------------------------------------------------
88.1.10.3       00:00:15:b9:6b:73 dynamic 03h52m50s

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 13                        Egress qos-policy : 13
Shared Q plcy      : n/a                       Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : 01/21/2009 14:19:23

                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 19961282                1556979996
Off. LowPrio          : 1840167                 143533026
Queueing Stats(Ingress QoS Policy 13)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 10730245                836959110
For. OutProf          : 11071204                863553912

Queueing Stats(Egress QoS Policy 13)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 600                     46800
-------------------------------------------------------------------------------
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 2 (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 3 (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 4 (Unicast) (Priority)
Off. HiPrio           : 6582217               513412926
Off. LoPrio           : 0                     0
Dro. HiPrio           : 0                     0
Dro. LoPrio           : 0                     0
For. InProf           : 4932647               384746466
For. OutProf          : 1649570               128666460

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

Egress Queue 2
For. InProf           : 0                     0
For. OutProf          : 200                   15600
Dro. InProf           : 0                     0
Dro. OutProf          : 0                     0

Egress Queue 3
For. InProf           : 0                     0
For. OutProf          : 200                   15600
Dro. InProf           : 0                     0
Dro. OutProf          : 0                     0

Egress Queue 4
For. InProf           : 0                     0
For. OutProf          : 200                   15600
Dro. InProf           : 0                     0
Dro. OutProf          : 0                     0

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:ALU-A#
```

Sample Output (Ipipe Frame Relay to Ethernet Service)

```
*A:7705:Dut-C# show service id 2 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 2
Service Type      : Ipipe
Name              : ipipe2
Description       : Default ipipe description for service id 2
Customer Id       : 1001
Last Status Change: 07/08/2011 07:29:50
Last Mgmt Change  : 07/08/2011 07:29:37
Admin State       : Up                  Oper State        : Up
MTU               : 1500
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
CE Addr Discovery : disabled
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 888:2  -(10.10.10.2)
-------------------------------------------------------------------------------
Description     : GRE_SdpToDut-B
SDP Id            : 888:2                 Type              : Spoke
Split Horiz Grp   : (Not Specified)
VC Type           : Ipipe                 VC Tag            : 0
Admin Path MTU    : 0                     Oper Path MTU     : 1530
Far End           : 10.10.10.2            Delivery          : GRE

Admin State       : Up                    Oper State        : Up
Acct. Pol         : None                  Collect Stats     : Disabled
Ingress Label     : 131070                Egress Label      : 131070
Ing mac Fltr      : n/a                   Egr mac Fltr      : n/a
Ing ip Fltr       : n/a                   Egr ip Fltr       : n/a
Admin ControlWord : Not Preferred         Oper ControlWord  : False
Admin BW(Kbps)    : 0                     Oper BW(Kbps)     : 0
Last Status Change : 07/08/2011 07:29:38  Signaling          : TLDP
Last Mgmt Change  : 07/08/2011 07:29:37
Endpoint          : N/A                   Precedence        : 4
Class Fwding State : Down
Flags             : None
Local Pw Bits     : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : mplsRouterAlertLabel

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

Statistics          :
I. Fwd. Pkts.     : 1022                  I. Dro. Pkts.     : 0
I. Fwd. Octs.     : 1565064               I. Dro. Octs.     : 0
E. Fwd. Pkts.     : 18648                 E. Fwd. Octets    : 1566432
```

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

-------------------------------------------------------------------------------
IPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Configured CE IP Addr  : 12.12.20.1           Peer CE IP Addr   : n/a

-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/10/9.1:101
-------------------------------------------------------------------------------
Service Id         : 2
SAP                : 1/10/9.1:101             Encap             : frRel
Description        : Default sap description for service id 2
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 07/08/2011 07:29:50
Last Mgmt Change   : 07/08/2011 07:29:37
Sub Type           : regular
Split Horizon Group: (Not Specified)

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

FRF-12             : Disabled


Acct. Pol          : None                     Collect Stats     : Disabled

-------------------------------------------------------------------------------
Ipipe SAP Configuration Information
-------------------------------------------------------------------------------
Ce IP Address      : 12.12.10.1
Configured CE IP   : 12.12.10.1               Discovered CE IP  : n/a

-------------------------------------------------------------------------------
Ipipe SAP ARP Entry Info
-------------------------------------------------------------------------------
No Ipipe SAP ARP entries
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
```

```
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 18648                   783216

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 18648                   783216

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 1020                    1521840
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 18648                   783216
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 18648                   783216

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 1020                    1521840
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:7705:Dut-C#
```

Sample Output (Ipipe cHDLC to Ethernet Service)

```
*A:7705:Dut-C# show service id 2 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 2
Service Type      : Ipipe
Name              : ipipe2
Description       : Default ipipe description for service id 2
```

```
Customer Id       : 1001
Last Status Change: 07/08/2011 07:50:55
Last Mgmt Change  : 07/08/2011 07:50:44
Admin State       : Up                  Oper State        : Up
MTU               : 1500
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1
CE Addr Discovery : disabled
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 333:2  -(10.10.10.2)
-------------------------------------------------------------------------------
Description     : RSVP_SdpToDut-B
SDP Id            : 333:2                     Type              : Spoke
Split Horiz Grp   : (Not Specified)
VC Type           : Ipipe                     VC Tag            : 0
Admin Path MTU    : 0                         Oper Path MTU     : 1550
Far End           : 10.10.10.2                Delivery          : MPLS

Admin State       : Up                        Oper State        : Up
Acct. Pol         : None                      Collect Stats     : Disabled
Ingress Label     : 131069                    Egress Label      : 131068
Ing mac Fltr      : n/a                       Egr mac Fltr      : n/a
Ing ip Fltr       : n/a                       Egr ip Fltr       : n/a
Admin ControlWord : Not Preferred             Oper ControlWord  : False
Admin BW(Kbps)    : 0                         Oper BW(Kbps)     : 0
Last Status Change : 07/08/2011 07:50:47      Signaling         : TLDP
Last Mgmt Change  : 07/08/2011 07:50:44
Endpoint          : N/A                       Precedence        : 4
Class Fwding State : Down
Flags             : None
Local Pw Bits     : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : mplsRouterAlertLabel

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

Statistics          :
I. Fwd. Pkts.     : 125                       I. Dro. Pkts.     : 0
I. Fwd. Octs.     : 31832                     I. Dro. Octs.     : 0
E. Fwd. Pkts.     : 489                       E. Fwd. Octets    : 31296

Associated LSP LIST :
Lsp Name          : LSPToDut-B
Admin State       : Up                        Oper State        : Up
Time Since Last Tr*: 00h03m46s

-------------------------------------------------------------------------------
IPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Configured CE IP Addr  : 12.12.20.1           Peer CE IP Addr   : n/a
```

```
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/10/9.1
-------------------------------------------------------------------------------
Service Id         : 2
SAP                : 1/10/9.1                 Encap             : cisco-hdlc
Description        : Default sap description for service id 2
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 07/08/2011 07:50:55
Last Mgmt Change   : 07/08/2011 07:50:44
Sub Type           : regular
Split Horizon Group: (Not Specified)

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

Acct. Pol          : None                     Collect Stats     : Disabled

-------------------------------------------------------------------------------
Ipipe SAP Configuration Information
-------------------------------------------------------------------------------
Ce IP Address      : 12.12.10.1
Configured CE IP   : 12.12.10.1               Discovered CE IP  : n/a

-------------------------------------------------------------------------------
Ipipe SAP ARP Entry Info
-------------------------------------------------------------------------------
No Ipipe SAP ARP entries

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A
                          Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 0                        0
Off. LowPrio          : 489                      20538
```

```
Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                     0
Dro. LowPrio          : 0                     0
For. InProf           : 0                     0
For. OutProf          : 489                   20538


Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                     0
Dro. OutProf          : 0                     0
For. InProf           : 0                     0
For. OutProf          : 123                   28782
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets               Octets

Ingress Queue 1 (Priority)
Off. HiPrio           : 0                     0
Off. LoPrio           : 489                   20538
Dro. HiPrio           : 0                     0
Dro. LoPrio           : 0                     0
For. InProf           : 0                     0
For. OutProf          : 489                   20538

Egress Queue 1
For. InProf           : 0                     0
For. OutProf          : 123                   28782
Dro. InProf           : 0                     0
Dro. OutProf          : 0                     0

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
===============================================================================
*A:7705:Dut-C#
```

base

Syntax **base**

Context show>service>id

Description This command displays basic information about the service specified by the ID, including service type, description, SAPs, SDPs, and SAP aggregation group (if present).

Output The following are examples of service-id base information. Table 45 describes the fields.

Sample Output (Apipe ATMVcc Base)

```
==============================
*A:ALU-12#  show service id 701 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 701
Service Type      : Apipe               VLL Type          : ATMVCC
Name              : apipe701
Description       : Default apipe description for service id 701
Customer Id       : 1
Last Status Change: 02/10/2008 03:30:03
Last Mgmt Change  : 02/10/2008 03:35:10
Admin State       : Up                  Oper State        : Down
MTU               : 1508
Vc Switching      : False
SAP Count         : 1                   SDP Bind Count    : 1

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                      Type         AdmMTU  OprMTU  Adm     Opr
-------------------------------------------------------------------------------
sap:1/1/9.1:10/50               atm          1572    1572    Up      Down
sdp:101:701 S(10.20.1.3)        n/a          0       1514    Up      Up
-------------------------------------------------------------------------------
[<sap-id>] indicates a Managed SAP
===============================================================================
```

Sample Output (Apipe with a SAP Aggregation Group)

```
*B:A:179_121# show service id 300 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 300
Service Type      : Apipe               VLL Type          : ATMVCC
Name              : apipe300
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 08/24/2011 19:47:16
Last Mgmt Change  : 08/24/2011 22:14:28
Admin State       : Up                  Oper State        : Up
MTU               : 1508
Vc Switching      : False
SAP Count         : 3                   SDP Bind Count    : 1

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm  Opr
-------------------------------------------------------------------------------
sap-agg-grp2
    sap:1/1/9.1:3/300                   atm          1524    1524    Up   Up
    sap:1/1/9.1:3/301                   atm          1524    1524    Up   Up
    sap:1/1/9.1:3/302                   atm          1524    1524    Up   Up
sdp:200:300 S(10.0.0.123)               n/a          1524    1524    Up   Up
===============================================================================
```

Table 45: Show Service-ID Base Output Fields

Label	Description
Service Basic Information	
Service Id	Identifies the service by its ID number
Service Type	Specifies the type of service
VLL Type	Specifies the VLL type
Name	Specifies the optional configured service name
Description	Displays generic information about the service
Customer Id	Identifies the customer by its ID number
Last Status Change	Displays the date and time of the most recent status change to this service
Last Mgmt Change	Displays the date and time of the most recent management-initiated change to this service
Admin State	Specifies the desired state of the service
Oper State	Specifies the operating state of the service

Table 45: Show Service-ID Base Output Fields (Continued)

Label	Description
MTU	Specifies the service MTU
SAP Count	Displays the number of SAPs specified for this service
SDP Bind Count	Displays the number of SDPs bound to this service
Service Access & Destination Points	
Identifier	Lists the SAP, SDP, and SAP group identifier (if present)
Type	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP
AdmMTU	Specifies the desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end edge services router (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
Adm	Indicates the operating state of the SAP or SDP
Opr	Indicates the operating state of the SAP or SDP

egress-label

Syntax **egress-label** *start-label* [*end-label*]

Context show>service

Description This command displays 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 this range of labels are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *start-label* — indicates the starting egress 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 to 131071

end-label — indicates the ending egress label value for which to display services using the label range

Values 2049 to 131071

Default the *start-label* value

Output The following output is an example of service egress-label information, and Table 46 describes the fields.

Sample Output

```
*A:ALU-12# show service egress-label 0 131071
==============================================================================
Martini Service Labels
==============================================================================
Svc Id     Sdp Binding        Type  I.Lbl                E.Lbl
------------------------------------------------------------------------------
1          101:1              Spok  131049               0
103        101:103            Spok  131067               131067
104        301:104            Spok  131066               131067
105        501:105            Spok  131065               131068
303        101:303            Spok  131064               131066
304        301:304            Spok  131063               131064
305        501:305            Spok  131062               131065
701        101:701            Spok  131059               131064
702        101:702            Spok  131058               131063
703        501:703            Spok  131057               131064
704        501:704            Spok  131056               131063
705        301:705            Spok  131055               131062
706        301:706            Spok  131054               131061
805        201:805            Spok  131053               131062
806        201:806            Spok  131052               131061
807        401:807            Spok  131051               131060
808        401:808            Spok  131050               131059
903        201:903            Spok  131061               131065
904        401:904            Spok  131060               131063
------------------------------------------------------------------------------
Number of Bindings Found : 19
------------------------------------------------------------------------------
```

Table 46: Show Service Egress Label Output Fields

Label	Description
Svc Id	Identifies the service
Sdp Binding	Identifies the SDP
Type	Specifies the SDP binding type (for example, spoke)
I. Lbl	Displays the VC label used by the far-end device to send packets to this device in this service by the SDP
E. Lbl	Displays the VC label used by this device to send packets to the far-end device in this service by the SDP

Table 46: Show Service Egress Label Output Fields (Continued)

Label	Description
Number of bindings found	Indicates the total number of SDP bindings that exist within the specified egress label range

id

Syntax **id** *service-id*

Context show>service

Description This command displays information for a particular service-id.

Parameters *service-id* — identifies the service in the domain

Values 1 to 2147483690 or *svc-name*

ingress-label

Syntax **ingress-label** *start-label* [*end-label*]

Context show>service

Description This command displays services using the range of ingress labels.

If only the mandatory *start-label* parameter is specified, only services using the specified label are displayed.

If both *start-label* and *end-label* parameters are specified, the services using this range of labels are displayed.

Use the **show router** *vprn-service-id* **ldp bindings** command to display dynamic labels.

Parameters *start-label* — indicates 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 to 131071

end-label — indicates the ending ingress label value for which to display services using the label range

Values 2049 to 131071

Default the *start-label* value

Output The following output is an example of service ingress-label information, and Table 47 describes the fields.

Sample Output

```
*A:ALU-12# show service ingress-label 0
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Binding        Type  I.Lbl        E.Lbl
-------------------------------------------------------------------------------
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 : 4
-------------------------------------------------------------------------------
*A:ALU-12#
```

Table 47: Show Service Ingress Label Output Fields

Label	Description
Svc ID	Identifies the service
SDP Binding	Identifies the SDP
Type	Specifies the SDP binding type (for example, spoke)
I.Lbl	Displays the ingress label used by the far-end device to send packets to this device in this service by the SDP
E.Lbl	Displays 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	Indicates the number of SDP bindings within specified the label range

endpoint

Syntax **endpoint** *endpoint-name*

Context show>service>id

Description This command displays the endpoint configuration status of the active spoke SDP and lists the primary and secondary spoke SDPs used by the service.

Output The following output is an example of service-id endpoint information, and Table 48 describes the fields.

Sample Output

```
*A:7705:Dut-C>show>service>id# endpoint Endpoint_Y

===============================================================================
Service 6 endpoints
===============================================================================
Endpoint name                : Endpoint_Y
Revert time                  : 0
Act Hold Delay               : 0
Ignore Standby Signaling     : false
Suppress Standby Signaling   : true
Tx Active                    : none
Tx Active Up Time            : 0d 00:00:00
Revert Time Count Down       : N/A
Tx Active Change Count       : 0
Last Tx Active Change        : 02/12/2009 19:16:37
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp                    : 6:6 Precedence:0
Spoke-sdp                    : 7:7 Precedence:1
===============================================================================
*A:7705:Dut-C>show>service>id# info
```

Table 48: Service-ID Endpoint Output Fields

Label	Description
Service endpoints	
Endpoint name	Identifies the endpoint
Revert time	Displays the revert time setting for the active spoke SDP
Act Hold Delay	Not applicable
Ignore Standby Signaling	Indicates whether standby signaling is ignored True: standby signaling is ignored False: standby signaling is not ignored
Suppress Standby Signaling	Indicates whether standby signaling is suppressed True: standby signaling is suppressed False: standby signaling is not suppressed
Tx Active	Identifies the actively transmitting spoke SDP
Tx Active Up Time	Indicates the length of time that the active spoke SDP has been up
Revert Time Count Down	Not applicable
Tx Active Change Count	Indicates the number of times that there has been a change of active spoke SDPs

Table 48: Service-ID Endpoint Output Fields (Continued)

Label	Description
Last Tx Active Change	Indicates the date and time when a different spoke SDP became the actively transmitting spoke SDP
Members	
Spoke-sdp	Identifies the primary and secondary spoke SDPs that are associated with this endpoint and shows their precedence value (0 precedence indicates the primary spoke SDP)

labels

Syntax **labels**

Context show>service>id

Description This command displays the labels being used by the service.

Output The following output is an example of service-id labels information, and Table 49 describes the fields.

Sample Output

```
*A:ALU-12# show service id 1 labels
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Binding       Type    I.Lbl        E.Lbl
-------------------------------------------------------------------------------
1          10:1              Spok    0            0
-------------------------------------------------------------------------------
Number of Bound SDPs : 1
-------------------------------------------------------------------------------
*A:ALU-12#
```

Table 49: Service-ID Labels Output Fields

Label	Description
Svc Id	Identifies the service
Sdp Binding	Identifies the SDP bound to the service
Type	Indicates the SDP binding type (for example, spoke)
I. Lbl	Displays the VC label used by the far-end device to send packets to this device in this service by the SDP
E. Lbl	Displays the VC label used by this device to send packets to the far-end device in this service by the SDP

sap

Syntax **sap** *sap-id* [**atm** | **base** | **detail** | **qos** | **sap-stats** | **stats**]

Context show>service>id

Description This command displays information for the SAPs associated with the service.

If no optional parameters are specified, a summary of all associated SAPs is displayed.

Parameters *sap-id* — identifies the SAPs for the service in the form *slot*/*mda*/*port*[.*channel*] for a port or *slot/mda/bridge-id.branch-id* for a SCADA bridge. See Table 39 for a full list of SAP IDs.

atm — displays configuration information about the ATM SAP

base — displays basic information for the SAP or an ATM VCC SAP that is a member of a SAP aggregation group

detail — displays detailed information for the SAP or an ATM VCC SAP that is a member of a SAP aggregation group

qos — displays detailed information for the SAP

This parameter cannot be used on a SAP that is a member of a SAP aggregation group; otherwise, the following error message appears:

```
MINOR: CLI QoS is not configurable for SAPs assigned to SAP
Aggregation Group.
```

See the **show>sap-aggregation-group** command for information about how to show statistics for aggregation group members.

sap-stats — displays detailed information for the SAP

This parameter cannot be used on a SAP that is a member of a SAP aggregation group; otherwise, the following error message appears:

```
MINOR: CLI Statistics are not supported for SAPs assigned to
SAP Aggregation Group.
```

See the **show>sap-aggregation-group** command for information about how to show statistics for aggregation group members.

stats — displays detailed information for the SAP

This parameter cannot be used on a SAP that is a member of a SAP aggregation group; otherwise, the following error message appears:

```
MINOR: CLI Statistics are not supported for SAPs assigned to
SAP Aggregation Group.
```

See the **show>sap-aggregation-group** command for information about how to show statistics for aggregation group members.

Output The Sample Output (Apipe) is an example of service-id SAP information. Table 50 describes the fields for the Apipe example. Table 50 includes additional fields that appear in the sample outputs of the other VLL services, which are listed below and are included after the table:

- Sample Output (Epipe)
- Sample Output (Fpipe)
- Sample Output (Hpipe)
- Sample Output (Ipipe)
- Sample Output (Ipipe with Frame Relay SAP)
- Sample Output (Ipipe with cHDLC SAP)
- Sample Output (Apipe in a SAP aggregation group)

Sample Output (Apipe)

```
*A:ALU-12>show>service>id# sap 1/4/1.1:2 detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 2
SAP                : 1/4/1.1:2               Encap             : atm
Description        : Apipe SAP
Admin State        : Up                      Oper State        : Down
Flags              : PortOperDown L2OperDown
Multi Svc Site     : None
Last Status Change : 04/30/2008 13:55:04
Last Mgmt Change   : 05/07/2008 15:51:51
Sub Type           : regular

Admin MTU          : 1572                    Oper MTU          : 1572
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
tod-suite          : None                    qinq-pbit-marking : both
Egr Agg Rate Limit : max
Endpoint           : N/A

Acct. Pol          : None                    Collect Stats     : Disabled

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       n/a
Off. HiPrio           : 21900                   n/a
Off. LowPrio          : n/a                     n/a
```

```
Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                      n/a
Dro. LowPrio          : n/a                    n/a
For. InProf           : 10950                  10950
For. OutProf          : 10950                  10950

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                      n/a
Dro. OutProf          : n/a                    n/a
For. InProf           : 21900                  21900
For. OutProf          : n/a                    n/a
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 21900                  n/a
Off. LoPrio           : n/a                    n/a
Dro. HiPrio           : 0                      n/a
Dro. LoPrio           : n/a                    n/a
For. InProf           : 10950                  10950
For. OutProf          : 10950                  10950

Egress Queue 1
For. InProf           : 21900                  21900
For. OutProf          : n/a                    n/a
Dro. InProf           : 0                      n/a
Dro. OutProf          : n/a                    n/a
-------------------------------------------------------------------------------
ATM SAP Configuration Information
-------------------------------------------------------------------------------
Ingress TD Profile : 1                      Egress TD Profile : 1
Alarm Cell Handling: Enabled
OAM Termination    : Disabled               Periodic Loopback : Disabled
===============================================================================

*A:ALU-12>show>service>id#
```

Table 50: Service-ID SAP Output Fields

Label	Description
Service Access Points	
Service Id	Identifies the service
SAP	Specifies the ID of the access port where this SAP is defined
Encap	Specifies the encapsulation type for this SAP on the access port
Admin State	Specifies the desired state of the SAP
Oper State	Specifies the operating state of the SAP

Table 50: Service-ID SAP Output Fields (Continued)

Label	Description
Flags	Specifies the conditions that affect the operating status of this SAP Display output includes ServiceAdminDown, PortOperDown, and so on
Last Status Change	Specifies the date and time of the most recent status change to this SAP
Last Mgmt Change	Specifies the date and time of the most recent management-initiated change to this SAP
Dot1Q Ethertype	Identifies the value of the dot1q Ethertype
QinQ Ethertype	Identifies the value of the qinq Ethertype
qinq-pbit-marking	Indicates the qinq P-bit marking for the service: both or top
LLF Admin State	Specifies the Link Loss Forwarding administrative state
LLF Oper State	Specifies the Link Loss Forwarding operational state
Admin MTU	Specifies the desired largest service frame size (in octets) that can be transmitted through this SAP to the far-end router, without requiring the packet to be fragmented
Oper MTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SAP to the far-end router, without requiring the packet to be fragmented
Ingr IP Fltr-Id	Specifies the ingress IP filter policy ID assigned to the SAP
Egr IP Fltr-Id	Specifies the egress IP filter policy ID assigned to the SAP
Ingr Mac Fltr-Id	Specifies the ingress MAC filter policy ID assigned to the SAP
Egr Mac Fltr-Id	Specifies the egress MAC filter policy ID assigned to the SAP
Ing Scheduler Mode	The scheduler mode for the SAP in the access ingress direction: 4-priority or 16-priority
Egr Scheduler Mode	The scheduler mode for the SAP in the access egress direction: 4-priority or 16-priority
Ing Agg Rate Limit	The PIR rate limit in the access ingress direction for the aggregate of the SAP queues
Egr Agg Rate Limit	The PIR rate limit in the access egress direction for the aggregate of the SAP queues
Ing Agg cir	The CIR rate limit in the access ingress direction for the aggregate of the SAP queues

Table 50: Service-ID SAP Output Fields (Continued)

Label	Description
Egr Agg cir	The CIR rate limit in the access egress direction for the aggregate of the SAP queues
Acct. Pol	Specifies the accounting policy applied to the SAP
Collect Stats	Specifies whether accounting statistics are collected on the SAP
QOS	
Ingress qos-policy	Displays the SAP ingress QoS policy ID
Egress qos-policy	Displays the SAP egress QoS policy ID
SAP Statistics	
Last Cleared Time	Displays the date and time that a clear command was issued on statistics
Forwarding Engine Stats	
Dropped	Indicates the number of packets or octets dropped by the forwarding engine
Off. HiPrio	Indicates the number of high-priority packets or octets offered to the forwarding engine
Off. LowPrio	Indicates the number of low-priority packets offered to the forwarding engine
Queueing Stats (Ingress QoS Policy)	
Dro. HiPrio	Indicates the number of high-priority packets or octets discarded, as determined by the SAP ingress QoS policy
Dro. LowPrio	Indicates the number of low-priority packets discarded, as determined by the SAP ingress QoS policy
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP ingress QoS policy
For. OutProf	Indicates the number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP ingress QoS policy
Queueing Stats (Egress QoS Policy)	
Dro. InProf	Indicates the number of in-profile packets or octets discarded, as determined by the SAP egress QoS policy
Dro. OutProf	Indicates the number of out-of-profile packets or octets discarded, as determined by the SAP egress QoS policy

Table 50: Service-ID SAP Output Fields (Continued)

Label	Description
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP egress QoS policy
For. OutProf	Indicates the number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP egress QoS policy
Sap per Queue stats	
Ingress Queue *n*	Specifies the index of the ingress QoS queue of this SAP, where *n* is the index number
Off. HiPrio	Indicates the number of packets or octets of high-priority traffic for the SAP (offered)
Off. LoPrio	Indicates the number or packets or octets of low-priority traffic for the SAP (offered)
Dro. HiPrio	Indicates the number of high-priority traffic packets or octets dropped
Dro. LoPrio	Indicates the number of low-priority traffic packets or octets dropped
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	Indicates the number of out-of-profile packets or octets (rate above CIR) forwarded
Egress Queue *n*	Specifies the index of the egress QoS queue of the SAP, where *n* is the index number
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	Indicates the number of out-of-profile packets or octets (rate above CIR) forwarded
Dro. InProf	Indicates the number of in-profile packets or octets dropped for the SAP
Dro. OutProf	Indicates the number of out-of-profile packets or octets discarded
Eth-Cfm Configuration Information	
Md-index	Displays the value of the MD index
Direction	Displays the direction of the MEP
Ma-index	Displays the value of the MA index
Admin	Displays the administrative state of the MEP (enabled or disabled)
MepId	Displays the MEP-ID

Table 50: Service-ID SAP Output Fields (Continued)

Label	Description
CCM-Enable	Displays the status of the Continuity Check Message (CCM)
LowestDefectPri	Displays a configured value that defects are evaluated against
HighestDefect	Displays the highest defect
Defect Flags	Indicates the defect flags
Mac Address	Displays the MAC address (the MAC address for a spoke SDP is the system MAC address; for a SAP, it is the port MAC address)
CcmLtmPriority	Displays the priority of the CCM Linktrace Message (LTM)
CcmTx	Displays the number of CCM transmissions
CcmSequenceErr	Displays the number of CCM sequence errors
DmrRepliesTx	Displays the number of delay measurement replies transmitted
LmrRepliesTx	Displays the number of loss measurement replies transmitted
Dual-Loss Test	Displays the status of the dual-ended loss measurement test (enabled or disabled)
Dual-Loss Thresh	Displays the frame error threshold beyond which an alarm will be raised. The threshold is expressed as a percentage.
Eth-Ais	Displays the status of the ETH-AIS test (enabled or disabled)
Eth-Ais Rx Ais	Indicates whether any ETH-AIS messages have been received
Eth-Ais Tx Priorit*	Displays the priority value of a transmitted ETH-AIS frame
Eth-Ais Rx Interv*	Indicates the interval of a received ETH-AIS frame
Eth-Ais Tx Interva*	Displays the interval of a transmitted ETH-AIS frame
Eth-Ais Tx Counte*	Displays the number of ETH-AIS frames that have been sent
Eth-Ais Tx Levels	Indicates the MD level of transmitted ETH-AIS frames
Eth-Tst	Indicates the status of the ETH-Test (enabled or disabled)
LbRxReply	Displays the number of received loopback (LB) replies
LbRxBadOrder	Displays the number of LB replies that have been received in the wrong order
LbRxBadMsdu	Displays the number of LB replies that have been received with the wrong destination MAC address (MSDU = MAC Service Data Unit)
LbTxReply	Displays the number of LBRs (loopback replies) transmitted out this MEP

Table 50: Service-ID SAP Output Fields (Continued)

Label	Description
LbNextSequence	Displays the sequence number of the next LB transmission
LtNextSequence	Displays the sequence number of the next Linktrace (LT) message transmitted
LtRxUnexplained	Displays the number of the unexplained Linktrace (LT) messages
ATM SAP Configuration Information	
Ingress TD Profile	The profile ID of the traffic descriptor applied to the ingress SAP
Egress TD Profile	The profile ID of the traffic descriptor applied to the egress SAP
Alarm Cell Handling	Indicates that OAM cells are being processed
OAM Termination	Indicates whether this SAP is an OAM termination point
CEM SAP Configuration Information	
Endpoint Type	Specifies the type of endpoint
Bit-rate	Specifies the number of DS0s or timeslots in the channel group
Payload Size	Specifies the number of octets contained in the payload of a TDM PW packet when the packet is transmitted
Jitter Buffer (ms)	Specifies the size of the receive jitter buffer, expressed in milliseconds
Playout Threshold (packets)	Indicates the number of packet buffers for the playout buffer packets threshold
Use RTP Header	Specifies whether RTP headers are used in CES packets (Yes or No)
Differential	Indicates whether DCR is active
Timestamp Freq	Specifies the optional timestamp frequency
CAS Framing	Specifies the type of CAS framing
Effective PVDT	Displays the peak-to-peak packet delay variation (PDV) used by the circuit emulation service. Since the operating system may adjust the jitter buffer setting in order to ensure no packet loss, the configured jitter buffer value may not be the value used by the system. The effective PVDT provides an indication that the PVD has been adjusted by the operating system (see Jitter Buffer)
AsymDelayControl	Specifies whether asymmetric delay control is enabled
RepeatPeriod	Specifies the number of times that the asymmetric delay control analysis is repeated

Table 50: Service-ID SAP Output Fields (Continued)

Label	Description
Samples	Specifies the number of packets that are analyzed (k = 1024)
Cfg Alarm	Specifies the alarms that have alarm reporting enabled
Alarm Status	Indicates the current alarm state (for example, stray, malformed, packet loss, overrun, underrun, remote packet loss, remote fault, or remote RDI)
CEM SAP Statistics	
Packets	(Column heading) Displays the number of packets counted for the statistic since the last counter reset
Seconds	(Column heading) Displays the number of seconds elapsed for the statistic since the last counter reset
Events	(Column heading) Displays the number of events counted for the statistic since the last counter reset
Egress Stats	Indicates that the following statistics are egress statistics
Forwarded	Displays the number of forwarded packets
Dropped	Displays the number of dropped packets
Missing	Displays the number of missing packets
Reordered Forwarded	Displays the number of packets that have been reordered and forwarded
Underrun	Displays the accumulated number of underrun packets for the number of underrun events
Overrun	Displays the accumulated number of overrun packets for the number of overrun events
Misordered Dropped	Displays the number of misordered packets that have been dropped
Malformed Dropped	Displays the number of malformed packets that have been dropped
LBit Dropped	Displays the number of L bit marked packets that have been dropped
Error	Displays the accumulated number of seconds that have passed while any error has occurred
Severely Error	Displays the accumulated number of seconds that have passed while severe errors have occurred
Unavailable	Displays the accumulated number of seconds that have passed while the Cpipe or MEF 8 Epipe is unavailable
Failure Count	Displays the accumulated number of failed events

Table 50: Service-ID SAP Output Fields (Continued)

Label	Description
Jitter Buffer Depth	Displays the number of packets sitting in the jitter buffer at that instant for a given Cpipe or MEF 8 Epipe
Jitter Buffer Sampling Complete	Displays the number of asymmetric delay control analysis periods completed
Jitter Buffer Adjust	Displays the number of jitter buffer adjustments made for asymmetric delay control
Ingress Stats	Indicates that the following statistics are ingress statistics
Forwarded	Displays the number of forwarded packets
Dropped	Displays the number of dropped packets

Sample Output (Epipe)

```
*A:csasim2>show>service>id# sap 1/3/1 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 3
SAP                : 1/3/1                    Encap             : q-tag
Admin State        : Up                       Oper State        : Down
LLF Admin State    : Up                       LLF Oper State    : Clear
Flags              : ServiceAdminDown
Multi Svc Site     : None
Last Status Change : 04/30/2008 13:55:04
Last Mgmt Change   : 05/07/2008 16:54:57
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Admin MTU          : 1518                     Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                     qinq-pbit-marking : both
Ing Scheduler Mode : 16-priority              Egr Scheduler Mode: 4-priority
Ing Agg Rate Limit : 50000
Ing Agg cir        : 5000
Egr Agg Rate Limit : max
Endpoint           : N/A
Q Frame-Based Acct : Disabled
Vlan-translation   : None

Acct. Pol          : None                     Collect Stats     : Disabled

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

```
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : 05/07/2008 21:32:32

                          Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                         0
Off. HiPrio           : 2655264                   2655264
Off. LowPrio          : 2655264                   2655264

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                         0
Dro. LowPrio          : 0                         0
For. InProf           : 3982896                   3982896
For. OutProf          : 1327632                   1327632

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                         0
Dro. OutProf          : 0                         0
For. InProf           : 2655264                   2655264
For. OutProf          : 2655264                   2655264
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------

                          Packets                 Octets

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

Egress Queue 1
For. InProf           : 0                         0
For. OutProf          : 0                         0
Dro. InProf           : 0                         0
Dro. OutProf          : 0                         0
-------------------------------------------------------------------------------
Eth-Cfm Configuration Information
-------------------------------------------------------------------------------
Md-index           : 2                        Direction         : Down
Ma-index           : 2                        Admin             : Disabled
MepId              : 2                        CCM-Enable        : Enabled
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : None
Mac Address        : 00:00:00:00:00:00
CcmTx              : 0                        CcmSequenceErr    : 0
DmrRepliesTx       : 0
LmrRepliesTx       : 0
Dual-Loss Test     : Enabled                  Dual-Loss Thresh  : 0.25%
Eth-Ais:           : Enabled                  Eth-Ais Rx Ais:   : No
```

```
Eth-Ais Tx Priorit*: 7                        Eth-Ais Rx Interv*: 1
Eth-Ais Tx Interva*: 1                        Eth-Ais Tx Counte*: 0
Eth-Ais Tx Levels  :
Eth-Tst:           : Disabled
LbRxReply          : 0                        LbRxBadOrder      : 0
LbRxBadMsdu        : 0                        LbTxReply         : 0
LbNextSequence     : 1                        LtNextSequence    : 1
LtRxUnexplained    : 0

Md-index           : 1                        Direction         : Down
Ma-index           : 1                        Admin             : Disabled
MepId              : 1                        CCM-Enable        : Enabled
LowestDefectPri    : macRemErrXcon            HighestDefect     : none
Defect Flags       : None
Mac Address        : 00:00:00:00:00:00
CcmTx              : 0                        CcmSequenceErr    : 0
DmrRepliesTx       : 0
LmrRepliesTx       : 0
Dual-Loss Test     : Disabled
Eth-Ais:           : Disabled
Eth-Tst:           : Disabled
LbRxReply          : 0                        LbRxBadOrder      : 0
LbRxBadMsdu        : 0                        LbTxReply         : 0
LbNextSequence     : 1                        LtNextSequence    : 1
LtRxUnexplained    : 0
===============================================================================
```

See Table 50 for Epipe field descriptions.

Sample Output (Fpipe)

```
*A:LL101210487>show>service# id 10021 sap 1/4/2:21 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1021
SAP                : 1/4/2:21                 Encap             : frRel
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 07/06/2011 18:06:20
Last Mgmt Change   : 07/06/2011 18:00:35
Sub Type           : regular
Split Horiz Group  : (Not Specified)

Admin MTU          : 1600                     Oper MTU          : 1600
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                     qinq-pbit-marking : both
Ing Agg Rate Limit : max                      Egr Agg Rate Limit: max
Endpoint           : N/A
FRF-12             : Disabled

Acct. Pol          : None                     Collect Stats     : Disabled
```

```
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 619549                  308535402

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 619549                  308535402

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 1803                    897894
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------

                        Packets                 Octets

Ingress Queue 1 (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 619549                  308535402
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 619549                  308535402

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 1803                    897894
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
===============================================================================

*A:LL101210487>show>service#
```

See Table 50 for Fpipe field descriptions.

Sample Output (Hpipe)

```
A:7705:Dut-C>config>service# show service id 501 sap 1/10/9.1 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 501
SAP                : 1/10/9.1                 Encap             : hdlc
Description        : Default sap description for service id 501
Admin State        : Up                       Oper State        : Down
Flags              : None
Multi Svc Site     : None
Last Status Change : 07/11/2011 14:16:53
Last Mgmt Change : 07/11/2011 14:16:14
Sub Type           : regular
Split Horizon Group: (Not Specified)
Admin MTU          : 1514                     Oper MTU          : 1514
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                     qinq-pbit-marking : n/a
Ing Agg Rate Limit : max                      Egr Agg Rate Limit: max
Endpoint           : N/A
Acct. Pol          : None                     Collect Stats     : Disabled
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A
                        Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 0                        0
Off. LowPrio          : 167                      15030
Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        0
Dro. LowPrio          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 167                      15030
Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 165                      14850
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                  Octets
Ingress Queue 1 (Priority)
Off. HiPrio           : 0                        0
Off. LoPrio           : 167                      15030
Dro. HiPrio           : 0                        0
Dro. LoPrio           : 0                        0
```

```
For. InProf           : 0                      0
For. OutProf          : 167                    15030
Egress Queue 1
For. InProf           : 0                      0
For. OutProf          : 165                    14850
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0
===============================================================================
*A:7705:Dut-C>config>service# exit all
```

See Table 50 for Hpipe field descriptions.

Sample Output (Ipipe)

```
*A:ALU-12# show service id 1301 sap 1/2/8:11 detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1301
SAP                : 1/2/8:11                Encap             : q-tag
Description        : Default sap description for service id 1301
Admin State        : Up                      Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 01/20/2009 16:44:14
Last Mgmt Change   : 01/21/2009 16:31:04
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100

Admin MTU          : 1572                    Oper MTU          : 1572
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
tod-suite          : None                    qinq-pbit-marking : both
Ing Scheduler Mode : 16-priority             Egr Scheduler Mode: 4-priority
Ing Agg Rate Limit : 50000
Ing Agg cir        : 5000
Egr Agg Rate Limit : max
Endpoint           : N/A
Q Frame-Based Acct : Disabled

Acct. Pol          : Default                 Collect Stats     : Enabled
Ce IP Address      : 88.1.10.3
SAP MAC Address    : 00:1a:f0:bd:ab:b0       Mac Refresh Inter*: 14400

-------------------------------------------------------------------------------
Ipipe SAP ARP Entry Info
-------------------------------------------------------------------------------
88.1.10.3       00:00:15:b9:6b:73 dynamic 03h50m24s

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 13                      Egress qos-policy : 13
Shared Q plcy      : n/a                     Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
```

```
-------------------------------------------------------------------------------
Last Cleared Time     : 01/21/2009 14:19:23
                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 20683584                1613319552
Off. LowPrio          : 1840167                 143533026


Queueing Stats(Ingress QoS Policy 13)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 11271525                879178950
For. OutProf          : 11252226                877673628
Queueing Stats(Egress QoS Policy 13)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 600                     46800
-------------------------------------------------------------------------------
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 2 (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 3 (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 4 (Unicast) (Priority)
Off. HiPrio           : 7304519                 569752482
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 5473927                 426966306
For. OutProf          : 1830592                 142786176

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

```
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0

Egress Queue 2
For. InProf           : 0                      0
For. OutProf          : 200                    15600
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0

Egress Queue 3
For. InProf           : 0                      0
For. OutProf          : 200                    15600
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0

Egress Queue 4
For. InProf           : 0                      0
For. OutProf          : 200                    15600
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:ALU-12#
```

See Table 50 for Ipipe field descriptions.

Sample Output (Ipipe with Frame Relay SAP)

```
*A:7705:Dut-C# show service id 2 sap 1/10/9.1:101 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 2
SAP                : 1/10/9.1:101             Encap             : frRel
Description        : Default sap description for service id 2
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 07/08/2011 07:29:50
Last Mgmt Change   : 07/08/2011 07:29:37
Sub Type           : regular
Split Horizon Group: (Not Specified)

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

FRF-12             : Disabled


Acct. Pol          : None                     Collect Stats     : Disabled

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

```
Ipipe SAP Configuration Information
-------------------------------------------------------------------------------
Ce IP Address      : 12.12.10.1
Configured CE IP   : 12.12.10.1              Discovered CE IP  : n/a

-------------------------------------------------------------------------------
Ipipe SAP ARP Entry Info
-------------------------------------------------------------------------------
No Ipipe SAP ARP entries

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 18648                   783216

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 18648                   783216

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 1020                    1521840
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 18648                   783216
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 18648                   783216

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 1020                    1521840
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
===============================================================================
*A:7705:Dut-C#
```

See Table 50 for Ipipe with Frame Relay field descriptions.

Sample Output (Ipipe with cHDLC SAP)

```
*A:7705:Dut-C# show service id 2 sap 1/10/9.1
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 2
SAP                : 1/10/9.1                 Encap             : cisco-hdlc
Description        : Default sap description for service id 2
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 07/08/2011 07:50:55
Last Mgmt Change   : 07/08/2011 07:50:44
-------------------------------------------------------------------------------
Ipipe SAP Configuration Information
-------------------------------------------------------------------------------
Ce IP Address      : 12.12.10.1
Configured CE IP   : 12.12.10.1               Discovered CE IP  : n/a

-------------------------------------------------------------------------------
Ipipe SAP ARP Entry Info
-------------------------------------------------------------------------------
No Ipipe SAP ARP entries

===============================================================================
*A:7705:Dut-C# show service id 2 sap 1/10/9.1 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 2
SAP                : 1/10/9.1                 Encap             : cisco-hdlc
Description        : Default sap description for service id 2
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 07/08/2011 07:50:55
Last Mgmt Change   : 07/08/2011 07:50:44
Sub Type           : regular
Split Horizon Group: (Not Specified)

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

Acct. Pol          : None                     Collect Stats     : Disabled

-------------------------------------------------------------------------------
Ipipe SAP Configuration Information
-------------------------------------------------------------------------------
Ce IP Address      : 12.12.10.1
Configured CE IP   : 12.12.10.1               Discovered CE IP  : n/a
```

```
-------------------------------------------------------------------------------
Ipipe SAP ARP Entry Info
-------------------------------------------------------------------------------
No Ipipe SAP ARP entries

-------------------------------------------------------------------------------

QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
Shared Q plcy      : n/a                     Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 0                        0
Off. LowPrio          : 489                      20538

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        0
Dro. LowPrio          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 489                      20538

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 123                      28782
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                  Octets

Ingress Queue 1 (Priority)
Off. HiPrio           : 0                        0
Off. LoPrio           : 489                      20538
Dro. HiPrio           : 0                        0
Dro. LoPrio           : 0                        0
For. InProf           : 0                        0
For. OutProf          : 489                      20538

Egress Queue 1
For. InProf           : 0                        0
For. OutProf          : 123                      28782
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
===============================================================================
```

See Table 50 for Ipipe with cHDLC field descriptions.

Sample Output (Apipe in a SAP aggregation group)

```
*A:7705custDoc:Sar18>show>service>id# sap 1/2/2.1:10/10
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 3
SAP                : 1/2/2.1:10/10            Encap            : atm
Description        : (Not Specified)
Admin State        : Up                       Oper State       : Down
Flags              : ServiceAdminDown
                     PortOperDown L2OperDown
Multi Svc Site     : None
Last Status Change : 08/24/2011 19:47:16
Last Mgmt Change   : 08/25/2011 15:07:47
Sap Aggregation Group: sap-agg-grp1
===============================================================================
*A:7705custDoc:Sar18>show>service>id#
```

See Table 50 for Apipe in a SAP aggregation group field descriptions.

sap-aggregation-group

Syntax **sap-aggregation-group** *group-id* [**base** | **detail** | **group-stats** | **stats**]

Context show>service>id

Description This command displays SAP aggregation group statistics.

Parameters *group-id* — identifies the aggregation group for the SAP

Values 1 to 32 characters

base — displays only the basic information about the SAP aggregation group

detail — displays the output for the SAP aggregation group in detailed format

group-stats — displays only the statistics for the SAP aggregation group

stats — displays only the statistics for the SAP aggregation group on a per-queue basis

Output The following is an example of the output for the **show sap-aggregation-group detail** command for an aggregation group named “GroupName”. Table 51 describes the fields.

Sample Output

```
*A:179_123>show>service>id# sap-aggregation-group GroupName detail
-------------------------------------------------------------------------------
SAP Aggregation Groups
-------------------------------------------------------------------------------
Group: GroupName
-------------------------------------------------------------------------------
Service Id         : 300
Acct. Pol          : None                     Collect Stats    : Disabled
-------------------------------------------------------------------------------
```

```
QOS
------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1

SAP Aggregation Group Statistics
------------------------------------------------------------------------------
Last Cleared Time     : 09/07/2009 10:03:57

Dropped Egress Cells (unconfigured vpi/vci): 0

                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       n/a
Off. HiPrio           : 77780484                n/a
Off. LowPrio          : n/a                     n/a
Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       n/a
Dro. LowPrio          : n/a                     n/a
For. InProf           : 0                       0
For. OutProf          : 77780484                5755755816
Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                       n/a
Dro. OutProf          : n/a                     n/a
For. InProf           : 77780417                4044581684
For. OutProf          : n/a                     n/a
------------------------------------------------------------------------------
SAP Aggregation Group per Queue stats
------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 77780484                n/a
Off. LoPrio           : n/a                     n/a
Dro. HiPrio           : 0                       n/a
Dro. LoPrio           : n/a                     n/a
For. InProf           : 0                       0
For. OutProf          : 77780484                5755755816

Egress Queue 1
For. InProf           : 77780417                4044581684
For. OutProf          : n/a                     n/a
Dro. InProf           : 0                       n/a
Dro. OutProf          : n/a                     n/a
==============================================================================
*A:179_123>show>service>id#
```

Table 51: Service-ID SAP Aggregation Group Output Fields

Label	Description
SAP Aggregation Groups	
Group *group-id*	
Service Id	Identifies the service
Acct. Pol	Specifies the accounting policy applied to the SAP aggregation group

Table 51: Service-ID SAP Aggregation Group Output Fields (Continued)

Label	Description
Collect Stats	Specifies whether accounting statistics are collected on the SAP aggregation group
QOS	
Ingress qos-policy	Displays the SAP ingress QoS policy ID
Egress qos-policy	Displays the SAP egress QoS policy ID
SAP Aggregation Group Statistics	
Last Cleared Time	Displays the date and time that a clear command was issued on the aggregation group statistics
Dropped Egress Cells (unconfigured vpi/vci)	Displays the number of unconfigured or unknown VPI/VCI cells that are received in an ATM PW payload from the network
Forwarding Engine Stats	
Dropped	Indicates the number of packets or octets dropped by the forwarding engine
Off. HiPrio	Indicates the number of high-priority packets or octets offered to the forwarding engine
Off. LowPrio	Indicates the number of low-priority packets offered to the forwarding engine
Queueing Stats (Ingress QoS Policy *n*)	
Dro. HiPrio	Indicates the number of high-priority packets or octets discarded, as determined by the SAP ingress QoS policy
Dro. LowPrio	Indicates the number of low-priority packets discarded, as determined by the SAP ingress QoS policy
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP ingress QoS policy
For. OutProf	Indicates the number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP ingress QoS policy
Queueing Stats (Egress QoS Policy *n*)	
Dro. InProf	Indicates the number of in-profile packets or octets discarded, as determined by the SAP egress QoS policy
Dro. OutProf	Indicates the number of out-of-profile packets or octets discarded, as determined by the SAP egress QoS policy
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP egress QoS policy

Table 51: Service-ID SAP Aggregation Group Output Fields (Continued)

Label	Description
For. OutProf	Indicates the number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP egress QoS policy
SAP Aggregation Group per Queue stats	
Ingress Queue *n* (Unicast) (Priority)	Specifies the index of the ingress QoS queue of the SAP, where *n* is the index number
Off. HiPrio	Indicates the number of packets or octets of high-priority traffic for the SAP (offered)
Off. LoPrio	Indicates the number of packets or octets of low-priority traffic for the SAP (offered)
Dro. HiPrio	Indicates the number of high-priority traffic packets or octets discarded
Dro. LoPrio	Indicates the number of low-priority traffic packets or octets discarded
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	Indicates the number of out-of-profile packets or octets (rate above CIR) forwarded
Egress Queue *n*	Specifies the index of the egress QoS queue of the SAP, where *n* is the index number
For. InProf	Indicates the number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	Indicates the number of out-of-profile packets or octets (rate above CIR) forwarded
Dro. InProf	Indicates the number of in-profile packets or octets discarded
Dro. OutProf	Indicates the number of out-of-profile packets or octets discarded

sap-using

Syntax **sap-using** [**sap** *sap-id*]
sap-using [**sap-aggregation-group** *group-id*]
sap-using description
sap-using [**ingress** | **egress**] **atm-td-profile** *td-profile-id*
sap-using [**ingress** | **egress**] **qos-policy** [*qos-policy-id* | *qos-policy-name*]
sap-using [**ingress** | **egress**] **scheduler-mode** {**4-priority** | **16-priority**}

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 *group-id* — identifies the SAP aggregation group for the SAP. The value applies to the whole node. For example, if the same identifier is used for SAP aggregation groups in multiple services, the **show** command, when issued with the identifier, displays all the associated services.

Values 1 to 32 characters

description — displays a SAP summary table with description information

ingress — specifies matching an ingress policy

egress — specifies matching an egress policy

scheduler-mode — specifies the scheduler mode for which to display matching SAPs

Values 4-priority or 16-priority

qos-policy-id — identifies the ingress or egress QoS policy ID for which to display matching SAPs

Values 1 to 65535

qos-policy-name — identifies the ingress or egress QoS policy name for which to display matching SAPs

Values up to 64 characters

td-profile-id — displays SAPs using this traffic description

sap-id — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

Output The following outputs are examples of service SAP-using information, and Table 52 describes the fields.

Sample Output

```
*A:ALU-48# show service sap-using
===============================================================================
Service Access Points
===============================================================================
PortId            SvcId     Ing.  Ing.    Egr.  Egr.   Adm  Opr
                            QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/2/7:1           103       1     none    1     none   Up   Up
1/2/7:2           104       1     none    1     none   Up   Up
1/2/7:3           105       1     none    1     none   Up   Up
1/1/1.1           303       1     none    1     none   Up   Up
1/1/1.2           304       1     none    1     none   Up   Up
1/1/1.3           305       1     none    1     none   Up   Up
1/1/9.1:10/50     701       1     none    1     none   Up   Down
1/1/9.1:20        702       1     none    1     none   Up   Down
1/1/9.1:10/51     703       1     none    1     none   Up   Down
1/1/9.1:30        704       1     none    1     none   Up   Down
1/1/9.1:10/52     705       1     none    1     none   Up   Down
1/1/9.1:40        706       1     none    1     none   Up   Down
1/1/9.1:11/50     805       1     none    1     none   Up   Down
1/1/9.1:21        806       1     none    1     none   Up   Down
1/1/9.1:12/52     807       1     none    1     none   Up   Down
1/1/9.1:41        808       1     none    1     none   Up   Down
1/1/1.9           903       1     none    1     none   Up   Up
1/1/1.10          904       1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 18
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48#

*A:ALU-48# show service sap-using sap 1/1/21:0
===============================================================================
Service Access Points Using Port 1/1/21:0
===============================================================================
PortId                SvcId      Ing.  Ing.    Egr.  Egr.    Adm  Opr
                                 QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/1/21:0              1          1     none    1     none    Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48#

*A:ALU-48# show service sap-using description
===============================================================================
Service Access Points
===============================================================================
PortId                            SvcId        Adm  Opr  Description
-------------------------------------------------------------------------------
1/1/2                             1            Down Down (Not Specified)
1/12/3:100.100                    1            Up   Down (Not Specified)
1/2/1.1                           4            Up   Down (Not Specified)
1/10/4                            5            Up   Down (Not Specified)
1/12/2:100                        5006         Up   Down (Not Specified)
```

```
-------------------------------------------------------------------------------
Number of SAPs : 5
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48#

*A:ALU-48# show service sap-using egress atm-td-profile 1
===============================================================================
Service Access Point Using ATM Traffic Profile 1
===============================================================================
PortId            SvcId     Ing.  Ing.    Egr.  Egr.   Adm  Opr
                            QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/1/9.1:10/50     701       1     none    1     none   Up   Down
1/1/9.1:20        702       1     none    1     none   Up   Down
1/1/9.1:10/51     703       1     none    1     none   Up   Down
1/1/9.1:30        704       1     none    1     none   Up   Down
1/1/9.1:10/52     705       1     none    1     none   Up   Down
1/1/9.1:40        706       1     none    1     none   Up   Down
1/1/9.1:11/50     805       1     none    1     none   Up   Down
1/1/9.1:21        806       1     none    1     none   Up   Down
1/1/9.1:12/52     807       1     none    1     none   Up   Down
1/1/9.1:41        808       1     none    1     none   Up   Down
-------------------------------------------------------------------------------
Saps : 10
===============================================================================
*A:ALU-12#

*A:7705custDoc:Sar18>show>service# sap-using ingress scheduler-mode 4-priority
===============================================================================
Service Access Points Using Ingress 4-priority Scheduler Mode
===============================================================================
PortId                          SvcId      Scheduler Mode    Adm  Opr
-------------------------------------------------------------------------------
1/12/5                          7000       4-priority        Down Down
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================

*A:7705custDoc:Sar18>show>service# sap-using ingress scheduler-mode 16-priority
===============================================================================
Service Access Points Using Ingress 16-priority Scheduler Mode
===============================================================================
PortId                          SvcId      Scheduler Mode    Adm  Opr
-------------------------------------------------------------------------------
1/12/7                          4005       16-priority       Down Down
1/12/8                          5005       16-priority       Up   Down
1/12/6                          6000       16-priority       Down Down
-------------------------------------------------------------------------------
Number of SAPs : 3
-------------------------------------------------------------------------------
===============================================================================
*A:7705custDoc:Sar18>show>service#
```

Sample Output (sap-using sap-aggregation-group)

```
A:SYS28>config>service>apipe# show service sap-using sap-aggregation-group SAG
===============================================================================
Sap Aggregation Groups
===============================================================================
GroupName                       SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
SAG                             1573       1     n/a     1     n/a    n/a  n/a
  bundle-ima-1/2.2:2/201                                                Up   Up
  1/2/9.1:2/201                                                         Up   Up
-------------------------------------------------------------------------------
SAG                             1574       1     n/a     1     n/a    n/a  n/a
  bundle-ima-1/2.2:2/301                                                Up   Up
  1/2/9.1:2/301                                                         Up   Up
-------------------------------------------------------------------------------
Number of Sap Aggregation Groups : 2
-------------------------------------------------------------------------------
===============================================================================
```

Table 52: Show Service SAP-Using Output Fields

Label	Description
GroupName	Displays the SAP aggregation group identifier
SvcId	Identifies the service
Scheduler Mode	Identifies the scheduler mode: 4-priority or 16-priority
Ing.QoS	Displays the SAP ingress QoS policy number specified on the ingress SAP
Ing.Fltr	Displays the SAP ingress filter data
Egr.QoS	Displays the SAP egress QoS policy number specified on the egress SAP
Egr.Fltr	Displays the SAP egress filter data
PortId	Identifies the access port where the SAP is defined
Adm	Specifies the desired state of the SAP
Opr	Indicates the actual state of the SAP
Description	Provides a description of the SAP
Number of SAPs/Saps	Number of SAPs using this service

scada

Syntax **scada** *bridge-id*

Context show>service

Description This command displays SAP information for a specified SCADA bridge

Parameters *bridge-id* — specifies the SCADA bridge ID, in the format *slot/mda/bridge-id*, where *bridge-id* is 1 to 16

Output The following output is an example of SCADA bridge SAP information, and Table 53 describes the fields.

Sample Output

```
*A:ALU-48# show service scada 1/8/1
===============================================================================
Service Access Points
===============================================================================
SapId           SvcId      Ing.  Ing.    Egr.  Egr.   Sap  Sap  Svc
                           QoS   Fltr    QoS   Fltr   Adm  Opr  Opr
-------------------------------------------------------------------------------
1/8/1.1         100        1     none    1     none   Up   Down Down
-------------------------------------------------------------------------------
Number of SAPs : 1
===============================================================================
```

Table 53: Show Service SCADA Bridge SAP Fields

Label	Description
Service Access Points	
SapId	Identifies the SAP
SvcId	Identifies the service
Ing.QoS	Displays the SAP ingress QoS policy number specified on the ingress SAP
Ing.Fltr	Displays the SAP ingress filter data
Egr.QoS	Displays the SAP egress QoS policy number specified on the egress SAP
Egr.Fltr	Displays the SAP egress filter data
SapAdm	Indicates the administrative state of the SAP
SapOpr	Indicates the operational state of the SAP
SvcOpr	Indicates the operational state of the service

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]

Context show>service>id

Description This command displays information for the SDPs associated with the service.

If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — displays only information for the specified SDP ID

Values 1 to 17407

ip-address — displays only SDPs matching the specified far-end IP address

Default SDPs with any far-end IP address

detail — displays detailed SDP information

keep-alive-history — displays the last 50 keepalive events for the SDP

Output The following output is an example of service-id SDP information, and Table 54 describes the fields.

Sample Output (Cpipe)

```
*A:csasim2>show>service>id# sdp 1 detail

===============================================================================
Service Destination Point (Sdp Id : 1) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 1:1  -(10.10.10.100)
-------------------------------------------------------------------------------
SDP Id             : 1:1                       Type             : Spoke
VC Type            : CESoPSN                   VC Tag           : 0
Admin Path MTU     : 0                         Oper Path MTU    : 0
Far End            : 10.10.10.100              Delivery         : LDP

Admin State        : Up                        Oper State       : Down
Acct. Pol          : None                      Collect Stats    : Disabled
Ingress Label      : 0                         Egress Label     : 0
Ing mac Fltr       : n/a                       Egr mac Fltr     : n/a
Ing ip Fltr        : n/a                       Egr ip Fltr      : n/a
Admin ControlWord  : Preferred                 Oper ControlWord : True
Admin BW(Kbps)     : 0                         Oper BW(Kbps)    : 0
Last Status Change : 04/30/2008 13:55:10       Signaling        : TLDP
Last Mgmt Change   : 05/02/2008 21:37:14
Endpoint           : N/A                       Precedence       : 4
Class Fwding State : Down
Flags              : SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
Mac Move           : Ukwn                      Blockable Level  : Unknown
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
```

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

Statistics           :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0
-------------------------------------------------------------------------------
CPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Local Bit-rate     : 1                        Peer Bit-rate     : n/a
Local Payload Size : 64                       Peer Payload Size : n/a
Local Jitter Buffer: 32                       Peer Jitter Buffer: n/a
Local Asym Delay   : enabled                  Peer Asym Delay   : n/a
Local Sig Pkts     : No Sig.                  Peer Sig Pkts     : No Sig.
Local CAS Framing  : No CAS                   Peer CAS Framing  : No CAS
Local RTP Header   : No                       Peer RTP Header   : No
Local Differential : No                       Peer Differential : No
Local Timestamp    : 0                        Peer Timestamp    : 0
===============================================================================

*A:csasim2>show>service>id#
```

Table 54: SDP Output Fields

Label	Description
Service Destination Points (SDPs)	
Description	Displays generic information about the SDP
SDP Id	Identifies the SDP
Type	Identifies the service SDP binding type (for example, spoke)
VC Type	Displays the VC type for the SDP (for example, CESoPSN)
VC Tag	The explicit dot1q value used when encapsulating to the SDP far end
Admin Path MTU	Specifies the desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented
Oper Path 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
Far End	Displays the IP address of the far end of the MPLS or GRE tunnel defined by this SDP
Delivery	Specifies the type of delivery used by the SDP (MPLS or GRE)
Admin State	Specifies the administrative state of this SDP

Table 54: SDP Output Fields (Continued)

Label	Description
Oper State	Specifies the operational state of this SDP
Acct. Pol	The accounting policy ID assigned to the SAP
Collect Stats	Specifies whether collect stats is enabled
Ingress Label	Displays the label used by the far-end device to send packets to this device in this service by this SDP
Egress Label	Displays the label used by this device to send packets to the far-end device in this service by this SDP
Admin ControlWord	Specifies the administrative state of the control word: Preferred (control word enabled) or Not Preferred (control word disabled)
Oper ControlWord	Specifies the operational state of the control word: True (control word enabled) or False (control word disabled)
Last Status Change	Specifies the time of the most recent operating status change to this spoke SDP
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this spoke SDP
Flags	Displays the conditions that affect the operating status of this spoke SDP. Display output includes PathMTUtooSmall, SdpOperDown, NoIngVCLabel, NoEgrVCLabel, and so on
Mac Move	Indicates the administrative state of the MAC movement feature associated with the service
Peer Pw Bits	Displays the setting of the pseudowire peer bits. Display output includes pwNotforwarding, psnIngressFault, psnEgressFault, IacIngressFault, lacEgressFault
Peer Fault Ip	N/A
Peer Vccv CV Bits	Displays the setting of the pseudowire peer VCCV control verification bits (lspPing)
Peer Vccv CC Bits	Displays the setting of the pseudowire peer VCCV control channel bits (pwe3ControlWord and/or mplsRouterAlertLabel)
Keepalive Information	
Admin State	Specifies the administrative state of the keepalive protocol
Oper State	Specifies the operational state of the keepalive protocol

Table 54: SDP Output Fields (Continued)

Label	Description
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
Statistics	
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets
I. Dro. Pkts.	Specifies the number of dropped ingress packets
I. Fwd. Octs.	Specifies the number of forwarded ingress octets
I. Dro. Octs.	Specifies the number of dropped ingress octets
E. Fwd. Pkts.	Specifies the number of forwarded egress packets
E. Fwd. Octets	Specifies the number of forwarded egress octets
Eth-Cfm Configuration Information	
Md-index	Displays the value of the MD index
Direction	Displays the direction of the MEP
Ma-index	Displays the value of the MA index
Admin	Displays the administrative state of the MEP (enabled or disabled)
MepId	Displays the MEP-ID
CCM-Enable	Displays the status of the Continuity Check Message (CCM)
LowestDefectPri	Displays a configured value that defects are evaluated against
HighestDefect	Displays the highest defect
Defect Flags	Indicates the defect flags
Mac Address	Displays the MAC address (the MAC address for a spoke SDP is the system MAC address; for a SAP, it is the port MAC address)
CcmLtmPriority	Displays the priority of the CCM Linktrace Message (LTM)
CcmTx	Displays the number of CCM transmissions

Table 54: SDP Output Fields (Continued)

Label	Description
CcmSequenceErr	Displays the number of CCM sequence errors
LbRxReply	Displays the number of received loopback (LB) replies
LbRxBadOrder	Displays the number of LB replies that have been received in the wrong order
LbRxBadMsdu	Displays the number of LB replies that have been received with the wrong destination MAC address (MSDU = MAC Service Data Unit)
LbTxReply	Displays the number of LBRs (loopback replies) transmitted out this MEP
LbNextSequence	Displays the sequence number of the next LB transmission
LtNextSequence	Displays the sequence number of the next Linktrace (LT) message transmitted
LtRxUnexplained	Displays the number of the unexplained Linktrace (LT) messages
Associated LSP LIST	
Lsp Name	Specifies the name of the static LSP
Admin State	Specifies the administrative state of the associated LSP
Oper State	Specifies the operational state of the associated LSP
Time Since Last Tr*	Specifies the time that the associated static LSP has been in service
APIPE Service Destination Point specifics	
Admin Concat Limit	Specifies the administrative (configured) value for the maximum number of cells for cell concatenation, as defined via the `max-cells` command
Oper Concat Limit	Specifies the operational value for the maximum number of cells for cell concatenation
Peer Concat Limit	Specifies the far-end value for the maximum number of cells for cell concatenation
Max Concat Delay	Specifies the amount of time to wait while cell concatenation is occurring, as defined via the `max-delay` command
CPIPE Service Destination Point specifics	
Local Bit-rate	Specifies the number of DS0s used by the local SDP
Peer Bit-rate	Specifies the number of DS0s used by the far-end SDP

Table 54: SDP Output Fields (Continued)

Label	Description
Local Payload Size	Specifies the local payload size, in bytes, used by the local SDP
Peer Payload Size	Specifies the peer payload size, in bytes, used by the far-end SDP
Local Jitter Buffer	Specifies the jitter buffer size, in milliseconds, used by the local SDP
Peer Jitter Buffer	Specifies the jitter buffer size, in milliseconds, used by the far-end SDP
Local Asym Delay	Specifies whether asymmetric delay control is enabled on the local SDP
Peer Asym Delay	Specifies whether asymmetric delay control is enabled on the far-end SDP
Local Sig Pkts	Specifies the type of signaling packets used by the local SDP
Peer Sig Pkts	Specifies the type of signaling packets used by the far-end SDP
Local CAS Framing	Specifies the type of CAS framing used by the local SDP
Peer CAS Framing	Specifies the type of CAS framing used by the far-end SDP
Local RTP Header	Specifies whether the local router inserts the RTP header
Peer RTP Header	Specifies whether the peer router inserts the RTP header
Number of SDPs	Specifies the number of SDPs bound to the service

sdp-using

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

Context show>service

Description Displays the services using SDP or far-end address options.

Parameters *sdp-id* — displays the services using the specified SDP

Values 1 to 17407

vc-id — displays the virtual circuit identifier

Values 1 to 4294967295

ip-address — displays the services that match the specified far-end IP address

Values 1 to 2147483647

Output The following output is an example of the show service sdp-using command, and Table 55 describes the fields.

Sample Output

```
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    131070
===============================================================================
Number of SDPs :4
```

Table 55: Sdp-using Output Fields

Label	Description
Sdp-using	
Svc Id	Identifies the service
Sdp Id	Identifies the SDP
Type	Identifies the type of SDP
Far End	Identifies the far-end IP address of SDP
Opr State	Identifies the operational state of this service
I.Label	Identifies the ingress label used by the far-end device to send packets to this device in this service by this SDP.
E.Label	Identifies the egress label used by this device to send packets to the far-end device in this service by this SDP.

service-using

Syntax **service-using** [**epipe**] [**ies**] [**vpls**] [**vprn**] [**apipe**] [**fpipe**] [**ipipe**] [**cpipe**] [**hpipe**] [**sdp-id** *sdp-id*] **customer-id** *customer-id*

Context show>service

Description Displays the services matching the specified 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 services

vpls — displays matching VPLS services

vprn — displays matching VPRN services

apipe — displays matching Apipe services

fpipe — displays matching Fpipe services

ipipe — displays matching Ipipe services

cpipe — displays matching Cpipe services

hpipe — displays matching Hpipe services

sdp-id — displays the services using the specified SDP

Values 1 to 17407

customer-id — displays the services using the specified customer ID

Values 1 to 2147483647

Output The following is an example of the output for the show service service-using command for an Fpipe. Table 56 describes the fields.

Sample Output (fpipe)

```
A:ALA-48# show service service-using fpipe

===============================================================================
Services [fpipe]
===============================================================================
ServiceID     Type      Adm     Opr       CustomerId        Service Name
-------------------------------------------------------------------------------
32            fpipe     Up      Down      9                 fpipe32_main
2011 1509:28
33            fpipe     Up      Down      9                 fpipe33_exec
2011 1509:28
102           fpipe     Up      Down      21                fpipe32_sales
2011 1509:28
364           fpipe     Up      Down      21
2011 1509:28
1005          fpipe     Down    Down      21
2011 1509:28
===============================================================================
Matching Services : 5
```

Table 56: Services-using Output Fields

Label	Description
ServiceID	Identifies the service
Type	Identifies the service type
Adm	Specifies the administrative state of this service

Table 56: Services-using Output Fields (Continued)

Label	Description
Opr	Specifies the operational state of this service
CustomerId	Specifies the ID of the customer
Service Name	The service name

Clear Commands

counters

Syntax **counters**

Context clear>service>statistics>id

Description This command clears all traffic queue counters associated with the service ID.

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears counters for a specific service.

Parameters *service-id* — uniquely identifies a service

Values 1 to 2147483690 or *svc-name*

sap

Syntax **sap** *sap-id* {**all** | **cem** | **counters**}

Context clear>service>statistics

Description This command clears SAP statistics for a SAP.

Parameters *sap-id* — specifies the physical port or bridge identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

all — clears all SAP queue statistics and STP statistics

cem — clears all queue statistics associated with a cem SAP

counters — clears all queue statistics associated with the SAP. Counters are not supported for members of SAP aggregation groups. QoS counters are not applicable to N > 1 SAPs.

sap-aggregation-group

Syntax **sap-aggregation-group** *group-id* {**all** | **counters**}

Context clear>service>statistics

Description This command clears all non-ATM layer SAP statistics for a SAP, that is, the QoS queue counters.

Parameters *group-id* — specifies the identifier for the SAP aggregation group

Values 1 to 32 characters

all — clears all QoS queue counters

counters — clears all QoS queue counters. The **counters** option performs the same function as the **all** option.

sdp

Syntax **sdp** *sdp-id* **keep-alive**

Context clear>service>statistics

Description This command clears keepalive statistics associated with the SDP ID.

Parameters *sdp-id* — identifies the SDP for which to clear keepalive statistics

Values 1 to 17407

arp

Syntax **arp**

Context clear>service>id

Description This command clears the ARP entries from an Ipipe service.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*:*vc-id* **ingress-vc-label**
spoke-sdp *sdp-id*:*vc-id* {**all** | **counters**}

Context clear>service>id
clear>service>statistics>id

Description This command clears and resets the spoke SDP bindings for the service.

Parameters *sdp-id* — the spoke SDP ID to be reset

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

all — clears all queue statistics and STP statistics associated with the SDP

counters — clears all queue statistics associated with the SDP

ingress-vc-label — clears the VC ingress value associated with the specified connection

VPLS

In This Chapter

This chapter provides an overview of Virtual Private LAN Service (VPLS) on the 7705 SAR. Topics in this chapter include:

- VPLS Overview
- VPLS Features
- Routed VPLS
- VPLS and Spanning Tree Protocol
- ATM PVC Access and Termination on a VPLS Service
- VPLS Service Considerations
- Configuration Notes
- Configuring a VPLS Service with CLI
- VPLS Command Reference

VPLS Overview

Topics in this section include:

- VPLS Packet Walkthrough
- Bridged Mobile Backhaul
- Multi-Tenant Unit (MTU) Termination

Virtual Private LAN Service (VPLS), as described in RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*, is a type of virtual private network service that allows the connection of multiple sites in a single bridged domain over a provider-managed IP/MPLS network or a Layer 2 Ethernet bridged network. The customer sites in a VPLS instance appear to be on the same LAN, regardless of their location. VPLS uses a native Ethernet SAP or a bridged PDU encapsulated SAP 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 pseudowire service (Epipe, Ipipe, etc.) and outsourced routed services (VPRN). Unlike VPRN service, 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 architecture 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 (local service) or more (distributed service) service routers. The connection appears to be a bridged domain to the customer sites so that protocols, including routing protocols, can traverse the VPLS service.

Other VPLS advantages include:

- VPLS is a transparent, protocol-independent bridged service
- no Layer 2 protocol conversion between LAN and WAN technologies
- no need to design, manage, configure, and maintain separate WAN access equipment, thereby eliminating the need to train personnel on WAN technologies such as ATM, IP over ATM, IP over PPP, and so on

VPLS is supported on the cards and platforms listed below. A VPLS SAP can reside on the following ports:

- any Ethernet port (null or tagged) in access mode
  - → on an 8-port Ethernet Adapter card with CLI identifier `a8-eth` or `a8-ethv2` installed on a 7705 SAR-8 or 7705 SAR-18
  - → on a 6-port Ethernet 10Gbps Adapter card with CLI identifier `a6-eth-10G` installed on a 7705 SAR-8 Shelf V2 with CSMv2 or a 7705 SAR-18
  - → on an 8-port Gigabit Ethernet Adapter card with CLI identifier `a8-1gb-sfp`, `a8-1gb-v2-sfp`, or `a8-1gb-v3-sfp` installed on a 7705 SAR-8 or 7705 SAR-18
  - → on a 10-port 1GigE/1-port 10GigE X-Adapter card in 10-port 1GigE mode with CLI identifier for mda-mode `x10-1gb-sfp` installed on a 7705 SAR-18
  - → on a 7705 SAR-F with CLI identifier `i8-eth`
  - → on a 7705 SAR-M (all variants) with CLI identifier `i7-1gb`
  - → on a GPON module and on DSL modules when the module is installed in the 7705 SAR-M (variants with module slots)
  - → on a 4-port SAR-H Fast Ethernet module when the module is installed in the 7705 SAR-H
  - → on a 6-port SAR-M Ethernet module when the module is installed in the 7705 SAR-M (variants with module slots)
  - → on a 7705 SAR-A (both variants)
  - → on a 7705 SAR-W
  - → on a 7705 SAR-Wx (all variants)
  - → on a 7705 SAR-H with CLI identifier `i8-1gb`
  - → on a 7705 SAR-Hc with CLI identifier `i6-1gb`
  - → on a 7705 SAR-X with CLI identifier `i7-mix-eth`
- any port using ATM encapsulation on a 4-port OC3/STM1 Clear Channel Adapter card installed on a 7705 SAR-8 or 7705 SAR-18

The transport of VPLS service is supported by LDP, GRE, and RSVP-TE tunnels, as well as static LSPs and dot1q-, qinq-, or null-encapsulated Ethernet SAPs at uplink.

VPLS Redundancy

Redundancy for a VPLS instance is provided using the endpoint concept to define primary and secondary spoke SDPs. This type of redundancy functions in a similar manner to PW redundancy. Refer to Pseudowire Redundancy for more information.

In addition, VPLS supports Spanning Tree Protocol (STP) on a per-VPLS instance basis, as well as management VPLS (mVPLS), where several VPLS instances are associated with a single STP instance running over redundant SAPs. The result of this STP is applied to the other VPLS services associated with the mVPLS instance. See VPLS and Spanning Tree Protocol for more information.

Access Control and Traffic Management

Access Control to and within VPLS is controlled via IP and MAC filter policies for ingress SAPs and SDPs (spoke and mesh), and IP filter policies for egress Ethernet SAPs. Traffic Management (TM) support at ingress and egress for unicast traffic is almost the same as TM support for an Ethernet PW SAP. The TM implementation is extended to support:

- at SAP ingress, queue selection for unicast and for broadcast, multicast, and unknown (BMU) traffic
- at network ingress, separate unicast and BMU queues
- at access ingress, ATM traffic (unicast and BMU) mapped to a single queue

Split Horizon

Within the context of VPLS services, a loop-free topology within a fully meshed VPLS core is achieved by applying a split-horizon forwarding concept whereby 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.

The 7705 SAR supports split horizon groups (SHGs) and residential SHGs, making it possible to control how traffic is propagated via configuring and applying forwarding directions to the received traffic. SHGs prevent multicast traffic from flowing within the same group, thereby preventing any loops.

In applications such as DSL aggregation, it is useful to extend the split-horizon concept to groups made up of SAPs and spoke SDPs. This extension is referred to as a split horizon group or residential bridging.

Traffic arriving on a SAP or a spoke SDP within a split horizon group is not copied to other SAPs and spoke SDPs in the same split horizon group; however, it is copied to SAPs and spoke SDPs in other split horizon groups, if these exist within the same VPLS.

Residential SHGs are only supported by ATM encapsulated SAPs. Residential (ATM) SAPs do not forward broadcast or unknown traffic; they only process known unicast traffic. Residential SAPs allow one queue per direction (ingress and egress) for all traffic types (unicast and BMU). In addition, OAM processing is allowed on residential (ATM) SAPs.

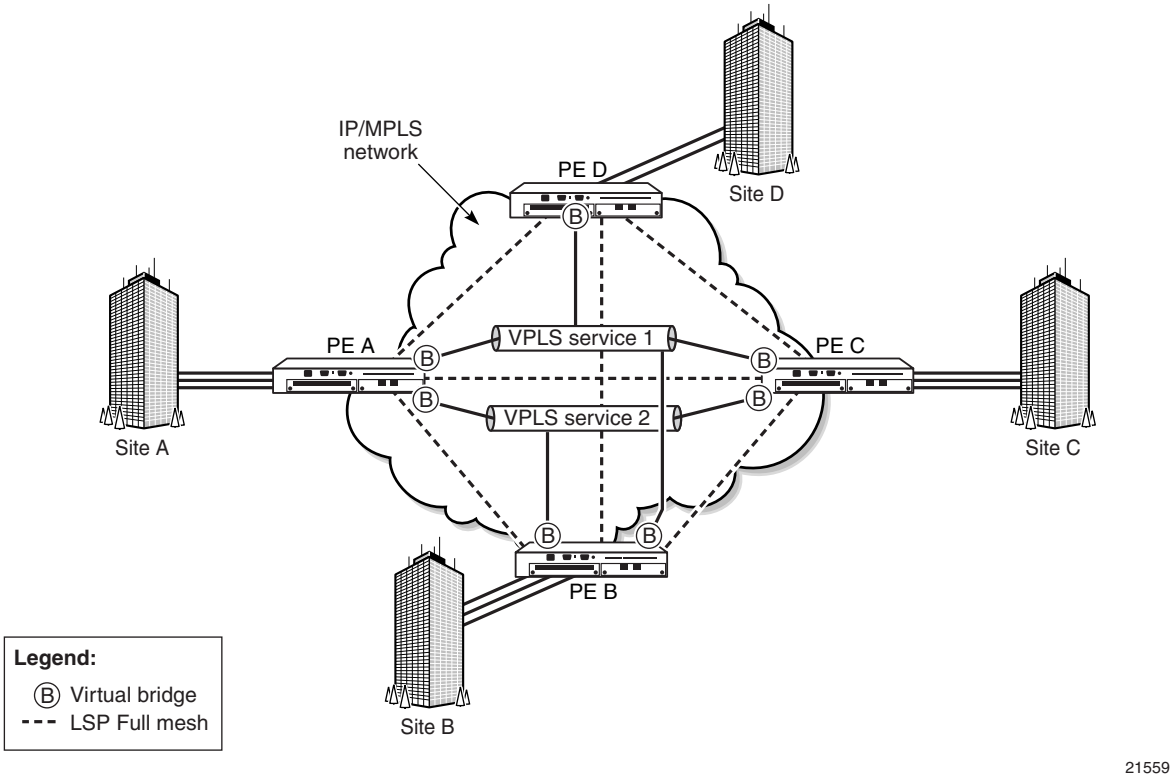
OAM support includes support for VPLS mac-ping, mac-trace, and cpe-ping.

Additional 7705 SAR support for VPLS service includes capabilities such as DHCP relay (on Ethernet SAPs). See VPLS Enhancements.

VPLS Packet Walkthrough

This section describes an example of VPLS processing based on the network shown in Figure 70.

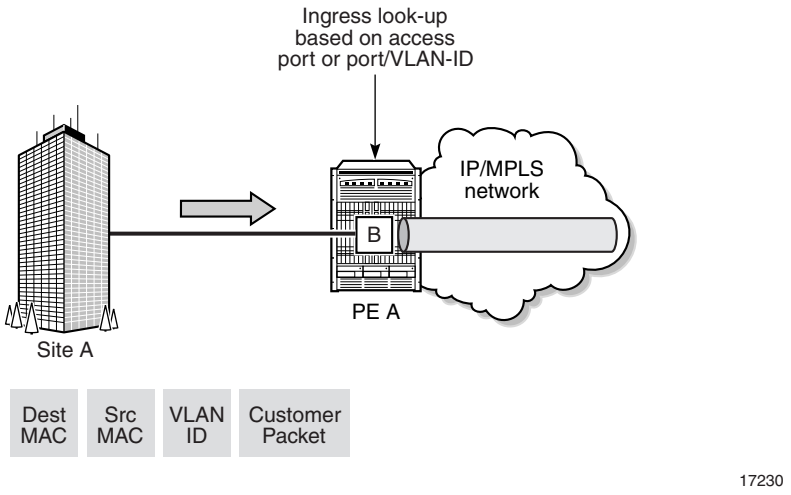
Figure 70: VPLS Service Architecture



1. **PE Router A** (Figure 71):
   a. Service packets arriving at PE A are associated with a VPLS service instance (VPLS service 2) based on the combination of the physical port and the IEEE 802.1Q tag (VLAN-ID) in the packet, if applicable.
   b. PE A learns the source MAC address in the packet and creates an entry in the Forwarding Database (FDB) table that associates the MAC address with the SAP on which it was received.

c. The destination MAC address in the packet is looked up in the FDB table for the VPLS instance. There are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address is not yet learned (unknown MAC address).

Figure 71: Access Port Ingress Packet Format and Lookup



For a Known MAC Address (Figure 72):

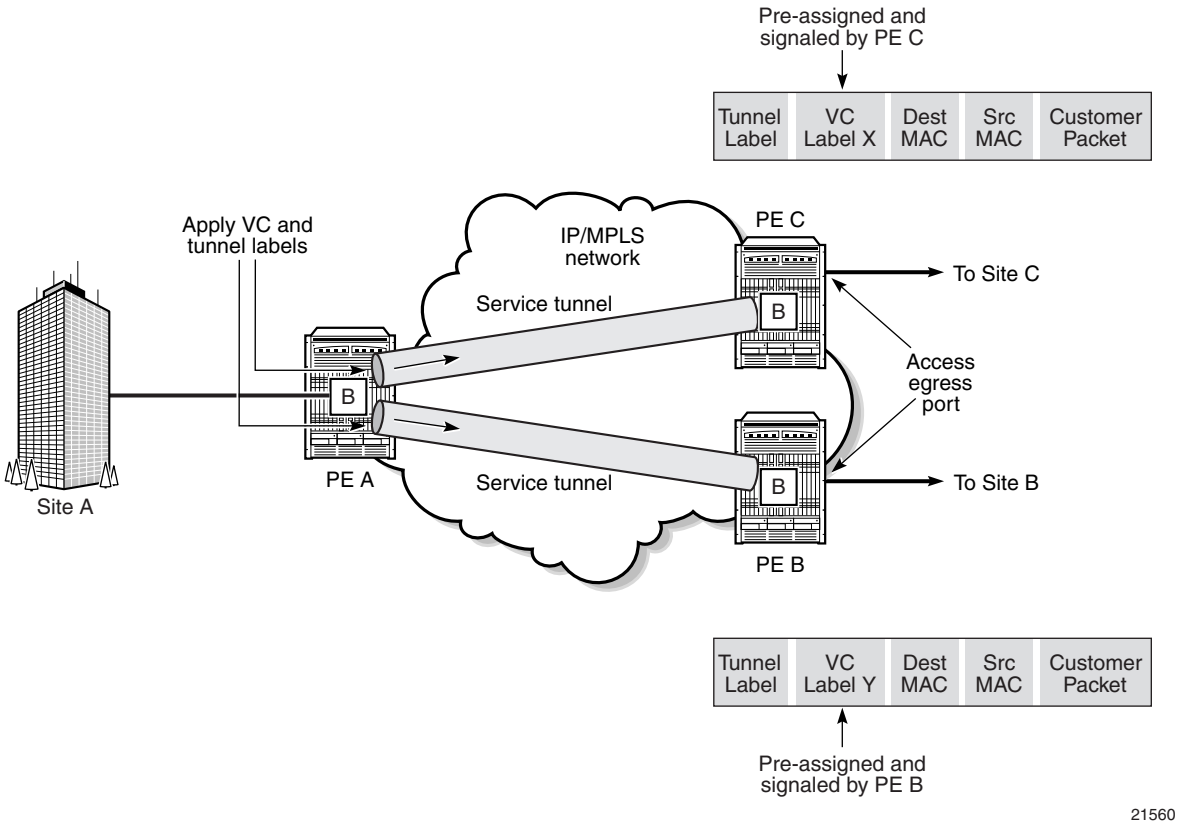
d. If the destination MAC address has already been learned by PE A, an existing entry in the FDB table identifies the far-end PE router and the service VC label (inner label) to be used before sending the packet to far-end PE C.

e. PE A chooses a transport LSP to send the customer packets to PE C. The customer packet is sent on this LSP once the IEEE 802.1Q tag is stripped and the service VC label (inner label) and the transport label (outer label) are added to the packet.

For an Unknown MAC Address (Figure 72):

f. If the destination MAC address has not been learned, PE A will flood the packet to both PE B and PE C that are participating in the service by using the VC labels that each PE router previously signaled for the VPLS instance. The packet is not sent to PE D since this VPLS service does not exist on that PE router.

Figure 72: Network Port Egress Packet Format and Flooding



2. **Core Router Switching**:

   a. All the core routers (P routers in IETF nomenclature) between PE A and routers PE B and PE C are Label Switch Routers (LSRs) that switch the packet based on the transport (outer) label of the packet until the packet arrives at the far-end PE router. All core routers are unaware of the content of the LSP payload (that is, the core routers do not know that this traffic is associated with a VPLS service).

3. **PE Router C**:
   a. PE C strips the transport label of the received packet to reveal the inner VC label. The VC label identifies the VPLS service instance to which the packet belongs.
   b. PE C learns the source MAC address in the packet and creates an entry in the FDB table that associates the MAC address to PE A and the VC label that PE A signaled for the VPLS service.
   c. The destination MAC address in the packet is looked up in the FDB table for the VPLS instance. Again, there are two possibilities: either the destination MAC address has already been learned (known MAC address) or it has not been learned on the access side of PE C (unknown MAC address).

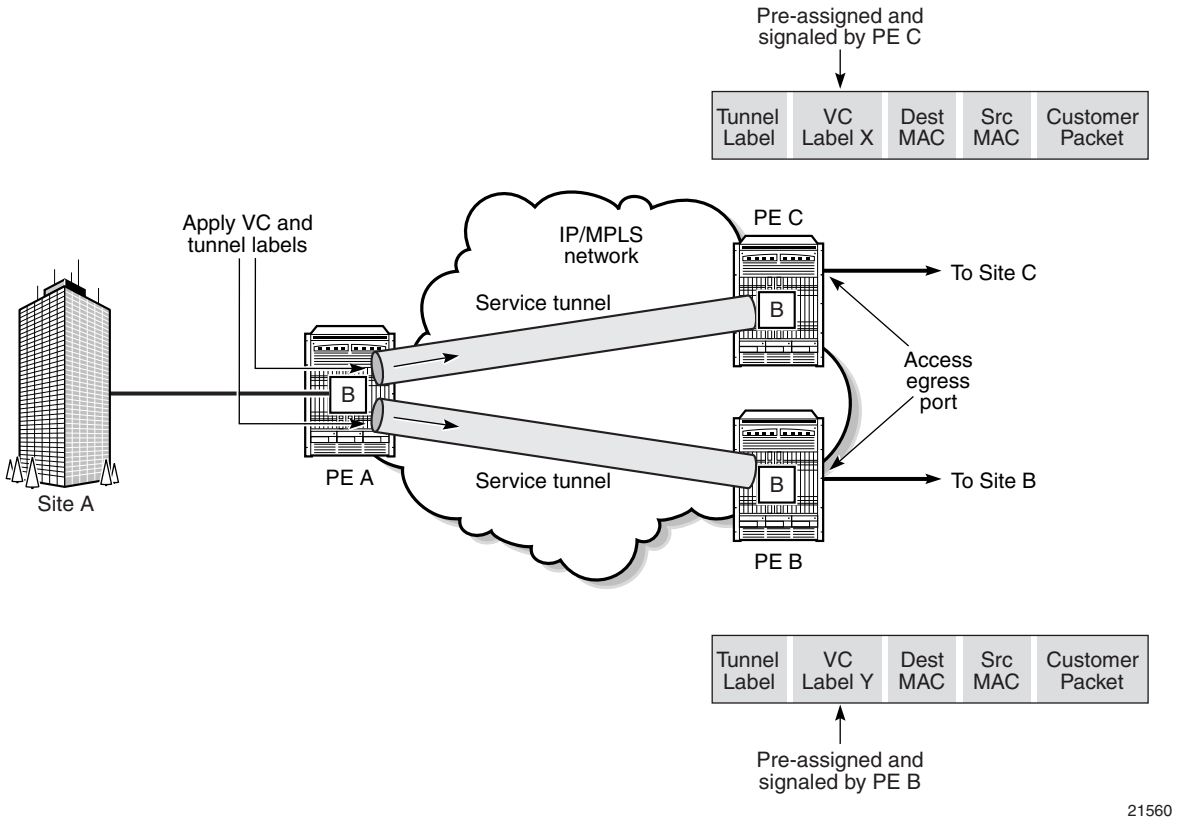
      For a Known MAC Address (Figure 73):

   d. If the destination MAC address has been learned by PE C, an existing entry in the FDB table identifies the local access port and the IEEE 802.1Q tag to be added before sending the packet to customer Location-C. The egress Q tag may be different from the ingress Q tag.

      For an Unknown MAC Address (Figure 73):

   e. PE C will flood packets, as applicable.

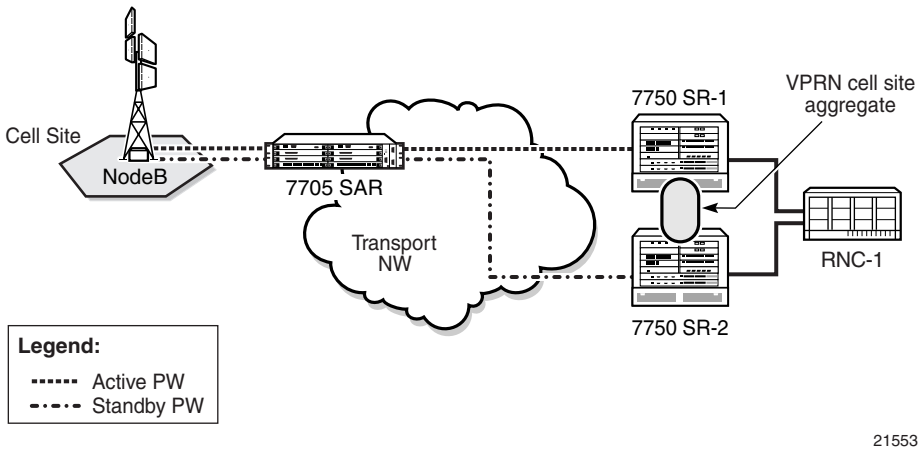
Figure 73: Access Port Egress Packet Format and Lookup



Bridged Mobile Backhaul

Figure 74 shows a PW-based backhaul option for mobile operators, where 7705 SAR-initiated Ethernet PWs terminate at 7750 SR nodes. In most cases, the 7705 SAR-initiated PWs terminate into a VPRN or IES service for routing purposes, or into a VPLS service for MAC forwarding purposes. PW termination into VPLS prevents unwanted exposure of IP addresses and eliminates concerns about the effect of IP addresses that change, thereby avoiding reconfiguration of the VPRN or IES access interfaces and routing entries.

Figure 74: Typical Pseudowire-based Mobile Backhaul



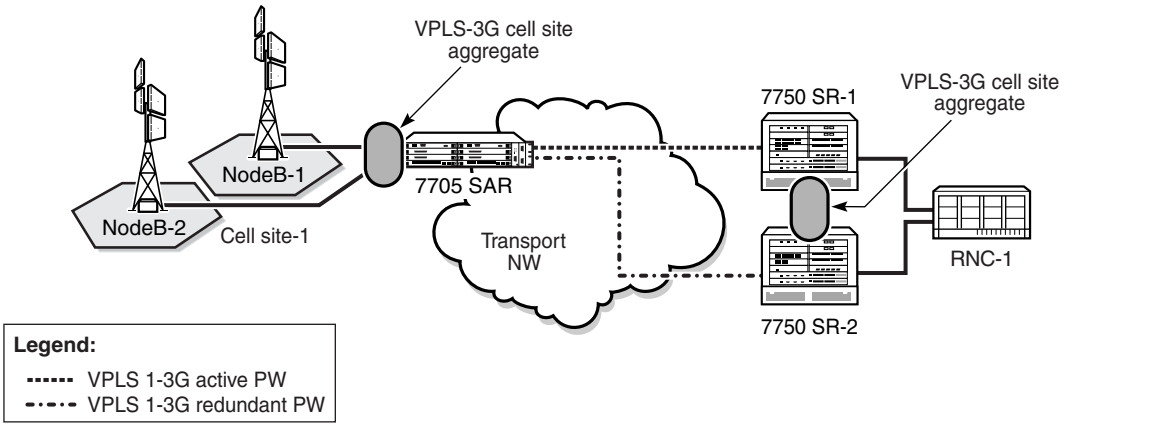
In addition, capacity changes in a radio network could make mobile operators shuffle their transmission links. A simple Layer 2-based backhaul could avoid this complication because the IP addresses are not required to be configured on SAPs (that is, the interfaces facing the base stations or similar equipment), meaning that the 7705 SAR and the backhaul network would not be impacted by mobile layer IP changes. Alternatively, the 7705 SAR implements VPLS to provide any-to-any connectivity at the Layer 2 level and an IP-agnostic network build-out option.

As is the case with VPRN, VPLS also supports multiple virtual forwarding instances. For example, in Figure 75, the 7705 SAR access SAPs facing NodeB-1 and NodeB-2 are bound to VPLS-3G. Another VPLS instance can be configured on the same 7705 SAR for handling eNB 4G traffic. In such a scenario, MAC addresses learned via these two different VPLS instances are stored in separate FDBs ensuring virtualization, which is similar to multiple IP-VPN instances.

Returning to the VPLS-3G example in Figure 75, upon receiving an Ethernet frame from a SAP, the 7705 SAR learns the MAC address and records it together with information from that SAP. If the destination MAC address is known, then the 7705 SAR switches the received Ethernet frame to its destination. If the destination MAC address is not known, then the 7705 SAR floods the frame to all possible destinations that are part of the same VPLS instance (that is, all the SAPs and the network site links).

On the network side, the 7705 SAR supports spoke SDPs to transport customer MAC frames. At ingress, the 7705 SAR strips off the dot1q or qinq header associated with the SAP and switches the Ethernet frame to its destination over the Ethernet PW. Loops can be avoided by using PW redundancy with standby signaling for spoke SDPs and mesh SDPs to ensure proper propagation of broadcast, multicast, and unknown (BMU) frames. Using standby signaling for spoke SDPs, the 7705 SAR ensures that only one spoke is active in the redundant PW deployment model. As a consequence, the 7750 SR disables the spoke SDP binding to VPLS for the standby PWs in order to ensure loop-free operation.

Figure 75: Local VPLS on 7705 SAR in Mobile Backhaul



In the case where the 7705 SAR receives an Ethernet frame from a SAP bound to a VPLS and the destination MAC address is not known, it replicates the frame to all other SAPs that are part of the same VPLS and switches a copy of the frame over all the active Ethernet spoke and mesh SDPs. In Figure 75, the 7705 SAR would switch the incoming frame over an Ethernet PW to 7750 SR-1 after stripping off the incoming dot1q or qinq header.

In terms of label activity, the inner label (the Ethernet PW label for VPLS) identifies the VPLS instance to which the frame belongs, and the outer label identifies the far-end LER node. Using a two-label model means that the traffic from multiple VPLS instances can be transported over a single tunnel between two LER nodes with unique PW labels on a per-VPLS-instance basis.

Upon receiving a VPLS packet, an LER uses the inner label to locate the correct FDB from which to perform MAC lookups. The associated FDB is checked against known and learned MAC addresses. If the lookup is successful, the frame is forwarded to the identified SAP with the appropriate dot1q or qinq header. If the lookup fails, the LER floods the frame to all SAPs that are members of the VPLS instance (that is, the VPLS instance designated by the inner PW label).

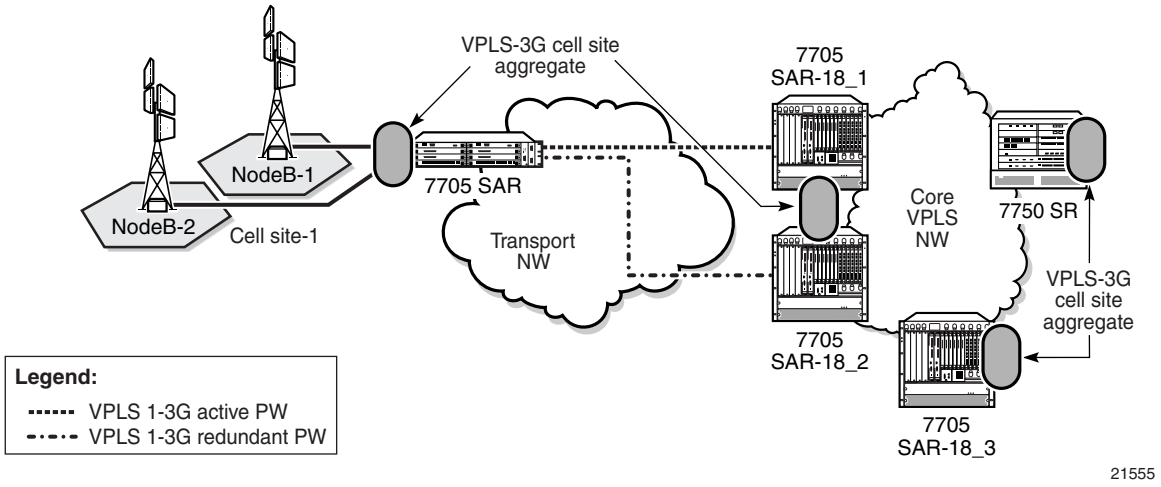
Multi-Tenant Unit (MTU) Termination

Figure 75 can also be used to show how the 7705 SAR can serve as an MTU as described in RFC 4762, section 10.2, to help the scalability of a VPLS core mesh architecture. To function as an MTU, the 7705 SAR is spoke SDP-terminated to a VPLS node (7750 SR node in Figure 75), eliminating the necessity to have a full mesh architecture for all VPLS-enabled nodes. Thus, the mesh requirement is "pushed" to the core nodes only (that is, to the 7750 SR nodes).

The 7750 SR nodes in Figure 75 can be replaced by 7705 SAR nodes running Release 4.0 or later of the 7705 SAR OS. Figure 76 illustrates this scenario, where a 7705 SAR MTU is spoke SDP-terminated to two 7705 SAR-18 nodes (7705 SAR-18_1 and 7705 SAR-18_2).

Using spoke-SDP termination means that it is important that the PW-signaling master node is a 7705 SAR (in Figure 76, the node that initiates the redundant PWs is the cell site 7705 SAR). Thus, only the 7705 SAR-18 that hosts the active spoke SDP will forward the Ethernet traffic to the 7705 SAR and the other 7705 SAR-18 will keep its spoke SDP in the operationally down state. If any failure of the active spoke SDP occurs (that is, if the PW activity switch takes place and the active endpoint of the PW moves from one 7705 SAR-18 to the other one), a mac-flush message is sent, which improves convergence times. In addition, the 7705 SAR-18 nodes can be configured to ignore standby signaling, which improves reconvergence times around failures for services that can tolerate dual-stream reception, such as broadcast TV.

Figure 76: Spoke-SDP Termination to VPLS using 7705 SAR-18 Routers



VPLS Features

Topics in this section include:

- VPLS Enhancements
- Fabric Mode
- Subscriber VLAN
- ATM Encapsulated Residential SAP
- VPLS over MPLS
- VPLS MAC Learning and Packet Forwarding
- Pseudowire Control Word
- Agent Circuit ID Insertion
- MAC Filters
- FDB Table Management
- VPLS and Rate Limiting Via QoS Policy
- MAC Move
- Split Horizon Groups (SAP and Spoke SDP)
- Multicast for VPLS and Routed VPLS (IGMP and MLD Snooping)

VPLS Enhancements

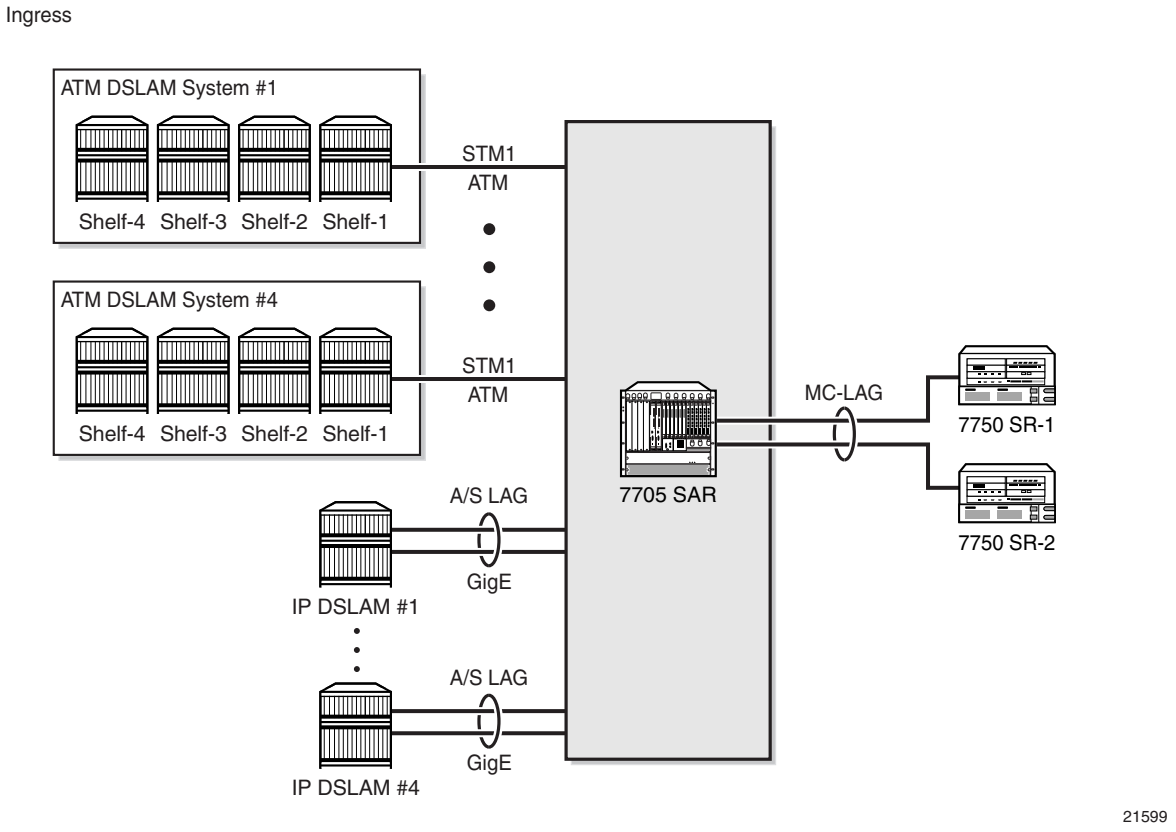
The Alcatel-Lucent VPLS implementation includes several enhancements to basic VPN connectivity. The following VPLS features can be configured individually for each VPLS:

- MAC and IP filter support (up to Layer 4). MAC and IP filters can be applied on a per-SAP ingress and per-SDP ingress (mesh and spoke) basis. IP filters can also be applied on a per-SAP egress basis (Ethernet SAPs only).
- FDB management features on a per-service level, including:
  - → configurable FDB size limit
  - → FDB size alarms
  - → MAC learning disable
  - → discard unknown
  - → separate aging timers for locally and remotely learned MAC addresses
- ingress rate limiting for broadcast, multicast, and (destination) unknown (BMU) flooding on a per-SAP basis
- implementation of STP parameters on a per-VPLS and per-SAP basis

- SHG on a per-SAP and per-spoke SDP basis
- DHCP snooping on a per-SAP basis
- IGMP and MLD snooping on a per-SAP and per-SDP basis
- optional spoke SDP redundancy to protect against node failure

Figure 77 illustrates VPLS enhancements using an example of ATM DSLAM backhaul, where the 7705 SAR might not be used solely for DSLAM backhaul purposes or not all the services might be bound to VPLS. In Figure 77, colocated IP DSLAM (ISAM) traffic can also be transported by the 7705 SAR.

Figure 77: ATM and IP DSLAM Backhaul



Fabric Mode

Similar to IES and VPRN services, to configure a VPLS instance, the fabric mode must be set to aggregate mode (not per-destination mode). Thus, VPLS service is only supported by aggregate-mode fabric profiles. The CLI blocks the creation of a VPLS instance when the fabric mode is set to per-destination. When a VPLS instance is configured, attempting to change the fabric mode to per-destination is blocked.

Subscriber VLAN

The subscriber VLAN feature can be enabled for ATM SAPs bound to a VPLS instance. Subscriber VLAN supports only residential ATM SAPs.

The subscriber VLAN pushes a VLAN tag onto the received bridged PDU on a per-subscriber basis, which helps to uniquely identify subscribers throughout the entire network. After ATM-layer VC termination—where each subscriber has a unique identifier (*port:vpi/vci*)—all the subscribers would be sharing the same uplink. This might present problems to CO IP nodes (such as a BRAS) that want to offer per-subscriber services and identify the subscribers based on dot1q and VLAN tags, which is compatible with the model offered in a native Ethernet model. In order to maintain the uniqueness of a subscriber, a subscriber VLAN tag can be pushed as per the configuration settings (commonly referred to as customer-tag, or c-tag).

A subscriber VLAN has the following characteristics.

- When subscriber VLAN is enabled, a VLAN tag (c-tag) is pushed at ATM ingress and removed at ATM egress. In other words, a symmetrical push/pop operation is supported and cannot be enabled/disabled on a per-direction basis. The exception to this occurs when the Ethernet frame received from the network side does not have any additional VLAN tags; in this case, the received frame is forwarded over the ATM SAP “as is”. That is, there is no pop operation or error message generated.
- The ATM port is always considered to be “NULL”, which means that when a frame is received at ATM egress from a dot1q port (Ethernet SAP to ATM SAP) or from a VLAN vc-type (network), the outer-most VLAN tag is removed (the subscriber tag, or s-tag). If subscriber VLAN is also enabled, the first two outer-most VLAN tags are removed (that is, the s-tag and the c-tag).

Since the ATM port is considered to be “NULL”, when a frame is received at ATM ingress and is going out on a dot1q Ethernet SAP (SAP-to-SAP) or VLAN vc-type (network), a new VLAN tag is pushed (s-tag). If the subscriber VLAN is also enabled, a c-tag and an s-tag are pushed. In short, Ethernet frames at ATM ingress or egress are manipulated in the same way as a null encapsulated Ethernet port.

ATM Encapsulated Residential SAP

For ATM encapsulated residential SAPs:

- the 7705 SAR always transmits the bridge PDU (BPDU) without FCS (PID = 0x00-07)
- the 7705 SAR supports reception of a BPDU both with FCS (PID = 0x00-01) and without FCS (PID = 0x00-07)

VPLS over MPLS

The VPLS architecture proposed in RFC 4664, *Framework for Layer 2 Virtual Private Networks (L2VPNs)* and RFC 4665, *Service Requirements for Layer 2 Provider-Provisioned Virtual Private Networks*, specifies the use of provider equipment (PE) that is capable of learning, bridging, and replicating on a per-VPLS basis. The PE routers that participate in the service are connected using MPLS LSP tunnels in a full mesh composed of mesh SDPs or based on an LSP hierarchy composed of mesh SDPs at the core and spoke SDPs as the access points.

Multiple VPLS instances can be offered over the same set of LSP tunnels. Signaling specified in RFC 4905, *Encapsulation Methods for Transport of Layer 2 Frames over MPLS Networks,* is used to negotiate a set of ingress and egress VC labels on a per-service basis. The VC labels are used by the PE routers for demultiplexing traffic arriving from different VPLS services over the same set of LSP tunnels.

VPLS is provided over MPLS by:

- connecting bridging-capable PE routers with a full mesh of MPLS LSP tunnels
- negotiating per-service VC labels using draft-Martini encapsulation
- replicating unknown and broadcast traffic in a service domain
- enabling MAC learning over tunnel and access ports (see VPLS MAC Learning and Packet Forwarding)
- using a separate forwarding database (FDB) per VPLS service

VPLS MAC Learning and Packet Forwarding

The 7705 SAR edge devices perform the packet replication required for broadcast and multicast traffic across the bridged domain. MAC address learning is performed by the 7705 SAR to reduce the amount of unknown destination MAC address flooding.

7705 SAR routers learn the source MAC addresses of the traffic arriving on their access and network ports. Each 7705 SAR maintains an FDB for each VPLS service instance, and learned MAC addresses are populated in the FDB table of the service. All traffic is switched based on MAC addresses and forwarded between all participating 7705 SAR routers using the LSP tunnels. Unknown destination packets (for example, the destination MAC address has not been learned) are forwarded on all LSPs to the participating 7705 SAR routers for that service until the target station responds and the MAC address is learned by the 7705 SAR associated with that service.

Pseudowire Control Word

The **control-word** command enables the use of the control word individually on each mesh SDP or spoke SDP. By default, the control word is disabled. When the control word is enabled, all VPLS packets are encapsulated along with the control word. The T-LDP control word signaling behavior is the same as that for the control word for VLL services. The configuration at the two endpoints of the VPLS service must match.

Agent Circuit ID Insertion

One of the main applications for VPLS is ATM DSLAM backhaul. DSL operators typically make use of PPPoE over ATM DSL lines for subscriber authentication, authorization, and accounting. When an ATM DSLAM is connected to VPLS service on a 7705 SAR such that the 7705 SAR offers an interworking function for ATM traffic to Ethernet traffic, the 7705 SAR can append the agent-circuit ID to the PPPoE frames received from the ATM DSLAM.

In accordance with RFC 4679, section 3.3.1: Agent-Circuit-ID, agent-circuit ID information can be appended to PPPoE Active Discovery Initiation (PADI) and PPPoE Active Discovery Request (PADR) frames on bridged llc-snap encapsulated SAPs bound to an ATM VPLS instance. Figure 78 illustrates the signaling.

Figure 78: PPPoE Initialization and Agent-Id Push Function

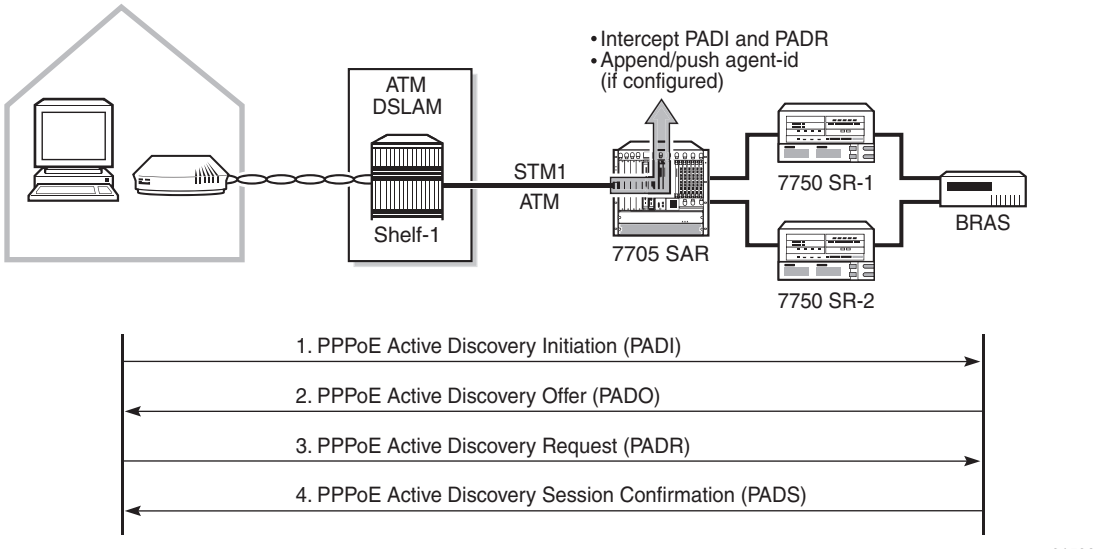
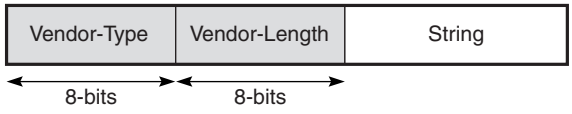


Figure 79 illustrates the agent circuit ID information, where the following definitions apply:

- vendor-type—is always the value 1
- vendor-length—is less than or equal to 65 bits
- string—is the access-node identifier (`atm` *`card/slot/port:vpi.vci`*), which is automatically assigned by the 7705 SAR to be the system-name (hostname)

Figure 79: Agent Circuit ID Information



Appending the agent circuit ID to a PADI or PADR frame is enabled and disabled via the `pppoe-circuit-id` command, which can be issued at the VPLS service and VPLS residential SAP levels. At the service level, the command sets the default value for all SAPs in the VPLS instance. At the SAP level, the command overrides the service level default. If there is a mix of enabled and disabled `pppoe-circuit-id` settings, reissuing the command at the service level will reset all SAPs to the new service level value.

As per the DSL Forum TR-101 April'06 specification, section 3.9.3, any PPPoE vendor-specific tag that may already be present in the received frame is replaced by the 7705 SAR client-id tag.

MAC Filters

MAC filters offer the ability to transport Ethernet frames that match certain criteria over the service to which the frames are bound. The 7705 SAR supports MAC filters at a VPLS ingress SAP and ingress SDP (spoke and mesh). MAC filters can be set to accept or reject the transport of filtered Ethernet frames over the VPLS. Via MAC filters, it is possible to filter traffic received from a defined source or destined for a certain host. MAC filters are the equivalent of IP ACLs, but apply to the Layer 2 MAC layer.

MAC filters support the following fields:

- source MAC address
- destination MAC address
- Ethertype

Any single item or combination of items can be used to define a MAC filter entry. For information on configuring MAC filters, see the "Filter Policies" chapter in the 7705 SAR OS Router Configuration Guide.

FDB Table Management

The following sections describe VPLS features related to management of the FDB, including:

- FDB size
- FDB size alarms
- local and remote aging timers
- Unknown MAC discard

FDB Size

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

- MAC FDB size limits—allows users to specify the maximum number of MAC FDB entries that are learned locally for a SAP or a spoke SDP. If the configured limit is reached, then no new addresses will be learned from the SAP until at least one FDB 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 via configuration). By default, if the destination MAC address is known, it is forwarded based on the FDB, and if the destination MAC address is unknown, it is flooded. Alternatively, if discard unknown is enabled at the VPLS service level, any packets from unknown source MAC addresses are discarded at the SAP.

→ The log event "SAP MAC limit reached" is generated when the limit is reached. When the condition is cleared, the log event "SAP MAC Limit Reached Condition Cleared" is generated.

→ `disable-learning` allows users to disable the dynamic learning function on a SAP or a spoke SDP of a VPLS instance

→ `disable-aging` allows users to turn off aging for learned MAC addresses on a SAP or a spoke SDP of a VPLS instance

FDB Size Alarms

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

Local and Remote Aging Timers

Similar to a Layer 2 switch, learned MACs within a VPLS instance can be aged out if no packets are sourced from the MAC address for a specified period of time (the aging time). In each VPLS instance, there are independent aging timers for locally learned MAC and remotely learned MAC entries in the FDB.

A local MAC address is a MAC address associated with a SAP because it ingressed on a SAP. A remote MAC address is a MAC address received via an SDP from another 7705 SAR 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 and remote MAC addresses, in some circumstances, can take up to two times their respective age timers 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 FDB. When aging is disabled, it is still possible to manually delete or flush learned MAC entries. Aging can be disabled for learned MAC addresses on a SAP or a spoke SDP of a VPLS instance.

Disable MAC Learning

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

Unknown MAC Discard

Unknown MAC discard is a feature that discards all packets that ingress the service whose destination MAC address is not in the FDB. The normal behavior is to flood these packets to all endpoints 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 Via QoS Policy

Traffic that is normally flooded throughout the VPLS can be rate-limited on SAP ingress through the use of service ingress QoS policies. In a service ingress QoS policy, individual queues can be defined per forwarding class to provide shaping of broadcast traffic, MAC multicast traffic and unknown destination MAC traffic.

For more information on QoS policies for broadcast, multicast, and unknown (BMU) traffic, see the “Filter Policies” chapter in the 7705 SAR OS Quality of Service Guide.

MAC Move

The MAC move feature is useful to protect against undetected loops in a VPLS topology when STP is not used. It also protects against the presence of duplicate MACs in a VPLS service.

A sustained high relearn 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 relearn rate of each MAC address. If the rate exceeds the configured maximum allowed limit, MAC move disables the SAP where the source MAC address 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. A SAP can be optionally configured as non-blockable, meaning that when the relearn rate has exceeded the limit, another (blockable) SAP will be disabled instead. By default, all SAPs and spoke SDPs are configured as blockable when MAC-move is enabled.

When MAC-move is enabled and the relearn rate exceeds the maximum limit, the 7705 SAR sends a “Mac move rate for MAC ... exceeded” alarm. This alarm is raised for both blockable and non-blockable SAPs and spoke SDPs. The alarm frequency for non-blockable SAPs and spoke SDPs decreases if the MAC-move condition persists.

The `mac-move` command enables the feature at the service level for SAPs and spoke SDPs. The operation of this feature is the same on the SAP and spoke SDP. For example, if a MAC address moves from SAP to SAP, SAP to spoke SDP, or between spoke SDPs, it will block one of them to prevent thrashing. The relearn rate is computed as the number of times a MAC address moves in a 5 s interval. Therefore, the fastest a loop can be detected and broken is 5 s.

MAC move allows sequential order port blocking. By configuration, some VPLS ports can be configured as “non-blockable”, which allows a simple level of control over which ports are being blocked during loop occurrence.

There are two control mechanisms that allow blocking of ports in a sequential order:

- configuration capabilities to group VPLS ports and to define the order in which they should be blocked
- criteria defining when individual groups should be blocked

For the first mechanism, the configuration CLI is extended by definition of “primary” and “secondary” ports. As the default, all VPLS ports are considered “tertiary” ports unless they are explicitly declared primary or secondary. The order of blocking will always follow a strict order, starting from tertiary, to secondary and then to primary.

The criterion for the second control mechanism is the number of periods during which the given relearn rate has been exceeded. The mechanism is based on the "cumulative" factor for every group of ports. Tertiary VPLS ports are blocked if the relearn rate exceeds the configured threshold during one period, while secondary ports are blocked only when relearn rates are exceeded during two consecutive periods, and so forth. The retry timeout period must be larger than the period before blocking the highest-priority port so it sufficiently spans across the period required to block all ports in sequence. The period before blocking the highest-priority port is the cumulative factor of the highest configured port multiplied by 5 s (the retry timeout can be configured through the CLI).

Split Horizon Groups (SAP and Spoke SDP)

Within the context of VPLS services, a loop-free topology within a fully meshed VPLS core is achieved by applying a split horizon forwarding concept whereby packets received from a mesh SDP are never forwarded to other mesh SDPs within the same service. The advantage of this approach is that no protocol is required to detect loops within the VPLS core network.

In applications such as DSL aggregation, it is useful to extend the split horizon concept to groups of SAPs and/or spoke SDPs. This extension is referred to as a split horizon SAP group or residential bridging.

Traffic arriving on a SAP or a spoke SDP within a split horizon group will not be copied to other SAPs and spoke SDPs in the same split horizon group (but will be copied to SAPs or spoke SDPs in other split horizon groups if these exist within the same VPLS).

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 reused in different VPLS instances. Up to 30 split horizon groups can be defined per VPLS instance. A split horizon group can contain a combination of spoke SDPs and SAPs.

A SAP or spoke SDP can only be added to a split horizon group during its creation. Similarly, a SAP or spoke SDP can be removed from a split horizon group only by its deletion. A split horizon group can be deleted only after all its members have been deleted.

Residential Split Horizon Groups

Residential split horizon groups are supported on ATM SAPs connected to VPLS on 4-port OC3/STM1 Clear Channel Adapter cards. While split horizon groups prevent multicast traffic from flowing within the same group, residential ATM SAPs do not forward broadcast or unknown traffic; they only process known unicast traffic. Residential split horizon groups allow one queue per direction (ingress and egress) for all traffic types (unicast and BMU). OAM processing is also allowed on residential ATM SAPs.

Multicast for VPLS and Routed VPLS (IGMP and MLD Snooping)

IGMP and MLD snooping allows a 7705 SAR to listen to the IGMP traffic between hosts and routers. The 7705 SAR extracts information from the traffic to create and maintain a multicast forwarding information base (MFIB) to track which hosts want which IP multicast streams. Multicasts may be filtered to control which ports receive specific multicast traffic.

For example, service providers that use a flat IP network to deliver video over a mobile backhaul network can take advantage of Layer 2 services (VPLS) to save IPv4 addresses. A Layer 2 domain with *n* nodes needs *n* IP addresses, whereas a point-to-point connections requires 2×*n* addresses. Service providers using VPLS and IGMP or MLD snooping to relay IGMP or MLD requests to uplink (network) PIM interfaces will save addresses.

This section contains information on the following topics:

- Sample Scenarios
- Group and Addressing Support
- IP Multicast in r-VPLS
- Multicast Router Ports
- Tagged Access Traffic
- Hardware Support

For more information on multicast, refer to the “IP Multicast” section in the 7705 SAR OS Routing Protocols Guide.

Sample Scenarios

Figure 80 shows a typical deployment. Host traffic arrives at the routed VPLS (r-VPLS), where IGMP or MLD snooping extracts all IGMP or MLD packets and sends them to the CSM, and from the CSM the packets are forwarded via PIM to the head-end 7750 SR nodes. The VPLS multicast FIB (MFIB) tracks all the IGMP or MLD join requests in an internal 7705 SAR forwarding table. On the upstream nodes, PIM builds the multicast tree from the 7705 SAR to the 7750 SR. In the reverse direction, the video source sends multicast traffic, which is forwarded by the PIM nodes to the addresses in the previously built multicast tree. As traffic from various sources arrives at the 7705 SAR, the r-VPLS MFIB directs each multicast stream to the correct eNodeB.

Figure 80: Video over Mobile Backhaul

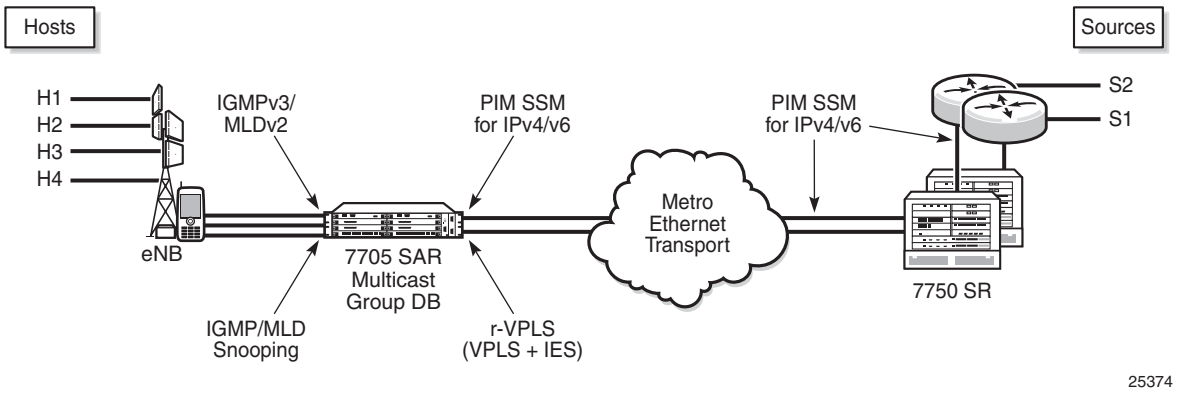


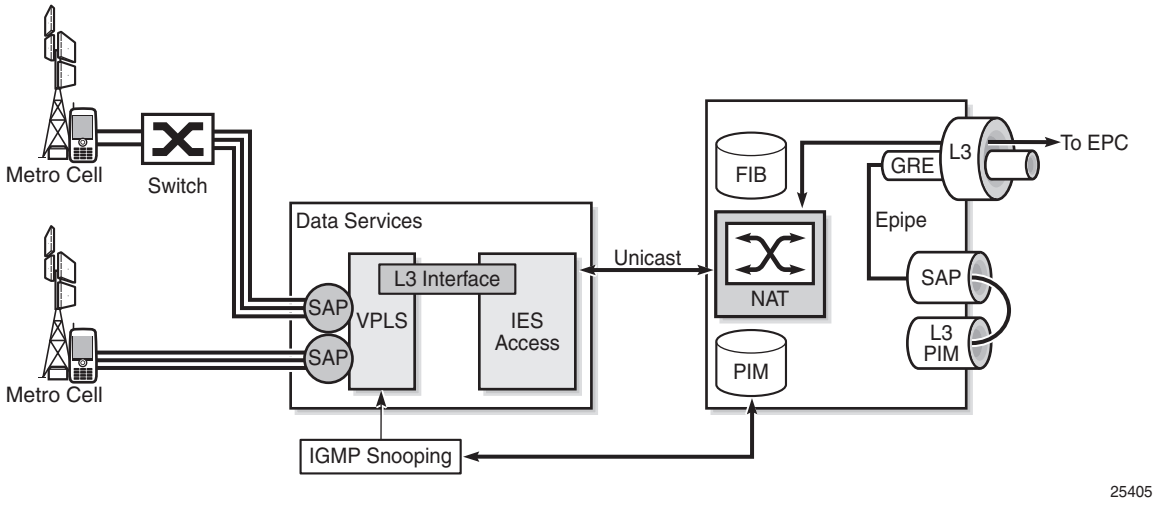
Figure 81 shows another example of IGMP and MLD snooping, where service providers can offer evolved multimedia broadcast multicast services (eMBMS) on their metro cell network (data services).

In Figure 81, the data VPLS performs IGMP or MLD snooping to build the MFIB. The extracted IGMP and MLD requests are forwarded via PIM over an Epipe or, preferably, via PIM over a Layer 3 spoke SDP to remove the external physical connection between two ports from the 7705 SAR. The traffic between IES access and NAT is unicast traffic. The Layer 3 spoke-SDP traffic is transported over a GRE tunnel via the Internet to the evolved packet core (EPC), where a secure gateway forwards the PIM join message to the multicast source servers. The GRE logical Layer 3 spoke SDP does not need to be part of the NAT function; if it is not, this logical interface must obtain its own public interface IP address.

Figure 81 also shows the typical metro cell deployment, where IGMP snooping is done on the r-VPLS of the data IES service. The IGMP join messages translate to PIM SSM via the uplink network interface, as described at the beginning of this section.

MLD snooping allows support for IPv6 addresses through the use of r-VPLS for IPv6, allowing the network design for IPv4 and IPv6 domains to be the same. MLD snooping allows support for IPv6 addresses through the use of r-VPLS for IPv6.

Figure 81: Metro Cell Multicast



Group and Addressing Support

This section contains information on the following topics:

- IPv4 Multicast Support
- IPv6 Layer 2 Multicast Support and Group Address

IPv4 Multicast Support

The 7705 SAR supports (S,G) and (*,G) for IPv4 multicast in Layer 2 services only, including Layer 2 services within the context of VPLS and r-VPLS.

7705 SAR supports PIM-SSM only. For IPv4 multicast services, PIM SSM requires SSM translation in the r-VPLS interface context for (*,G) joins.

IPv6 Layer 2 Multicast Support and Group Address

The 7705 SAR supports (S,G) and (*,G) for IPv6 multicast in Layer 2 services only, including Layer 2 services within the context of VPLS and r-VPLS, and uses the MAC-format group-addressing scheme to minimize the size of the MFIB.

The multicast MAC-format group address consists of the IPv6 multicast prefix (33:33) and the four least significant bytes of the IPv6 address. In the sample CLI display below, for the IPv6 address FF04::1:FFFF:0011, the representation of the group address is 33:33:FF:FF:00:11.

```
Multicast FIB, Service 50
===============================================================================
Source Address  Group Address         Sap/Sdp Id               Svc Id   Fwd/Blk
-------------------------------------------------------------------------------
*               33:33:FF:FF:00:11     sap:1/1/1:50             Local    Fwd
===============================================================================
```

7705 SAR supports PIM-SSM only. For IPv6 multicast services, PIM SSM requires SSM translation in the r-VPLS interface context for (*,G) joins.

IP Multicast in r-VPLS

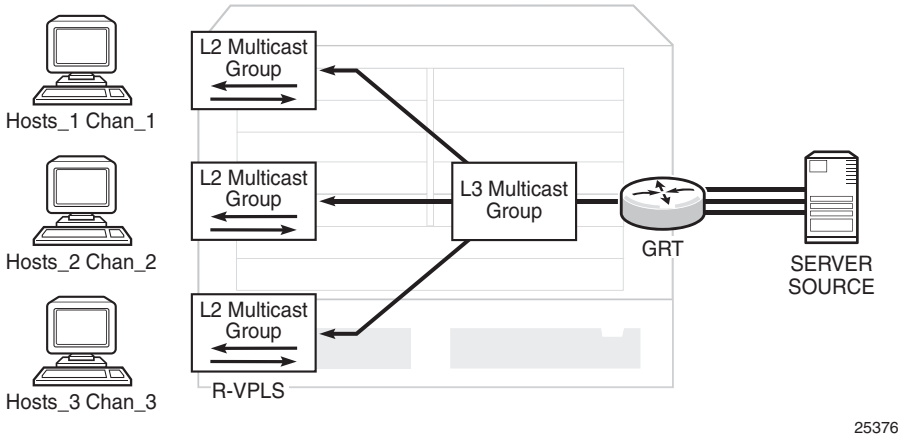
When creating a Layer 2 multicast service in the context of an r-VPLS with PIM configuration on the r-VPLS Layer 3 interface, the 7705 SAR creates two multicast groups: one Layer 2 multicast group and one Layer 3 multicast group. Once the Layer 2 group is created, the Layer 3 group is created automatically. The 7705 SAR uses one Layer 3 multicast group per source, and one Layer 2 multicast group per source per VPLS.

Note: IP multicast on r-VPLS is supported only if IGMP or MLD snooping is enabled on the VPLS associated with the IES. That is, configuring IGMP or MLD on the IES that is associated with the VPLS without configuring snooping on the VPLS will not flood multicast traffic out the VPLS.

Figure 82 and Figure 83 illustrate how Layer 2 multicast interacts with Layer 3 multicast.

In Figure 82, there are three hosts and one channel. The Layer 3 multicast group forwards source traffic to each port on which there is a corresponding Layer 2 multicast group. All three Layer 2 groups are within the same r-VPLS. To configure this scenario, create three Layer 2 groups on the 7705 SAR (one for each host), and one Layer 3 group for the source. The single Layer 3 multicast group streams the multicast traffic to all three hosts.

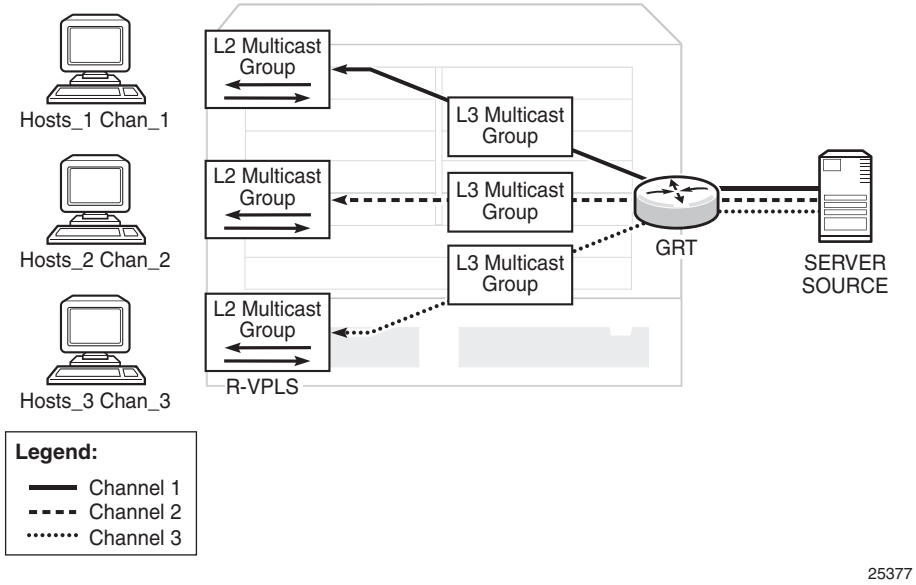
Figure 82: Multiple Hosts in an r-VPLS Receiving the Same Channel



In Figure 83, there are three hosts and three channels. Each host connects to a different port and wants to receive a different (S,G) group (that is, a different channel). Thus, three layer 3 (S,G) groups are needed. To achieve this scenario, create three Layer 2 multicast groups (one for each host) and three Layer 3 multicast groups (one for each channel).

Note: For r-VPLS, the 7705 SAR supports multicast data flows only from a Layer 3 to a Layer 2 domain. For example, in Figure 83, multicast data can only flow from the server source in the Layer 3 domain to the hosts in the Layer 2 domain. Currently, the 7705 SAR does not support multicast data flow from a Layer 2 to a Layer 3 domain.

Figure 83: Multiple Hosts in an r-VPLS Receiving Different Channels



IPv6 Multicast Forwarding Behavior in r-VPLS

In general, the behavior of IPv6 multicast in r-VPLS is as follows.

- Multicast in Layer 2 is only supported by (*,G). That is, (S,G) in Layer 3 gets translated to (*,G) in Layer 2.
- Multicast in Layer 3 is only supported by (S,G).

The IPv6 multicast control plane behavior is as follows.

- (S,G) is forwarded from the Layer 2 snooping interface to PIM without translation to (*,G).
- The Layer 2 multicast forwarding table is built based on (*,G) and the MAC-format multicast-group address scheme, using the four least significant bytes of the IPv6 address (see IPv6 Layer 2 Multicast Support and Group Address)
- The Layer 3 multicast forwarding table is built based on (S,G) and the full IPv6 multicast group address.

The IPv6 multicast data plane behavior is as follows.

- The Layer 2 forwarding table (see bullet 2 above) is downloaded to the data plane.
- The Layer 3 forwarding table (see bullet 3 above) is downloaded to the data plane.
- Layer 3 multicast for IPv6 supports the entire range of multicast addresses.
- The Layer 2 multicast address is limited and unique to prefix /96. That is, only the four least significant bytes of the IPv6 multicast address are used. Note the following items.
  - → It is useful to keep the multicast table small at the network edge, where multicast groups can be effectively managed via 32-bit (4-byte) addressing.
  - → A 32-bit multicast address can provide 4 bytes of multicast group addressing.
  - → Optionally, multicast zones can be created on the access side with overlapping 32-bit addresses, but in the core—where the entire IPv6 multicast group is available—multicast zones can guide traffic correctly to the corresponding access group.

Multicast Router Ports

Membership reports are only sent to multicast router (mrouter) ports. An mrouter port is a port through which membership queries are received. An mrouter port can be configured manually on a VPLS SAP or SDP using the `mrouter-port` command under `igmp-snooping` or `mld-snooping`.

Tagged Access Traffic

The 7705 SAR processes tagged querier requests arriving on a null-encapsulated port and installs the querier message. This means that an IGMP or MLD router is recognized to exist on that port and reports (joins and leaves) will be forwarded out that port.

Similarly, for multicast data, the 7705 SAR processes tagged multicast traffic arriving on a null-encapsulated port according to the MFIB.

Hardware Support

Multicast VPLS and r-VPLS are supported on the following 7705 SAR hardware:

- 8-port Ethernet Adapter card, version 2
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card
- 6-port Ethernet 10Gbps Adapter card
- Packet Microwave Adapter card
- standalone platforms, including the 7705 SAR-M, 7705 SAR-A, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-W, 7705 SAR-Wx, and 7705 SAR-X
- 7705 SAR-M modules, including the GPON module, 8-port xDSL module, and 6-port DSL Combination module
- 4-port SAR-H Fast Ethernet module
- hardware that supports PPP/MLPPP for uplink spoke SDPs

Routed VPLS

Topics in this section include:

- IES or VPRN IP Interface Binding
- IP Interface MTU and Fragmentation
- ARP/ND and VPLS FIB Interactions
- The allow-ip-int-binding VPLS Flag
- DSCP Marking
- VPLS SAP Ingress IP Filter Override
- Routed VPLS Supported Routing-Related Protocols

Hosts within the same subnet communicate directly with each other without the need for a router, but any communication with a host that is external to the subnet requires routing. With routed VPLS, you can use bridging for local destinations when possible and routing for non-local destinations that cannot be reached directly.

Routed VPLS appears similar to a LAN Ethernet switch and a router. The VPLS instance on the 7705 SAR node grants Ethernet switch functionality among connected nodes. When the destination IP is not local, the 7705 SAR routes the traffic by means of the VPRN or the IES instance.

Routed VPLS is enabled for IPv4 and IPv6 forwarding.

IES or VPRN IP Interface Binding

A standard IP interface within an existing IES or VPRN service context can be bound to a VPLS service. A VPLS service only supports binding for a single IP interface.

Although an IP interface can only be bound to a single VPLS service, the routing context containing the IP interface (IES or VPRN) can have other IP interfaces bound to other VPLS service contexts.

Topics in this section include:

- Assigning a Service Name to a VPLS Service
- Binding a Service Name to an IP Interface
- Removing a Bound VPLS Service or Service Name
- IP Interface and VPLS Operational State Coordination

Assigning a Service Name to a VPLS Service

If a service name is applied to a VPLS 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.

If the `config>service>vpls>allow-ip-int-binding` command is enabled on the VPLS service, when the service name is applied to the VPLS service, the system will scan the existing IES and VPRN services for an IP interface that is bound to the specified service name. If found, the IP interface will be attached to the VPLS service associated with the service name. Only one interface can be bound to the specified service name.

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

Binding a Service Name to an IP Interface

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

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

- if the service name is not currently in use by a service, the IP interface will be placed in an operationally down: Non-existent service name or inappropriate service type state
- if the service name is currently in use by a VPLS service without the allow-ip-int-binding flag set, the IP interface will be placed in the operationally down: VPLS service allow-ip-int-binding flag not set state
- if the service name is currently in use by a valid VPLS service and the allow-ip-int-binding flag is set, the IP interface will be attached to the VPLS service

Removing a Bound VPLS Service or 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.

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

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

IP Interface and VPLS Operational State Coordination

If 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 remains down until at least one virtual port (SAP, spoke SDP or mesh SDP) is operational.

IP Interface MTU and Fragmentation

The VPLS service is affected by two MTU values: port MTUs and the VPLS service MTU. The MTU on each physical port defines the largest Layer 2 packet (including all DLC headers) that may be transmitted out of a port. The VPLS itself has a service-level MTU that defines the largest packet supported by the service. This MTU does not include the local encapsulation overhead for each port (dot1q, qinq, or SDP service delineation fields and headers) but does include the remainder of the packet. As virtual ports are created in the system, any given virtual port cannot become operational unless the configured port MTU minus the virtual port service delineation overhead is greater than or equal to the configured VPLS service MTU. This ensures that an operational virtual port can 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.

If 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 the VPLS service MTU. The operational IP MTU cannot be greater than the VPLS service MTU minus 14, 18, or 22 bytes (for null, dotq1, or qinq port encapsulations, respectively) to account for the Layer 2 headers and tags.

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.

The VPLS service MTU and the IP interface MTU parameters can be changed at any time.

ARP/ND and VPLS FIB Interactions



Note: References to ARP in this section also apply to Neighbor Discovery (ND) protocol. ARP applies to IPv4. ND protocol applies to IPv6.

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. If the destination IP address in the routed packet is a host on the local subnet represented by the VPLS instance, the destination IP address 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. 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 ARP cache can be populated with a static ARP entry for the next-hop IP address. Dynamically populated ARP entries will age out according to the ARP aging timer; static ARP entries never age out.

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

Routed VPLS Specific ARP/ND Cache Behavior

Note: References to ARP in this section also apply to Neighbor Discovery (ND) protocol. ARP applies to IPv4. ND protocol applies to IPv6.

In typical routing behavior, the system uses the IP routing table to select the egress interface, and at the egress forwarding engine, an ARP entry is used to forward the packet to the appropriate Ethernet MAC address. With routed VPLS, the egress IP interface can be represented by multiple egress forwarding engines (wherever the VPLS service virtual ports exist). To optimize routing performance, the ingress forwarding engine performs an ingress ARP lookup in order to resolve which VPLS MAC address the IP frame must be routed towards.Table 57 describes how the ARP cache and MAC FIB entry states interact at ingress and Table 58 describes the corresponding egress behavior.

Table 57: Ingress Behavior for VPLS Next-Hop Routing

Layer 3 Next-Hop ARP Cache Entry	Next-Hop MAC FIB Entry	Ingress Behavior
ARP Cache Miss (No Entry)	Known or Unknown	Flood to all egress forwarding engines associated with the VPLS context
ARP Cache Hit	Known	Forward to specific egress forwarding engine associated with VPLS virtual port
	Unknown	Flood to all egress forwarding engines associated with the VPLS for forwarding out to all VPLS virtual ports

Table 58: Egress Behavior for VPLS Next-Hop Routing

Layer 3 Next-Hop ARP Cache Entry	Next-Hop MAC FIB Entry	Egress Behavior
ARP Cache Miss (No Entry)	Known	No ARP entry. The MAC address is unknown and the ARP request is flooded out to all virtual ports of the VPLS instance.
	Unknown	ARP processing request transmitted out to all virtual ports associated with the VPLS service. Only the first egress forwarding engine ARP processing request triggers the egress ARP request.
ARP Cache Hit	Known	Forward out to specific egress VPLS virtual port where MAC address has been learned
	Unknown	Flood to all egress VPLS virtual ports on forwarding engine

The allow-ip-int-binding VPLS Flag

The allow-ip-int-binding flag on a VPLS service context informs the system that the VPLS service is enabled for routing support. The system uses the setting of the flag as a key to determine what type of ports and forwarding planes the VPLS service can 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 if routing support is enabled.

Once 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. ATM SAPs are not supported.

VPLS Feature Restrictions (due to allow-ip-int-binding)

If the allow-ip-int-binding flag is set on a VPLS service, the following features are disabled:

- MAC filters
- residential SHG
- DHCP
- mVPLS
- mac-subnet-length
- GRE SDP (cannot be bound to the VPLS)

Note: The DHCP relay functionality under `service>ies>if>dhcp` can be used to dynamically assign IP addresses to the clients connected to routed VPLS SAPs.

DSCP Marking

Note: The 7705 SAR does not support ingress DSCP marking.

Egress DSCP re-marking is supported on routed VPLS service for bridged packets only. It is not supported for packets routed out from a VPLS SAP.

The egress re-marking defined in the SAP egress QoS policy is not performed for packets that are routed out an egress VPLS SAP.

VPLS SAP Ingress IP Filter Override

If an IP interface is attached to a VPLS service context, the VPLS SAP provisioned IP filter for ingress routed packets can be optionally overridden in order to provide special ingress filtering for routed packets. This allows different filtering for routed packets and non-routed packets. The filter override is defined on the IP interface bound to the VPLS service name. A separate override filter can be specified for IPv4 and IPv6 packet types.

If a filter for an IPv4 or IPv6 packet type is not overridden, the SAP-specified filter is applied to the packet (if defined).

Routed VPLS Supported Routing-Related Protocols

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

- BGP (IPv4 and IPv6)
- OSPF (IPv4 only)
- static routing (IPv4 and IPv6)
- BFD (IPv4 and IPv6)
- VRRP (IPv4 only)
- ARP
- ND protocol
- DHCP (IPv4 and IPv6)

VPLS and Spanning Tree Protocol

Topics in this section include:

- VPLS Redundancy
- VPLS Access Redundancy
- MAC Flush Message Processing

The Alcatel-Lucent VPLS service provides a bridged or switched Ethernet Layer 2 network. Equipment connected to SAPs or spoke SDPs forward Ethernet packets into the VPLS service. The 7705 SAR participating in the service learns where the customer MAC addresses reside on ingress SAPs or ingress SDPs.

Unknown destinations, broadcasts, and multicasts are flooded to all other SAPs or spoke SDPs in the service. If SAPs or spoke SDPs are connected together, either through misconfiguration or for redundancy purposes, loops can form and flooded packets can keep flowing through the network. The Alcatel-Lucent implementation of STP is designed to remove these loops from the VPLS topology. This is done by putting one or more SAPs or spoke SDPs in the discarding state.

STP parameters allow a balance between resiliency and speed of convergence extremes. Modifying particular parameters can affect the behavior. For information on command usage, descriptions, and CLI syntax, refer to Configuring a VPLS Service with CLI.

Each VPLS instance on the 7705 SAR operates in Rapid Spanning Tree Protocol (RSTP) mode and is compliant with IEEE 802.1D-2004 - default mode.

VPLS Redundancy

The VPLS standard (RFC 4762, *Virtual Private LAN Services Using LDP Signalling*) includes provisions for hierarchical VPLS using point-to-point spoke SDPs. Two applications have been identified for spoke SDPs:

- connecting MTUs to PEs in a metro area network
- interconnecting the VPLS nodes of two metro networks

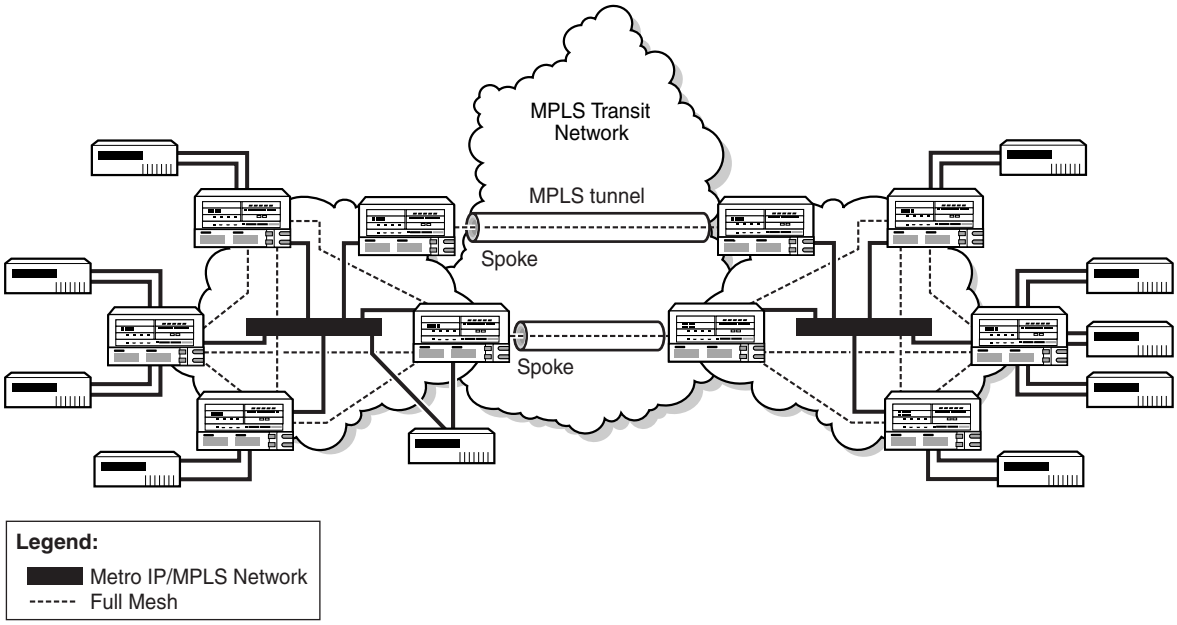
In both applications, the spoke SDPs improve the scalability of VPLS. Since node redundancy is implicit in non-hierarchical VPLS services (using a full mesh of SDPs between PEs), node redundancy for spoke SDPs needs to be provided separately.

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

Spoke SDP Redundancy for Metro Interconnection

When two or more meshed VPLS instances, such as in Figure 84, are interconnected by redundant spoke SDPs, a loop in the topology results. In order to remove a topology loop, STP can be run over the SDPs (links) that form the loop, which then blocks one of the SDPs.

Figure 84: H-VPLS with Spoke Redundancy



To avoid the inefficiency of running STP separately in every VPLS in the topology, the node can associate a number of VPLS services with a single STP instance running over redundant SDPs. Node redundancy is achieved by running STP in one VPLS, and then 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.

If the active node fails, STP on the mVPLS on the standby node changes the link state from disabled to active. The standby node broadcasts a MAC flush LDP control message in each of the protected VPLS instances so that the address of the newly active node can be relearned by all PEs in the VPLS.

It is possible to configure two mVPLS services, where both mVPLS services have different active spokes (this is achieved by changing the path cost in STP). Load balancing across the spokes is achieved by associating different user VPLS services with the two mVPLS services.

Spoke SDP-Based Redundant Access

This feature provides the ability to have a node deployed as MTU switches to be multi-homed for VPLS to multiple routers deployed as PEs without requiring the use of mVPLS.

In the configuration example displayed in Figure 84, the MTUs have spoke SDPs to two PE devices. One is designated as the primary and one as the secondary spoke SDP. This is based on a precedence value associated with each spoke.

The secondary spoke is in a blocking state (both on receive and transmit) as long as the primary spoke is available. If the primary spoke becomes unavailable (due to link failure, PEs failure, and so on), the MTU immediately switches traffic to the backup spoke and starts receiving traffic from the standby spoke. Optional revertive operation (with configurable switch-back delay) is supported. Forced manual switchover is also supported.

To speed up the convergence time during a switchover, MAC flush is configured. The MTUs generate a MAC flush message over the newly unblocked spoke when a spoke change occurs. As a result, the PEs receiving the MAC flush will flush all MACs associated with the impacted VPLS instance and forward the MAC flush to the other PEs in the VPLS network if `propagate-mac-flush` is enabled.

VPLS Access Redundancy

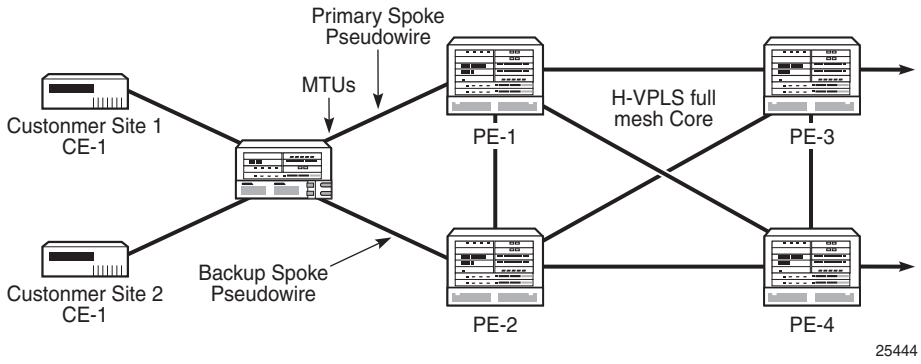
Another application of hierarchical VPLS uses MTUs that are not MPLS-enabled and must have Ethernet links to the closest PE node. To protect against failure of the PE node, an MTU can be dual-homed and have two SAPs on two PE nodes.

On the 7705 SAR, the mechanism used to resolve a loop in an access circuit uses STP-based access, with or without mVPLS.

STP-Based Redundant Access to VPLS

In the configuration shown in Figure 85, STP is activated on the MTU and two PEs in order to resolve a potential loop. STP needs to run only in a single VPLS instance, and the results of the STP calculations are applied to every VPLS on the link. In this configuration, the scope of the STP domain is limited to the MTU and PEs but any topology change must be propagated across the whole VPLS domain, including mesh SDPs. This is done with MAC flush messages defined by RFC 4762.

Figure 85: Dual Homed MTUs in Two-Tier Hierarchical VPLS



When STP is used as a loop resolution mechanism, every Topology Change Notification (TCN) received in an STP instance is translated into a MAC flush message request to clear all FDB entries except the ones learned from the originating PE. MAC flush messages are sent to all PE peers connected through mesh and spoke SDPs in the context of the VPLS services managed by the given STP instance.

MAC Flush Message Processing

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

The 7705 SAR 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 the 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-mine requests the clearing of all FDB entries learned from the originating PE. This means that this message has the opposite effect of the flush-all-but-mine message. This type is not included in the RFC 4762 definition and is implemented using vendor-specific TLV.

Upon reception of MAC flush messages (regardless of the type), the 7705 SAR PE takes the following actions:

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

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

- The flush-all-but-mine message is received from an LDP peer and the propagate-mac-flush flag is enabled. The message is sent to all LDP peers in the context of the VPLS service in which it was received.
- A flush-all-but-mine message is generated when a switchover occurs between spoke SDPs of the same endpoint. The message is sent to the LDP peer connected through the newly active spoke SDP.
- A flush-all-but-mine message is generated when a TCN message is received in an STP context and the propagate-mac-flush flag is enabled. The message is sent to all LDP peers connected by spoke and mesh SDPs in the context of the VPLS service controlled by the STP instance, as determined by the mVPLS definition.

  If all LDP peers are in the STP domain, it means that the mVPLS and the VPLS both have the same topology and the 7705 SAR will not send any flush-all-but-mine messages. If there are VPLS LDP peers outside the STP domain, the router sends flush-all-but-mine messages to all its VPLS peers. When a TCN occurs in the Layer 2 domain, the MAC flush message is propagated over spoke SDPs.

  The 7705 SAR will only send a withdrawal request if the mVPLS contains a mesh SDP.

The flush-mine message is generated under the following conditions:

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

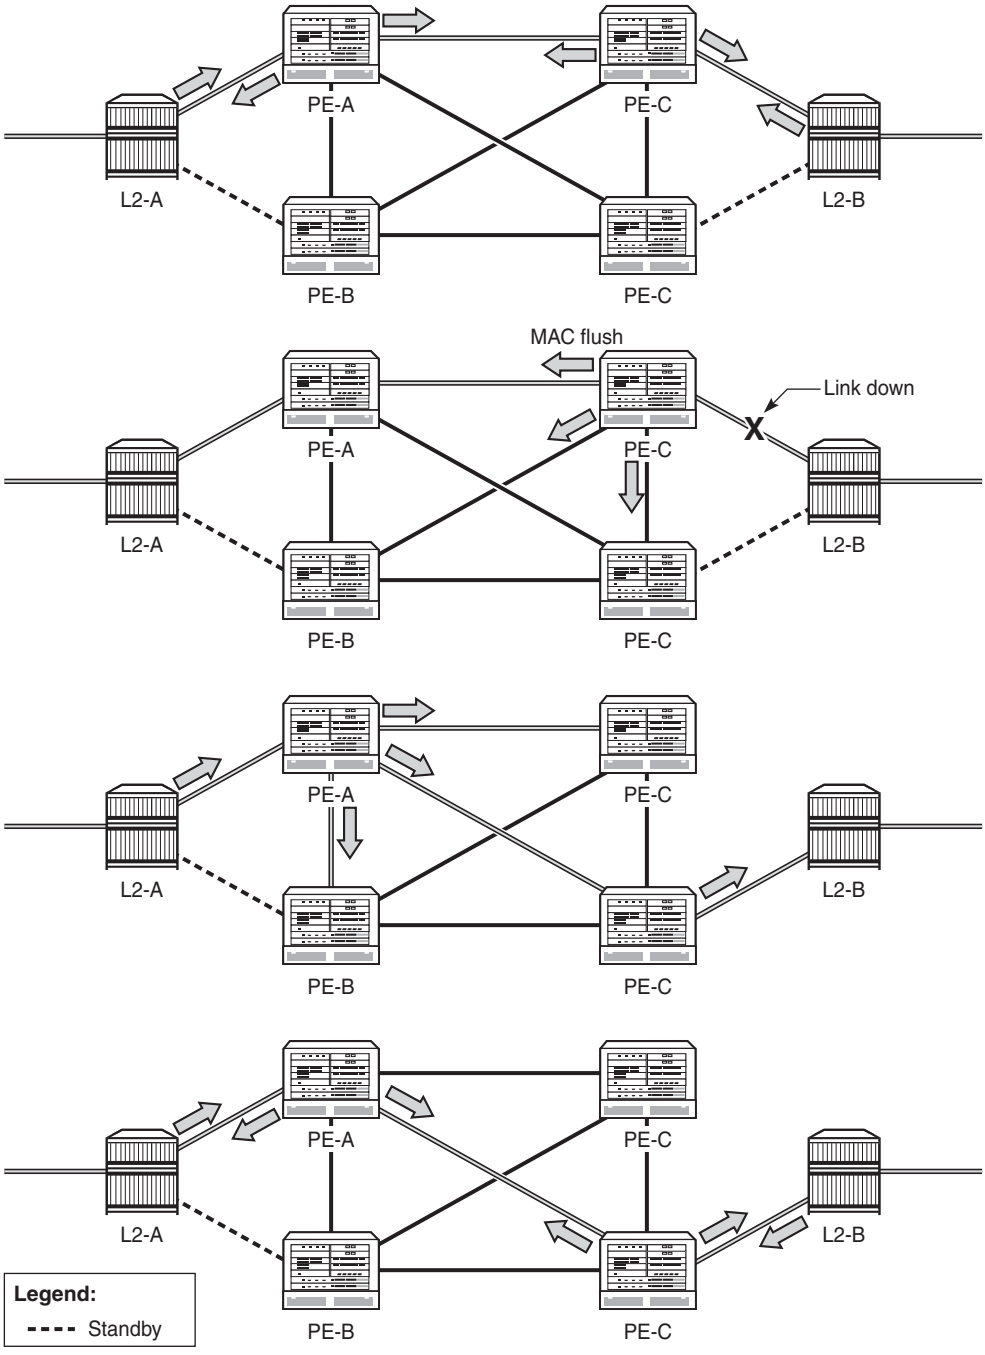
Dual Homing to a VPLS Service

Figure 86 illustrates a dual-homed connection to a VPLS service (PE-A, PE-B, PE-C, PE-D) and the operation in case of link failure (between PE-C and L2-B). Upon detection of a link failure, PE-C sends MAC Address Withdraw messages, which indicate to all LDP peers that they should flush all MAC addresses learned from PE-C. This triggers packets to be broadcast addressing the affected hosts and a relearning process in case an alternative route exists.

The MAC Address Withdraw message is different from the message described in draft-ietf-l2vpn-vpls-ldp-xx.txt, *Virtual Private LAN Services over MPLS*. The difference is in the interpretation and action performed at the receiving PE. According to the draft definition, upon receipt of a MAC withdraw message, all MAC addresses, except the ones learned from the source PE, are flushed. In the 7705 SAR implementation, upon receipt of the MAC Address Withdraw message, all MAC addresses learned from the source are flushed. In this implementation, this message is an LDP address message with vendor-specific TLV, and is called the flush-all-from-ME message.

The message as defined in the draft definition is currently used in any mVPLS that is using RSTP to recover from failures in Layer 2 topologies. The advantage of the alternative messaging behavior over RSTP-based methods is that only MAC-affected addresses are flushed, and not the full forwarding database. This method does not provide a mechanism to secure alternative loop-free topology. However, the convergence time depends on how quickly the given CE device opens the alternative link (L2-B switch in Figure 86) as well as how quickly the PE routers flush their FDBs. Additionally, since this method relies on reacting to the physical failure of the link, it is effective only if the PE and CE are directly connected with no hub or bridge.

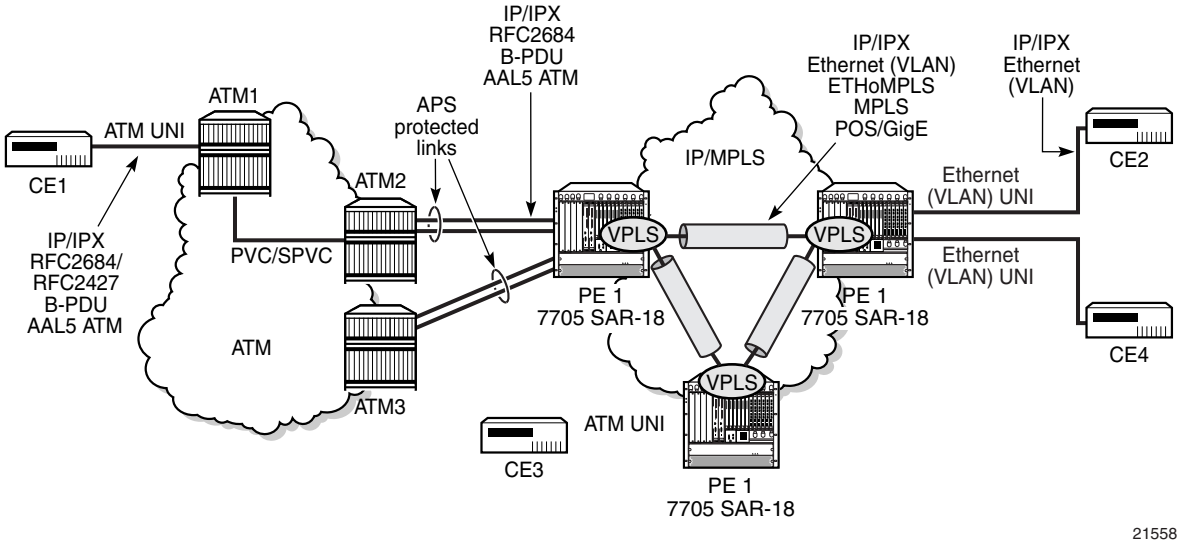
Figure 86: Dual-Homed CE Connection to VPLS



ATM PVC Access and Termination on a VPLS Service

The application shown in Figure 87 provides access to a VPLS service for ATM users connected either directly or through an ATM access network to a 7705 SAR PE node. The 7705 SAR supports an ATM VC-delimited SAP terminating on a VPLS service.

Figure 87: Example of ATM PVC Access and Termination on a VPLS



RFC 2427-encapsulated or RFC 2684-encapsulated untagged Ethernet/802.3 frames (with or without Frame Check Sequence (FCS)) or BPDUs from a customer's bridge device are received on a given SAP over an ATM interface on the 7705 SAR. The ATM-related encapsulation is stripped, and the frames (without FCS) are forwarded towards destination SAPs either locally or using SDPs associated with the VPLS service (as dictated by destination MAC address VPLS processing). In the egress direction, the received untagged frames are encapsulated into RFC 2427 or RFC 2684 (no Q-tags are added, no FCS in the forwarded frame) and sent over the ATM VC towards the customer CPE.

When AAL5 RFC 2427/2684 encapsulated tagged frames are received from the customer's bridge on an ATM SAP, the tags are transparent and the frames are processed as described above, with the exception that the frames forwarded towards the destination(s) will have the received tags preserved. Similarly, in the egress direction, the received tagged Ethernet frames are encapsulated as is (Q-tags are again transparent and preserved) into RFC 2427/2684 and sent over the ATM PVC towards the customer CPE.

Since the tagging is transparent, the 7705 SAR performs unqualified MAC learning (for example, MAC addresses are learned without reference to the VLANs they are associated with). Hence, MAC addresses used must be unique across all the VLANs used by the customer for a given VPLS service instance. If a customer wants a per-VLAN separation, then the VLAN traffic that needs to be separated must travel on different VCs (different SAPs) associated with different VPLS service instances.

All VPLS functionality available on the 7705 SAR is applicable to ATM-delimited VPLS SAPs. For example, bridged PDUs received over an ATM SAP can be tunneled through or dropped, all Forwarding Database (FDB) functionality applies, packet-level QoS and MAC filtering applies. Also, split horizon groups are applicable to ATM SAPs terminating on VPLS. In other words, frame forwarding between ATM SAPs, also referred to as VCI-to-VCI forwarding, is disabled within the same group.

The Ethernet pseudowire is established using Targeted LDP (T-LDP) signaling and uses the `ether`, `vlan`, or `vpls` VC type on the SDP. The SDP can be an MPLS or a GRE type.

VPLS Service Considerations

This section describes general 7705 SAR service features and any special capabilities or considerations as they relate to VPLS services:

- SAP Encapsulations
- VLAN Processing
- QinQ (VPLS)

SAP Encapsulations

VPLS services are designed to carry Ethernet frame payloads; therefore, VPLS can provide connectivity between any SAPs and SDPs that pass Ethernet frames. The following SAP encapsulations are supported on the 7705 SAR VPLS service:

- Ethernet null
- Ethernet dot1q
- Ethernet qinq
- ATM VC with RFC 2684 llc-snap bridged encapsulation (see ATM PVC Access and Termination on a VPLS Service)

VLAN Processing

The SAP encapsulation definition on Ethernet ingress ports defines which VLAN tags are used to determine the service that the packet belongs to:

- **null** encapsulation defined at ingress—any VLAN tags are ignored and the packet goes to a default service for the SAP
- **dot1q** encapsulation defined at ingress—only the first label is considered
- **qinq** encapsulation defined at ingress—only the topmost two labels are considered

Note: The SAP can be defined with a wildcard (*) for the inner label (for example, “SAP 100:100.*”). In this example, 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 by the QinQ encapsulation of SAP 100:100.1, then traffic with 100:1 will go to that SAP while all other traffic with 100 as the first label will go to the SAP with the SAP 100:100.* definition.

For dot1q and qinq encapsulations, traffic encapsulated with tags for which there is no definition are discarded.

Tagging Rules for VPLS

VLAN tagging rules for VPLS SAPs are the same as those for Epipe SAPs except that VPLS includes the `force-c-vlan-forwarding` command.

The `force-c-vlan-forwarding` command provides users with the ability to preserve a customer VLAN tag and append a configured egress SAP VLAN ID on top of the customer tag. See the force-c-vlan-forwarding command for details.

For information on tagging rules, see Tagging Rules for Epipe.

QinQ (VPLS)

VPLS supports QinQ functionality. For details, see QinQ Support.

Configuration Notes

The following caveats apply to the implementation of VPLS:

- fabric mode must be set to aggregate mode (not per-destination mode)
- associating a service with a filter policy other than the default policy is optional

Reference Sources

For information on supported IETF drafts and standards, as well as standard and proprietary MIBs, refer to Standards and Protocol Support.

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
- Common Configuration Tasks
- Configuring VPLS Components
- Configuring a VPLS SAP
- Configuring SDP Bindings
- Configuring Routed VPLS
- Configuring IP Multicast in VPLS
- Configuring IP Multicast in r-VPLS
- Configuring Multicast Parameters for VPLS and r-VPLS
- Configuring a Static Multicast Group
- Service Management Tasks

Basic Configuration

The following fields require specific input (there are no defaults) to configure a basic VPLS service:

- customer ID (refer to Configuring Customer Accounts)
- for a local service, configure two SAPs, specifying local access ports and encapsulation values
- for a distributed service, configure a SAP and an SDP for each far-end node

The following example displays a configuration of a local VPLS service on ALU-1.

```
*A:ALU-1>config>service>vpls# info
----------------------------------------------
...
    vpls 9001 customer 6 create
        description "Local VPLS"
        sap 1/2/2:0 create
            description "SAP for local service"
        exit
        sap 1/1/5:0 create
            description "SAP for local service"
        exit
        no shutdown
----------------------------------------------
*A:ALU-1>config>service>vpls#
```

The following example displays a configuration of a distributed VPLS service between ALU-1, ALU-2, and ALU-3. The *vc-id* for all mesh SDPs must match the *service-id*.

```
*A:ALU-1>config>service# info
----------------------------------------------
...
    vpls 9000 customer 6 create
        shutdown
        description "This is a distributed VPLS."
        sap 1/1/5:16 create
            description "VPLS SAP"
        exit
        spoke-sdp 2:22 create
        exit
        mesh-sdp 7:9000 create
        exit
    exit
...
----------------------------------------------
*A:ALU-1>config>service#
```

```
*A:ALU-2>config>service# info
----------------------------------------------
...
    vpls 9000 customer 6 create
        description "This is a distributed VPLS."
        sap 1/1/5:16 create
            description "VPLS SAP"
        exit
        spoke-sdp 2:22 create
        exit
        mesh-sdp 8:9000 create
        exit
        no shutdown
    exit
...
----------------------------------------------
*A:ALU-2>config>service#

*A:ALU-3>config>service# info
----------------------------------------------
...
    vpls 9000 customer 6 create
        description "This is a distributed VPLS."
        sap 1/1/3:33 create
            description "VPLS SAP"
        exit
        spoke-sdp 2:22 create
        exit
        mesh-sdp 8:9000 create
        exit
        no shutdown
    exit
...
----------------------------------------------
*A:ALU-3>config>service#
```

Common Configuration Tasks

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

For VPLS services:

1. Associate a VPLS service with a customer ID. For management VPLS, include the `m-vpls` keyword when creating the VPLS.
2. Define SAPs:
   i. Select node(s) and port(s)
   ii. (optional) Select QoS policies other than the default (configured in `config>qos` context)
   iii. (optional) Select filter policies (configured in `config>filter` context)
   iv. (optional) Select accounting policy (configured in `config>log` context)
3. Associate SDPs (for distributed services).
4. (optional) Modify STP default parameters (see VPLS and Spanning Tree Protocol).
5. Enable the service.

Configuring VPLS Components

Topics in this section include:

- Creating a VPLS Service
- Creating a Split Horizon Group
- Enabling MAC Move
- Configuring STP Bridge Parameters in a VPLS

Creating a VPLS Service

Use the following CLI syntax to create a VPLS service:

CLI Syntax:

```
config>service#
vpls service-id [customer customer-id] [m-vpls] [create]
     description description-string
     no shutdown
```

The following example displays a VPLS configuration:

```
*A:ALU-1>config>service>vpls# info
----------------------------------------------
...
    vpls 9000 customer 6 create
        description "This is a distributed VPLS."
        shutdown
        exit
    exit
...
----------------------------------------------
*A:ALU-1>config>service>vpls#
```

Creating a Split Horizon Group

Use the following CLI syntax to create a split horizon group for a VPLS instance. Including the `residential-group` parameter creates a residential split horizon group.

CLI Syntax:
```
config>service>vpls#
     split-horizon-group group-name [residential-group]
        [create]
```

The following example displays a VPLS configuration:

```
*A:ALU-1>config>service>vpls# info
----------------------------------------------
...
    vpls 9000 customer 6 create
        description "VPLS with split horizon"
        split-horizon-group "SHG-group1" residential-group create
            description "Residential Split horizon group"
        exit
        no shutdown
    exit
...
----------------------------------------------
```

Enabling MAC Move

The MAC move feature is useful to protect against undetected loops in the 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 relearn rate for the MAC and will shut down the SAP or spoke SDP when the threshold is exceeded.

Use the following CLI syntax to configure MAC move parameters:

CLI Syntax:
```
config>service
vpls service-id [customer customer-id] [m-vpls] [create]
     mac-move
          primary-ports
               spoke-sdp spoke-id
               cumulative-factor cumulative-factor
          exit
          secondary-ports
               spoke-sdp spoke-id
               sap sap-id
          exit
          move-frequency frequency
          retry-timeout timeout
          no shutdown
```

The following example displays a MAC move configuration:

```
*A:ALU-2009>config>service>vpls>mac-move# show service id 500 mac-move
===============================================================================
Service Mac Move Information
===============================================================================
Service Id         : 500                      Mac Move         :Enabled
Primary Factor     : 4                        Secondary Factor : 2
Mac Move Rate      : 2                        Mac Move Timeout : 10
Mac Move Retries   : 3
-------------------------------------------------------------------------------
SAP Mac Move Information: 1/1/3:501
-------------------------------------------------------------------------------
Admin State        : Up                       Oper State       : Down
Flags              : RelearnLimitExceeded
Time to come up    : 1 seconds                Retries Left     : 1
Mac Move           : Blockable                Blockable Level  : Tertiary
-------------------------------------------------------------------------------
SAP Mac Move Information: 1/1/3:502
-------------------------------------------------------------------------------
Admin State        : Up                       Oper State       : Up
Flags              : None
Time to RetryReset : 267 seconds              Retries Left     : none
Mac Move           : Blockable                Blockable Level  : Tertiary
-------------------------------------------------------------------------------
SDP Mac Move Information: 21:501
-------------------------------------------------------------------------------
Admin State        : Up                       Oper State       : Up
Flags              : None
Time to RetryReset : never                    Retries Left    : 3
Mac Move           : Blockable                Blockable Level : Secondary
-------------------------------------------------------------------------------
SDP Mac Move Information: 21:502
-------------------------------------------------------------------------------
Admin State        : Up                       Oper State      : Down
Flags              : RelearnLimitExceeded
Time to come up    : never                    Retries Left    : none
Mac Move           : Blockable                Blockable Level : Tertiary
===============================================================================
*A:*A:ALU-2009>config>service>vpls>mac-move#
```

Configuring STP Bridge Parameters in a VPLS

Modifying some of the STP parameters allows the operator to balance STP between resiliency and speed of convergence extremes.

The following STP parameters can be modified at the VPLS level:

- Bridge STP Admin State
- Mode
- Bridge Priority
- Hello Time
- Hold Count

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

For the parameters 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 that Hello Time is always taken from the locally configured parameter.

Bridge STP Admin State

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

For SAPs, if STP on the VPLS is administratively disabled, any BPDUs are forwarded transparently through the 7705 SAR. If STP on the VPLS is administratively enabled, but the administrative state of a SAP is down, BPDUs received on such a SAP are discarded.

The 7705 SAR does not support BPDU extraction over spoke SDPs. If STP on the VPLS instance is disabled, BPDUs are forwarded transparently over the spoke SDP. If STP is enabled, the spoke SDP discards all BPDUs received.

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

Mode

The 7705 SAR operates in the Rapid Spanning Tree Protocol (RSTP) mode and is compliant with IEEE 802.1D-2004 - default mode.

CLI Syntax: `config>service>vpls` *`service-id`*`# stp`
`mode {rstp}`

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`

Hello Time

The `hello-time` command configures the STP hello time for the 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.

On the 7705 SAR, the hello time for the spanning tree is determined by the locally configured parameter.

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

Configuring a VPLS SAP

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

Topics in this section include:

- Local VPLS SAPs
- Distributed VPLS SAPs
- Configuring SAP-Specific STP Parameters
- STP SAP Operational States
- Configuring VPLS SAPs with Split Horizon

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.

All supported service types and corresponding uplink SAPs are specified in the following examples.

The following example displays a local VPLS configuration:

```
*A:ALU-1>config>service# info
----------------------------------------------
...
    vpls 9000 customer 6 create
        description "Local VPLS"
        sap 1/2/2:0 create
            description "SAP for local service"
        exit
        sap 1/1/5:0 create
            description "SAP for local service"
        exit
        no shutdown
    exit
----------------------------------------------
*A:ALU-1>config>service#
```

Distributed VPLS SAPs

To configure a distributed VPLS service, you must configure service entities on originating and far-end nodes. You must use the same service ID on all ends (for example, create a VPLS service ID 9000 on ALU-1, ALU-2, and ALU-3). A distributed VPLS consists of a SAP on each participating node and an SDP bound to each participating node.

For SDP configuration information, see Configuring an SDP. For SDP binding information, see Configuring SDP Bindings.

The following example displays a configuration of VPLS SAPs configured for ALU-1, ALU-2, and ALU-3:

```
*A:ALU-1>config>service# info
----------------------------------------------
...
    vpls 9000 customer 6 create
        description "Distributed VPLS services."
        shutdown
        exit
        sap 1/2/5:0 create
            description "VPLS SAP"
        exit
...
----------------------------------------------
*A:ALU-1>config>service#


*A:ALU-2>config>service# info
----------------------------------------------
...
    vpls 9000 customer 6 create
        description "Distributed VPLS services."
        shutdown
        exit
        sap 1/1/2:22 create
            description "VPLS SAP"
        exit
...
----------------------------------------------
*A:ALU-2>config>service#


*A:ALU-3>config>service# info
----------------------------------------------
...
    vpls 9000 customer 6 create
        description "Distributed VPLS services."
        shutdown
        exit
        sap 1/1/3:33 create
            description "VPLS SAP"
        exit
...
----------------------------------------------
*A:ALU-3>config>service#
```

Configuring SAP-Specific STP Parameters

When a VPLS has STP enabled, each SAP within the VPLS has STP enabled by default. The operation of STP on each SAP is governed by:

- SAP STP Administrative State
- SAP Virtual Port Number
- SAP Priority
- SAP Path Cost
- SAP Edge Port
- SAP Auto Edge
- SAP Link Type

SAP STP Administrative State

The administrative state of STP within a SAP controls how BPDUs are transmitted and handled when received. The allowable states are:

- **SAP Admin Up**

  The default administrative state is up for STP on a SAP. BPDUs are handled in the normal STP manner on a SAP that is administratively up.

- **SAP Admin Down**

  An administratively down state allows a service provider to prevent a SAP from becoming operationally blocked. BPDUs will not originate out the SAP towards the customer.

  If STP is enabled on the VPLS level, but disabled on the SAP, received BPDUs are discarded. Discarding the incoming BPDUs allows STP to continue to operate normally within the VPLS service while ignoring the down SAP. The specified SAP will always be in an operationally forwarding state.



Note: The administratively down state allows a loop to form within the VPLS.

CLI Syntax:

```
config>service>vpls>sap>stp#
[no] shutdown

Range: shutdown or no shutdown
Default: no shutdown (SAP admin up)
```

SAP Virtual Port Number

The virtual port number uniquely identifies a SAP within configuration BPDUs. The internal representation of a SAP is unique to a system and has a reference space much bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a SAP and identifies it with its own virtual port number, which is unique to every other SAP defined on the VPLS. The virtual port number is assigned at the time that the SAP is added to the VPLS.

Since the order in which SAPs are added to the VPLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance. To achieve consistency after a reboot, the virtual port number can be specified explicitly.

CLI Syntax:

```
config>service>vpls>sap# stp
port-num number

Range: 1 to 2047
Default: (automatically generated)
Restore Default: no port-num
```

SAP Priority

SAP priority allows a configurable “tie-breaking” parameter to be associated with a SAP. When configuration BPDUs are being received, the configured SAP priority will be used in some circumstances to determine whether a SAP will be designated or blocked.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16-bit value.

In the latest STP standard (802.1D-2004), only the upper 4 bits of the port priority field are used to encode the SAP priority. The remaining 4 bits are used to extend the port ID field into a 12-bit virtual port number field. The virtual port number uniquely references a SAP within the STP instance. See SAP Virtual Port Number for details on the virtual port number.

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

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

CLI Syntax: `config>service>vpls>sap>stp#`
`priority` *`stp-priority`*

Range: 0 to 255 (240 largest value, in increments of 16)
Default: 128
Restore Default: no priority

SAP Path Cost

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

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 7705 SAR 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 200000000, 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. 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 reinitialized 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 `auto-edge` 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 false (see SAP Edge Port).

Valid values for SAP `auto-edge` are enabled and disabled, with enabled being the default.

CLI Syntax:
```
config>service>vpls>sap>stp#
[no] auto-edge

Default: auto-edge
```

SAP Link Type

The SAP `link-type` parameter instructs STP on the maximum number of bridges behind this SAP.

If there is only a single bridge, transitioning to the forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected by a shared media, their SAPs should all be configured as shared, and timer-based transitions are used.

Valid values for SAP `link-type` are **shared** and **pt-pt**, with **pt-pt** being the default.

CLI Syntax:
```
config>service>vpls>sap>stp#
link-type {pt-pt|shared}

Default: link-type pt-pt
Restore Default: no link-type
```

STP SAP Operational States

The operational state of STP within a SAP controls how BPDUs are transmitted and handled when received. Defined states are:

- Operationally Disabled
- Operationally Discarding
- Operationally Learning
- Operationally Forwarding

Operationally Disabled

Operationally disabled is the normal operational state for STP on a SAP in a VPLS that has any of the following conditions:

- VPLS state administratively down
- SAP state administratively down
- SAP state operationally down

If the SAP enters the operationally up state with the STP administratively up and the SAP STP state is up, the SAP will transition to the STP SAP discarding state.

When, during normal operation, the router detects a downstream loop behind a SAP, BPDUs can be received at a very high rate. To recover from this situation, STP will transition the SAP to the disabled state for the forward-delay duration of 15 s.

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.



Note: In previous versions of the STP standard, the discarding state was called a blocked state.

Operationally Learning

The learning state allows for the 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.

Configuring VPLS SAPs with Split Horizon

To configure a VPLS service with a split horizon group, add the `split-horizon-group` parameter when creating the SAP. Traffic arriving on a SAP within a split horizon group will not be copied to other SAPs in the same split horizon group.

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

```
*A:ALU-1>config>service# info
----------------------------------------------
...
    vpls 800 customer 6001 create
        description "VPLS with split horizon for DSL"
        sap 1/1/3:1/100 split-horizon-group "DSL-group1" create
            description "SAP for residential bridging"
        exit
        sap 1/1/3:1/200 split-horizon-group "DSL-group1" create
            description "SAP for residential bridging"
        exit
        split-horizon-group "DSL-group1" residential-group create
            description "Split horizon group for DSL"
        exit
        no shutdown
    exit
...
----------------------------------------------
*A:ALU-1>config>service#
```

Configuring SDP Bindings

This section contains the following topics:

- Configuring Mesh SDP Bindings
- Configuring Spoke SDP
- Configuring VPLS Spoke SDPs with Split Horizon
- Configuring Selective MAC Flush

VPLS provides scaling and operational advantages. A hierarchical configuration eliminates the need for a full mesh of VCs between participating devices. Hierarchy is achieved by enhancing the base VPLS core mesh of VCs with access VCs (spoke) to form two tiers. Spoke SDPs are generally created between Layer 2 switches and placed at the Multi-Tenant Unit (MTU). The PE routers are placed at the service provider's Point of Presence (POP). Signaling and replication overhead on all devices is considerably reduced.

A spoke SDP is treated like the equivalent of a traditional bridge port, where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received on (unless a split horizon group was defined on the spoke SDP; see Configuring VPLS Spoke SDPs with Split Horizon).

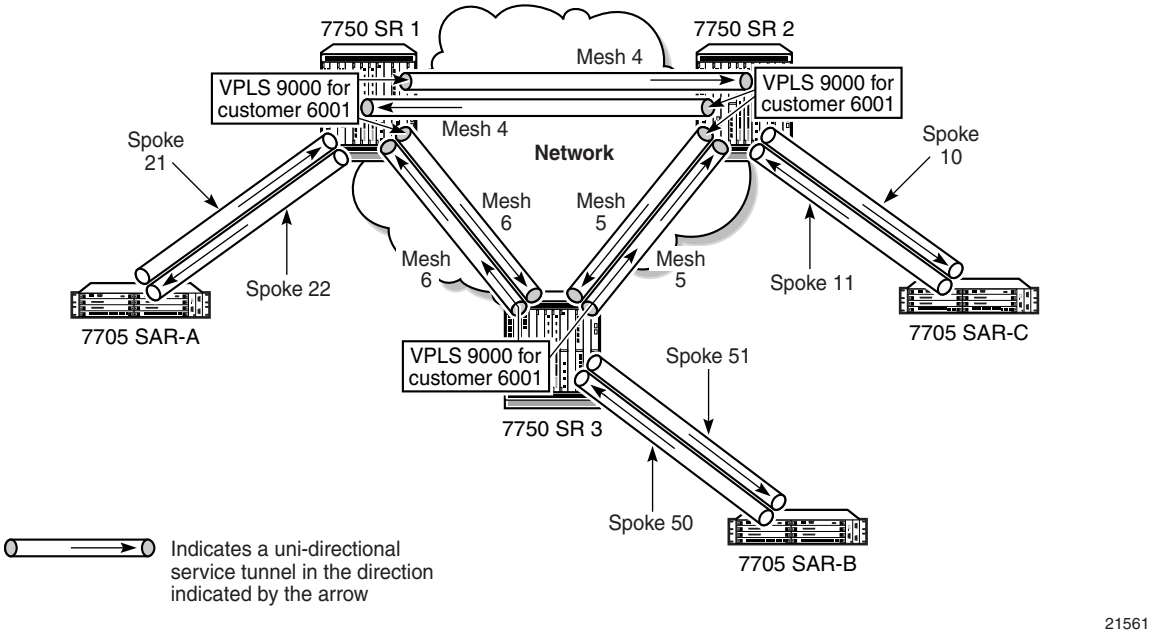
A spoke SDP connects a VPLS service between two sites and, in its simplest form, could be a single tunnel LSP. A set of ingress and egress VC labels are exchanged for each VPLS service instance to be transported over this LSP. The PE routers at each end treat this as a virtual spoke connection for the VPLS service in the same way as the PE-MTU connections. This architecture minimizes the signaling overhead and avoids a full mesh of VCs and LSPs between the two metro networks.

A mesh SDP bound to a service is logically treated like a single bridge “port” for flooded traffic, where flooded traffic received on any mesh SDP on the service is replicated to other “ports” (spoke SDPs and SAPs) and not transmitted on any mesh SDPs.

A VC-ID can be specified with the SDP-ID. The VC-ID is used instead of a label to identify a virtual circuit. The VC-ID is significant between peer 7705 SAR routers on the same hierarchical level. The value of a VC-ID is conceptually independent from the value of the label or any other datalink-specific information of the VC.

Figure 88 displays an example of a distributed VPLS service configuration of spoke and mesh SDPs (unidirectional tunnels) between 7750 SR routers and 7705 SAR MTUs.

Figure 88: SDPs—Unidirectional Tunnels



Configuring Mesh SDP Bindings

Use the following CLI syntax to create a mesh SDP binding with a distributed VPLS service. SDPs must be configured prior to binding. Refer to Configuring an SDP for information about creating SDPs.

Use the following CLI syntax to configure mesh SDP bindings:

CLI Syntax:

```
config>service# vpls service-id
    mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
        egress
            vc-label egress-vc-label
        ingress
            filter {ip ip-filter-id | mac mac-filter-
              id}
            vc-label ingress-vc-label
        no shutdown
        static-mac ieee-address
        vlan-vc-tag 0..4094
```

Configuring Spoke SDP

Topics in this section include:

- Configuring Spoke SDP Binding
- Configuring Spoke SDP-Specific STP Parameters

Configuring Spoke SDP Binding

Use the following CLI syntax to create a spoke SDP binding with a distributed VPLS service. SDPs must be configured prior to binding. Refer to Configuring an SDP for information about creating SDPs.

Use the following CLI syntax to configure spoke SDP bindings:

CLI Syntax:

```
config>service# vpls service-id
     spoke-sdp sdp-id:vc-id [vc-type {ether | vlan}]
        [split-horizon-group group-name]
          egress
               vc-label egress-vc-label
          ingress
               filter {ip ip-filter-id | mac mac-filter-
                  id}
               vc-label ingress-vc-label
          limit-mac-move [non-blockable]
          no shutdown
          static-mac ieee-address
          vlan-vc-tag [0..4094]
```

The following displays SDP binding configurations for ALU-1, ALU-2, and ALU-3 for VPLS service ID 9000 for customer 6:

```
*A:ALU-1>config>service# info
----------------------------------------------
    vpls 9000 customer 6 create
        description "This is a distributed VPLS."
        sap 1/2/5:0 create
        exit
        spoke-sdp 2:22 create
        exit
        mesh-sdp 5:9000 create
        exit
        mesh-sdp 7:9000 create
        exit
        no shutdown
    exit
----------------------------------------------
```

```
*A:ALU-1>config>service#

*A:ALU-2>config>service# info
----------------------------------------------
...
    vpls 9000 customer 6 create
        description "This is a distributed VPLS."
        sap 1/1/2:22 create
        exit
        spoke-sdp 2:22 create
        exit
        mesh-sdp 5:9000 create
        exit
        mesh-sdp 7:9000 create
        exit
        no shutdown
13.0 is     exit
----------------------------------------------

*A:ALU-3>config>service# info
----------------------------------------------
...
    vpls 9000 customer 6 create
        description "This is a distributed VPLS."
        sap 1/1/3:33 create
        exit
        spoke-sdp 2:22 create
        exit
        mesh-sdp 5:9000 create
        exit
        mesh-sdp 7:9000 create
        exit
        no shutdown
    exit
----------------------------------------------
*A:ALU-3>config>service#
```

Configuring Spoke SDP-Specific STP Parameters

When a VPLS has STP enabled, each spoke SDP within the VPLS has STP enabled by default. The operation of STP on each spoke SDP is governed by:

- Spoke SDP STP Administrative State
- Spoke SDP Virtual Port Number
- Spoke SDP Priority
- Spoke SDP Path Cost
- Spoke SDP Edge Port
- Spoke SDP Auto Edge
- Spoke SDP Link Type

Spoke SDP STP Administrative State

The administrative state of STP within a spoke SDP controls how BPDUs are transmitted and handled when received. The allowable states are:

- **spoke SDP admin up**

  The default administrative state is up for STP on a spoke SDP. BPDUs are handled in the normal STP manner on a spoke SDP that is administratively up.

- **spoke SDP admin down**

  An administratively down state allows a service provider to prevent a spoke SDP from becoming operationally blocked. BPDUs will not originate out the spoke SDP towards the customer.

  If STP is enabled on the VPLS level, but disabled on the spoke SDP, received BPDUs are discarded. Discarding the incoming BPDUs allows STP to continue to operate normally within the VPLS service while ignoring the down spoke SDP. The specified spoke SDP will always be in an operationally forwarding state.



Note: The administratively down state allows a loop to form within the VPLS.

CLI Syntax:
```
config>service>vpls>spoke-sdp>stp#
[no] shutdown
```

Spoke SDP Virtual Port Number

The virtual port number uniquely identifies a spoke SDP within configuration BPDUs. The internal representation of a spoke SDP is unique to a system and has a reference space much larger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a spoke SDP and identifies it with its own virtual port number, which is unique to any other spoke SDP defined on the VPLS. The virtual port number is assigned at the time that the spoke SDP is added to the VPLS.

Since the order in which spoke SDPs are added to the VPLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance. To achieve consistency after a reboot, the virtual port number can be specified explicitly.

CLI Syntax:
```
config>service>vpls>spoke-sdp# stp
port-num virtual-port-number
```

Spoke SDP Priority

Spoke SDP priority allows a configurable “tie-breaking” parameter to be associated with a spoke SDP. When configuration BPDUs are being received, the configured spoke SDP priority will be used in some circumstances to determine whether a spoke SDP will be designated or blocked.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (also 0 to 255) to create a 16-bit value.

In the latest STP standard (802.1D-2004), only the upper 4 bits of the port priority field are used to encode the spoke SDP priority. The remaining 4 bits are used to extend the port ID field into a 12-bit virtual port number field. The virtual port number uniquely references a spoke SDP within the STP instance. See Spoke SDP Virtual Port Number for details on the virtual port number.

STP computes the actual spoke SDP priority by taking the configured priority value and masking out the lower 4 bits. The result is the value that is stored in the spoke SDP 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 spoke SDP priority is 128. This parameter can be configured within a range of 0 to 255, with 0 being the highest priority. Masking causes the values actually stored and displayed to be 0 to 240, in increments of 16.

CLI Syntax: `config>service>vpls>spoke-sdp>stp#`
`priority` *`stp-priority`*

Spoke SDP Path Cost

The spoke SDP path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that spoke SDP. When BPDUs are sent out other egress spoke SDPs, the newly calculated root path cost is used.

STP suggests that the path cost is defined as a function of the link bandwidth. Since spoke SDPs are controlled by complex queuing dynamics, in the 7705 SAR the STP path cost is a purely static configuration.

The default value for spoke SDP path cost is 10. This parameter can be configured within a range of 1 to 200000000, with 1 being the lowest cost.

CLI Syntax: `config>service>vpls>spoke-sdp>stp#`
`path-cost` *`stp-path-cost`*

Spoke SDP Edge Port

The spoke SDP `edge-port` command is used to reduce the time it takes a spoke SDP to reach the forwarding state when the spoke SDP is on the edge of the network and thus has no further STP bridge to handshake with.

The `edge-port` command is used to initialize the internal OPER_EDGE variable. At any time, when OPER_EDGE is false on a spoke SDP, the normal mechanisms are used to transition to the forwarding state. When OPER_EDGE is true, STP assumes that the remote end agrees to transition to the forwarding state without actually receiving a BPDU with an agreement flag set.

The OPER_EDGE variable will dynamically be set to false if the spoke SDP receives BPDUs (the configured `edge-port` value does not change). The OPER_EDGE variable will dynamically be set to true if `auto-edge` is enabled and STP concludes there is no bridge behind the spoke SDP.

When STP on the spoke SDP is administratively disabled and re-enabled, the OPER_EDGE is reinitialized to the value configured for `edge-port`.

Valid values for spoke SDP edge port are enabled and disabled, with disabled being the default.

CLI Syntax:
```
config>service>vpls>spoke-sdp>stp#
[no] edge-port
```

Spoke SDP Auto Edge

The spoke SDP `auto-edge` command is used to instruct STP to dynamically decide whether the spoke SDP is connected to another bridge.

If `auto-edge` is enabled, and STP concludes there is no bridge behind the spoke SDP, the OPER_EDGE variable will dynamically be set to true. If `auto-edge` is enabled and a BPDU is received, the OPER_EDGE variable will dynamically be set to false (see Spoke SDP Edge Port).

Valid values for spoke SDP `auto-edge` are enabled and disabled, with enabled being the default.

CLI Syntax:
```
config>service>vpls>spoke-sdp>stp#
[no] auto-edge
```

Spoke SDP Link Type

The spoke SDP `link-type` parameter instructs STP on the maximum number of bridges behind this spoke SDP.

If there is only a single bridge, transitioning to the forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected by a shared media, their spoke SDPs should all be configured as shared, and timer-based transitions are used.

Valid values for spoke SDP `link-type` are `shared` and `pt-pt`, with `pt-pt` being the default.

CLI Syntax: `config>service>vpls>spoke-sdp>stp#`
`link-type {pt-pt|shared}`

Configuring VPLS Spoke SDPs with Split Horizon

To configure spoke SDPs with a split horizon group, add the `split-horizon-group` parameter when creating the spoke SDP. Traffic arriving on a SAP or spoke SDP within a split horizon group will not be copied to other SAPs or spoke SDPs in the same split horizon group.

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

```
*A:ALU-1>config>service# info
----------------------------------------------
...
    vpls 800 customer 6001 create
        description "VPLS with split horizon for DSL"
        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:ALU-1>config>service#
```

Configuring Selective MAC Flush

Use the following CLI syntax to enable selective MAC flush in a VPLS instance:

CLI Syntax: `config>service# vpls` *`service-id`*
`send-flush-on-failure`

Use the following CLI syntax to disable selective MAC flush in a VPLS instance:

CLI Syntax: `config>service# vpls` *`service-id`*
`no send-flush-on-failure`

Configuring Routed VPLS

To establish routed VPLS, a VPLS service must be bound to a standard IP interface within an IES or VPRN service. This is done by giving the VPLS a *service-name* and setting the VPLS `allow-ip-int-binding` flag. The binding is completed when the IES or VPRN interface is associated with the VPLS *service-name*. See Routed VPLS for details.

A VPLS service only supports binding for a single IP interface.

Additionally, an ingress IPv4 or IPv6 filter can be assigned to the VPLS SAP and the IES or VPRN interface. Use the `v4-` and `v6-routed-override-filter` commands to give the IP interface filter precedence over the VPLS SAP filter. See IES Command Reference and VPRN Services Command Reference for command descriptions.

Use the following CLI syntax to set up routed VPLS in a VPLS instance:

CLI Syntax:
```
config>service# vpls service-id
     allow-ip-int-binding
     service-name service-name
```

Use the following CLI syntax to bind an IES or VPRN interface to the routed VPLS instance and to configure an override filter:

CLI Syntax:
```
config>service>ies# interface ip-interface-name [create]
     vpls service-name
          ingress
               v4-routed-override-filter ipv4-filter-id
               v6-routed-override-filter ipv6-filter-id
```

CLI Syntax:
```
config>service>vprn# interface ip-interface-name
  [create]
     vpls service-name
          ingress
               v4-routed-override-filter ipv4-filter-id
               v6-routed-override-filter ipv6-filter-id
```

Configuring IP Multicast in VPLS

Use the `config>service>vpls>igmp-snooping (or mld-snooping)` command to enable IP multicast in VPLS. The `igmp-snooping` and `mld-snooping` commands stop the default flooding of multicast traffic and allow the creation of a multicast forwarding database (MFIB) on a per-port basis.

The following displays a VPLS configuration with IGMP snooping. Configuring MLD snooping is similar except that the `mld-snooping` command and IPv6 addresses are used instead of the `igmp-snooping` command and IPv4 addresses:

```
*A:ALU-1>config>service>vpls# info
----------------------------------------------
            description "Default tls description for service id 1"
            service-mtu 1400
            stp
                shutdown
            exit
            igmp-snooping
                no shutdown
            exit
            service-name "snooper"
            sap 1/1/5:12 create
                description "Default sap description for service id 1"
            exit
            sap 1/1/5:34 create
                description "Default sap description for service id 1"
            exit
            mesh-sdp 21:1 create
                no shutdown
            exit
            spoke-sdp 23:1 create
                no shutdown
            exit
            no shutdown
----------------------------------------------
*A:ALU-1>config>service>vpls#
```

Configuring IP Multicast in r-VPLS

Configuring IP multicast in a routed VPLS requires several steps.

Creating a Layer 2 multicast service in the context of an r-VPLS with PIM translation configured on the r-VPLS Layer 3 interface creates two multicast groups: one Layer 2 multicast group and one Layer 3 multicast group. Creating the Layer 2 multicast group automatically creates the Layer 3 group. It is not necessary to create both groups. The 7705 SAR uses one Layer 3 multicast group per source, and one Layer 2 multicast group per source per VPLS. See IP Multicast in r-VPLS for details.

Perform the following steps to create Layer 2 and Layer 3 multicast groups on a SAP or SDP.

1. Create an r-VPLS by using the `vpls>allow-ip-int-binding` command and the `ies>interface>vpls` *`service-name`* command.
2. Configure IGMP or MLD on IES to form the link between Layer 3 and Layer 2.
3. Configure PIM on a network interface to allow the propagation of multicast join messages into the network.
4. (Optional) Configure the Layer 2 multicast service parameters, as described in Configuring Multicast Parameters for VPLS and r-VPLS.

The following displays illustrate step 1 to step 3 for an r-VPLS configuration with IGMP snooping. Configuring MLD snooping is similar except that the `mld-snooping` command and IPv6 addresses are used instead of the `igmp-snooping` command and IPv4 addresses.

To create the r-VPLS:

```
*A:ALU-1>config>service>vpls# info
----------------------------------------------
            description "Default tls description for service id 1"
            service-mtu 1400
            allow-ip-int-binding
            stp
                shutdown
            exit
            igmp-snooping
                no shutdown
            exit
            service-name "snooper"
            sap 1/1/5:12 create
                description "Default sap description for service id 1"
            exit
            sap 1/1/5:34 create
                description "Default sap description for service id 1"
            exit
            mesh-sdp 21:1 create
                no shutdown
            exit
            spoke-sdp 23:1 create
                no shutdown
```

```
            exit
            no shutdown
-----------------------------------------------
*A:ALU-1>config>service>vpls#

*A:ALU-1>config>service>ies# info
-----------------------------------------------
            description "Default Ies description for service id 2"
            interface "rvpls_ies" create
                address 19.77.0.77/24
                ipv6
                    address fc77::77/64
                exit
                vpls "snooper"
                exit
            exit
            service-name "XYZ Ies 2"
            no shutdown
-----------------------------------------------
*A:ALU-1>config>service>ies#
```

To link Layer 3 and Layer 2:

```
*A:ALU-1>config>router>igmp# info
-----------------------------------------------
            interface "rvpls_ies"
                no shutdown
            exit
            no shutdown
-----------------------------------------------
*A:ALU-1>config>router>igmp#
```

To configure PIM on a network interface:

```
*A:ALU-1>config>router>pim# info
-----------------------------------------------
            no ipv6-multicast-disable
            interface "PimtoDut4"
            exit
            rp
            exit
            no shutdown
-----------------------------------------------
*A:ALU-1config>router>pim#
```

Configuring Multicast Parameters for VPLS and r-VPLS

The 7705 SAR supports multicast for VPLS and r-VPLS through IGMP and MLD snooping at the VPLS service level, as well as at the VPLS SAP and SDP (mesh and spoke) levels. Note the following considerations for IGMP and MLD snooping on a SAP or SDP.

- A filter policy can be imported on a SAP or SDP. Import policies are defined in the `config>router>policy-options` context. Refer to the "Filter Policies" section in the 7705 SAR OS Router Configuration Guide for details.
- A SAP or SDP can be configured as a multicast router port (`mrouter-port`), meaning that a multicast router is attached to this port. However, the `mrouter-port` and the `send-queries` commands are mutually exclusive commands.
- A static multicast group can be configured on a SAP or SDP (see Configuring a Static Multicast Group).
- The `send-queries` command must be enabled for the following commands to be operational: `query-interval`, `query-response-interval`, `robust-count`, and `version`.

Use the following CLI syntax to configure IGMP snooping parameters for VPLS and r-VPLS. Configuring MLD snooping parameters is similar except that the `mld-snooping` command and IPv6 addresses are used instead of the `igmp-snooping` command and IPv4 addresses.

CLI Syntax:

```
config>service>vpls# igmp-snooping
     query-interval seconds
     query-src-ip ip-address
     report-src-ip ip-address
     robust-count robust-count
```

The following displays IGMP snooping configuration for a VPLS service:

```
*A:ALU-1>config>service>vpls>igmp-snooping# info detail
----------------------------------------------
                no shutdown
                query-interval 125
                robust-count 2
                report-src-ip 0.0.0.0
                no query-src-ip
----------------------------------------------
*A:ALU-1>config>service>vpls>igmp-snooping#
```

Use the following CLI syntax to configure IGMP snooping on a SAP. Configuring IGMP snooping on an SDP is similar. Configuring MLD snooping on a SAP or SDP is also similar, except that the `mld-snooping` command and IPv6 addresses are used instead of the `igmp-snooping` command and IPv4 addresses, and the `max-num-grp-sources` and `max-num-sources` commands do not apply.

CLI Syntax:

```
config>service>vpls# sap sap-id
config>service>vpls>sap# igmp-snooping
config>service>vpls>sap>igmp-snooping#
     [no] disable-router-alert-check
     [no] fast-leave
     [no] import policy-name
     [no] last-member-query-interval interval
     [no] max-num-groups max-num-groups
     [no] max-num-grp-sources 1..512
     [no] max-num-sources max-num-sources
     [no] mrouter-port
     [no] query-interval seconds
     [no] query-response-interval seconds
     [no] robust-count robust-count
     [no] send-queries
     [no] static
          [no] group grp-ip-address
               [no] source ip-address
               [no] starg
     [no] version version
```

The following displays IGMP snooping configuration for a VPLS service:

```
*A:ALU-1>config>service>vpls>sap>igmp-snooping# info detail
----------------------------------------------
                    no fast-leave
                    no import
                    no max-num-groups
                    no max-num-sources
                    no max-num-grp-sources
                    last-member-query-interval 10
                    mrouter-port
                    query-interval 125
                    query-response-interval 10
                    robust-count 2
                    version 3
                    no send-queries
                    no disable-router-alert-check
----------------------------------------------
*A:ALU-1>config>service>vpls>sap>igmp-snooping#
```

Configuring a Static Multicast Group

A static multicast group is not created until the `source` or `starg`—(*,G)—is specified. More than one group can be created per SAP or SDP, and more than one source can be added to a group. A static source cannot be added to a group if a starg already exists in the group.

Use the following CLI syntax to configure a static group for IGMP snooping on a VPLS SAP. Configuring a static group for IGMP snooping on an SDP is similar. Configuring a static group for MLD snooping on a SAP or SDP is also similar, except that the `mld-snooping` command and IPv6 addresses are used instead of the `igmp-snooping` command and IPv4 addresses.

CLI Syntax:

```
config>service>vpls>sap# igmp-snooping
config>service>vpls>sap>igmp-snooping#
     [no] static
          [no] group grp-ip-address
               [no] source ip-address
               [no] starg
```

The following displays a static group configuration for IGMP snooping on a VPLS SAP (multiple groups and multiple sources):

```
*A:ALU-1>config>service>vpls>sap>igmp-snooping# info
----------------------------------------------
                    send-queries
                    static
                        group 225.0.0.2
                            starg
                        exit
                        group 225.0.0.3
                            source 11.12.12.12
                            source 11.12.13.13
                        exit
                    exit
----------------------------------------------
*A:ALU-1>config>service>vpls>sap>igmp-snooping#
```

Service Management Tasks

This section discusses the following service management tasks:

- Modifying VPLS Service Parameters
- Modifying Management VPLS Parameters
- Deleting a Management VPLS
- Disabling a Management VPLS
- Deleting a VPLS Service
- Disabling a VPLS Service
- Re-enabling a VPLS Service

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 parameters such as description, SAP, SDP, and/or service-MTU command syntax, and then enter the new information.

The following displays a modified VPLS configuration:

```
*A:ALU-1>config>service>vpls# info
----------------------------------------------
    description "This is a different description."
    disable-learning
    disable-aging
    discard-unknown
    local-age 500
    remote-age 1000
    stp
         shutdown
    exit
    sap 1/1/5:22 create
        description "VPLS SAP"
    exit
    spoke-sdp 2:22 create
    exit
    no shutdown
----------------------------------------------
*A:ALU-1>config>service>vpls#
```

Modifying Management VPLS Parameters

To modify the range of VLANs on an access port that are to be managed by an existing management VPLS, first the new range should be entered and then 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 shut down, and the service is shut down 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 unmanage 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 and SDPs are unbound (deleted), interfaces are shut down, and the service is shut down on the service level.

Use the following CLI syntax to delete a VPLS service:

CLI Syntax:
```
config>service
     [no] vpls service-id
          shutdown
          [no] mesh-sdp sdp-id
               shutdown
          sap sap-id [split-horizon-group group-name]
          no sap sap-id
               shutdown
```

Disabling a VPLS Service

Use the following CLI syntax to shut down a VPLS service without deleting the service parameters:

CLI Syntax:
```
config>service> vpls service-id
     [no] shutdown
```

Example:
```
config>service# vpls 1
config>service>vpls# shutdown
config>service>vpls# exit
```

Re-enabling a VPLS Service

To re-enable a VPLS service that was shut down:

CLI Syntax:
```
config>service> vpls service-id
     [no] shutdown
```

Example:
```
config>service# vpls 1
config>service>vpls# no shutdown
config>service>vpls# exit
```

VPLS Command Reference

Command Hierarchies

- VPLS Service Configuration Commands
  - → Global Commands
  - → SAP Commands
  - → Mesh SDP Commands
  - → Spoke SDP Commands
  - → Routed VPLS Commands
- Show Commands
- Clear Commands
- Debug Commands

VPLS Service Configuration Commands

Global Commands

```
config
    — service
        — vpls service-id [customer customer-id] [m-vpls] [create]
        — no vpls service-id
            — description description-string
            — no description
            — [no] disable-aging
            — [no] disable-learning
            — [no] discard-unknown
            — endpoint endpoint-name [create]
            — no endpoint
                — [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]
            — igmp-snooping
                — query-interval seconds
                — no query-interval
                — query-src-ip ip-address
                — no query-src-ip
                — report-src-ip ip-address
                — no report-src-ip
                — robust-count robust-count
                — no robust-count
                — [no] shutdown
            — local-age aging-timer
            — no local-age
            — [no] mac-move
                — move-frequency frequency
                — no move-frequency
                — number-retries number-retries
                — no number-retries
                — primary-ports
                    — cumulative-factor cumulative-factor
                    — no cumulative-factor
```

```
                    — [no] sap sap-id
                    — [no] spoke-sdp spoke-id
                — retry-timeout timeout
                — no retry-timeout
                — secondary-ports
                    — cumulative-factor cumulative-factor
                    — no cumulative-factor
                    — [no] sap sap-id
                    — [no] spoke-sdp spoke-id
                — [no] shutdown
            — mac-subnet-length subnet-length
            — no mac-subnet-length
            — mld-snooping
                — query-interval seconds
                — no query-interval
                — query-src-ip ipv6-address
                — no query-src-ip
                — report-src-ip ipv6-address
                — no report-src-ip
                — robust-count robust-count
                — no robust-count
                — [no] shutdown
            — [no] pppoe-circuit-id
            — [no] propagate-mac-flush
            — remote-age aging-timer
            — no remote-age
            — [no] send-flush-on-failure
            — service-mtu octets
            — no service-mtu
            — service-name service-name
            — no service-name
            — [no] shutdown
            — [no] split-horizon-group group-name [residential-group]
                — description description-string
                — no description
            — stp
                — hello-time hello-time
                — no hello-time
                — hold-count BDPU tx hold count
                — no hold-count
                — mode rstp
                — no mode
                — priority bridge-priority
                — no priority
                — [no] shutdown
```

SAP Commands

```
config
    — service
        — vpls service-id [customer customer-id] [create]
        — no vpls service-id
            — sap sap-id [split-horizon-group group-name] [create]
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy
                — atm
                    — egress
                        — traffic-desc traffic-desc-profile-id
                        — no traffic-desc
                    — encapsulation atm-encap-type
                    — oam
                        — [no] alarm-cells
                    — subscriber-vlan [vlan-id]
                    — no subscriber-vlan
                — [no] collect-stats
                — description description-string
                — no description
                — dhcp
                    — description description-string
                    — no description
                    — [no] option
                        — action [dhcp-action]
                        — no action
                        — circuit-id [ascii-tuple | vlan-ascii-tuple] | no circuit-id
                        — [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
                — [no] disable-aging
                — [no] disable-learning
                — [no] discard-unknown-source
                — egress
                    — agg-rate-limit agg-rate [cir cir-rate]
                    — no agg-rate-limit
                    — filter ip ip-filter-id
                    — filter ipv6 ipv6-filter-id
                    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
                    — [no] qinq-mark-top-only
                    — qos policy-id
                    — no qos
                    — scheduler-mode {4-priority | 16-priority}
                    — [no] shaper-group shaper-group-name
```

- [no] force-c-vlan-forwarding
- igmp-snooping
  - [no] disable-router-alert-check
  - [no] fast-leave
  - import *policy-name*
  - no import
  - last-member-query-interval *interval*
  - no last-member-query-interval
  - max-num-groups *max-num-groups*
  - no max-num-groups
  - max-num-grp-sources [*1..512*]
  - no max-num-grp-sources
  - max-num-sources *max-num-sources*
  - no max-num-sources
  - [no] mrouter-port
  - query-interval *seconds*
  - no query-interval
  - query-response-interval *seconds*
  - no query-response-interval
  - robust-count *robust-count*
  - no robust-count
  - [no] send-queries
  - static
    - [no] group *grp-ip-address*
      - [no] source *src-ip-address*
      - [no] starg
  - version *version*
  - no version
- ingress
  - agg-rate-limit *agg-rate* [**cir** *cir-rate*]
  - no agg-rate-limit
  - filter **ip** *ip-filter-id*
  - filter **ipv6** *ipv6-filter-id*
  - filter **mac** *mac-filter-id*
  - no filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*] [**mac** *mac-filter-id*]
  - match-qinq-dot1p {**top** | **bottom**}
  - no match-qinq-dot1p
  - qos *policy-id*
  - no qos
  - scheduler-mode {**4-priority** | **16-priority**}
  - [no] shaper-group *shaper-group-name*
- limit-mac-move [**blockable** | **non-blockable**]
- no limit-mac-move
- [no] mac-pinning
- managed-vlan-list
  - [no] range *vlan-range*
- max-nbr-mac-addr *table-size*
- no max-nbr-mac-addr
- mld-snooping
  - [no] disable-router-alert-check
  - [no] fast-leave
  - import *policy-name*
  - no import
  - last-member-query-interval *interval*

```
                        — no last-member-query-interval
                        — max-num-groups max-num-groups
                        — no max-num-groups
                        — [no] mrouter-port
                        — query-interval seconds
                        — no query-interval
                        — query-response-interval seconds
                        — no query-response-interval
                        — robust-count robust-count
                        — no robust-count
                        — [no] send-queries
                        — static
                            — [no] group grp-ipv6-address
                                — [no] source src-ipv6-address
                                — [no] starg
                        — version version
                        — no version
                    — [no] pppoe-circuit-id
                    — [no] shutdown
                    — [no] static-mac ieee-address [create]
                    — stp
                        — [no] auto-edge
                        — [no] edge-port
                        — link-type {pt-pt | shared}
                        — no link-type
                        — path-cost path-cost
                        — no path-cost
                        — [no] port-num virtual-port-number
                        — priority stp-priority
                        — no priority
                        — [no] root-guard
                        — [no] shutdown
```

Mesh SDP Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id] [create]
            — mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
            — no mesh-sdp sdp-id[:vc-id]
                — [no] control-word
                — egress
                    — vc-label egress-vc-label
                    — no vc-label [egress-vc-label]
                — igmp-snooping
                    — [no] disable-router-alert-check
                    — [no] fast-leave
                    — import policy-name
                    — no import
                    — last-member-query-interval interval
                    — no last-member-query-interval
```

```
            — max-num-groups max-num-groups
            — no max-num-groups
            — max-num-grp-sources [1..512]
            — no max-num-grp-sources
            — max-num-sources max-num-sources
            — no max-num-sources
            — [no] mrouter-port
            — query-interval seconds
            — no query-interval
            — query-response-interval seconds
            — no query-response-interval
            — robust-count robust-count
            — no robust-count
            — [no] send-queries
            — static
                — [no] group grp-ip-address
                    — [no] source src-ip-address
                    — [no] starg
            — version version
            — no version
        — ingress
            — filter ip ip-filter-id
            — filter ipv6 ipv6-filter-id
            — filter mac mac-filter-id
            — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
            — vc-label ingress-vc-label
            — no vc-label [ingress-vc-label]
        — [no] mac-pinning
        — mld-snooping
            — [no] disable-router-alert-check
            — [no] fast-leave
            — import policy-name
            — no import
            — last-member-query-interval interval
            — no last-member-query-interval
            — max-num-groups max-num-groups
            — no max-num-groups
            — [no] mrouter-port
            — query-interval seconds
            — no query-interval
            — query-response-interval seconds
            — no query-response-interval
            — robust-count robust-count
            — no robust-count
            — [no] send-queries
            — static
                — [no] group grp-ipv6-address
                    — [no] source src-ipv6-address
                    — [no] starg
            — version version
            — no version
        — [no] shutdown
        — [no] static-mac ieee-address
        — vlan-vc-tag 0..4094
```

```
— no vlan-vc-tag [0..4094]
```

Spoke SDP Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id] [create]
            — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [split-horizon-group group-name]
                [create] [no-endpoint]
            — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [split-horizon-group group-name]
                [create] endpoint endpoint-name
            — no spoke-sdp sdp-id[:vc-id]
                — [no] block-on-mesh-failure
                — [no] control-word
                — [no] disable-aging
                — [no] disable-learning
                — [no] discard-unknown-source
                — egress
                    — vc-label egress-vc-label
                    — no vc-label [egress-vc-label]
                — igmp-snooping
                    — [no] disable-router-alert-check
                    — [no] fast-leave
                    — import policy-name
                    — no import
                    — last-member-query-interval interval
                    — no last-member-query-interval
                    — max-num-groups max-num-groups
                    — no max-num-groups
                    — max-num-grp-sources [1..512]
                    — no max-num-grp-sources
                    — max-num-sources max-num-sources
                    — no max-num-sources
                    — [no] mrouter-port
                    — query-interval seconds
                    — no query-interval
                    — query-response-interval seconds
                    — no query-response-interval
                    — robust-count robust-count
                    — no robust-count
                    — [no] send-queries
                    — static
                        — [no] group grp-ip-address
                            — [no] source src-ip-address
                            — [no] starg
                    — version version
                    — no version
                — [no] ignore-standby-signaling
                — ingress
                    — filter ip ip-filter-id
                    — filter ipv6 ipv6-filter-id
```

```
            — filter mac mac-filter-id
            — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
            — vc-label ingress-vc-label
            — no vc-label [ingress-vc-label]
        — limit-mac-move [blockable | non-blockable]
        — no limit-mac-move
        — [no] mac-pinning
        — max-nbr-mac-addr table-size
        — no max-nbr-mac-addr
        — mld-snooping
            — [no] disable-router-alert-check
            — [no] fast-leave
            — import policy-name
            — no import
            — last-member-query-interval interval
            — no last-member-query-interval
            — max-num-groups max-num-groups
            — no max-num-groups
            — [no] mrouter-port
            — query-interval seconds
            — no query-interval
            — query-response-interval seconds
            — no query-response-interval
            — robust-count robust-count
            — no robust-count
            — [no] send-queries
            — static
                — [no] group grp-ipv6-address
                    — [no] source src-ipv6-address
                    — [no] starg
            — version version
            — no version
        — precedence [precedence-value | primary]
        — no precedence
        — [no] pw-status-signaling
        — [no] shutdown
        — [no] static-mac ieee-address
        — stp
            — [no] auto-edge
            — [no] edge-port
            — link-type {pt-pt | shared}
            — no link-type
            — path-cost path-cost
            — no path-cost
            — [no] port-num virtual-port-number
            — priority stp-priority
            — no priority
            — [no] root-guard
            — [no] shutdown
        — vlan-vc-tag 0..4094
        — no vlan-vc-tag [0..4094]
```

Routed VPLS Commands

```
config
    — service
        — vpls service-id
            — allow-ip-int-binding
            — no allow-ip-int-binding
            — service-name service-name
            — no service-name
```

Show Commands

```
show
    — service
        — egress-label start-label [end-label]
        — fdb-info
        — fdb-mac ieee-address [expiry]
        — id service-id
            — all
            — base
            — dhcp
                — statistics [sap sad-id] | [sdp sdp-id:vc-id]
                — statistics [interface interface-name | ip-address]
                — summary
            — endpoint
            — fdb [sap sap-id | sdp sdp-id | mac ieee-address | endpoint endpoint | detail ] [expiry]
            — igmp-snooping
                — all
                — base
                — port-db sap sap-id [detail]
                — port-db sap sap-id group grp-address
                — port-db sdp sdp-id:vc-id [detail]
                — port-db sdp sdp-id:vc-id group grp-address
                — proxy-db [detail]
                — proxy-db group grp-ip-address
                — querier
                — static [sap sap-id | sdp sdp-id:vc-id]
                — statistics [sap sap-id | sdp sdp-id:vc-id]
            — labels
            — mac-move
            — mld-snooping
                — all
                — base
                — port-db sap sap-id [detail]
                — port-db sap sap-id group grp-ipv6-address
                — port-db sdp-id:vc-id [detail]
                — port-db sdp sdp-id:vc-id group grp-ipv6-address
                — proxy-db [detail]
                — proxy-db group grp-ipv6-address
                — querier
```

```
                        — static [sap sap-id | sdp sdp-id:vc-id]
                        — statistics [sap sap-id | sdp sdp-id:vc-id]
                    — sap
                        — pppoe-circuit-id statistics
                    — sap sap-id [atm | base | detail | qos | sap-stats | stats]
                    — sdp
                    — sdp {sdp-id[:vc-id] | far-end ip-address} [detail]
                    — split-horizon-group [group-name]
                    — stp [detail]
                — ingress-label start-label [end-label]
                — sap-using [sap sap-id]
                — sap-using authentication-policy auth-plcy-name
                — sap-using interface [ip-address | ip-int-name]
                — sap-using description
                — sap-using [ingress | egress] atm-td-profile td-profile-id
                — sap-using [ingress | egress] filter filter-id
                — sap-using [ingress | egress] qos-policy qos-policy-id
                — sap-using [ingress | egress] scheduler-mode {4-priority | 16-priority}
```

Clear Commands

```
clear
    — service
        — id service-id
            — dhcp
                — statistics [sap sap-id | sdp sdp-id[:vc-id] | interface [ip-address | ip-int-name]]
            — fdb {all | mac ieee-address | sap sap-id | mesh-sdp sdp-id[:vc-id] | spoke-sdp
              sdp-id:vc-id}
            — igmp-snooping
                — port-db sap sap-id [group grp-address [source src-ip-address]]
                — port-db sdp sdp-id:vc-id [group grp-address [source src-ip-address]]
                — querier
                — statistics [all | sap sap-id | sdp sdp-id:vc-id]
            — mld-snooping
                — port-db sap sap-id [group grp-ipv6-address]
                — port-db sap sap-id group grp-ipv6-address source src-ipv6-address
                — port-db sdp sdp-id:vc-id [group grp-ipv6-address]
                — port-db sdp sdp-id:vc-id group grp-ipv6-address source src-ipv6-address
                — querier
                — statistics all
                — statistics sap sap-id
                — statistics sdp sdp-id:vc-id
            — mesh-sdp sdp-id[:vc-id] ingress-vc-label
            — sap
                — pppoe-circuit-id statistics
            — spoke-sdp sdp-id:vc-id ingress-vc-label
```

clear
- **service**
  - **statistics**
    - **id** *service-id*
      - **cem**
      - **counters**
      - **spoke-sdp** *sdp-id*:*vc-id* {**all** | **counters** | **l2pt** | **mrp**}
      - **stp**
    - **sap** *sap-id* {**all** | **cem** | **counters** | **l2pt** | **stp** | **mrp**}
    - **sdp** *sdp-id* **keep-alive**

clear
- **router**
  - **dhcp**
    - **statistics** [**interface** *ip-int-name* | *ip-address*]

Debug Commands

```
debug
    — service
        — id service-id
            — [no] igmp-snooping
                — detail-level {low | medium | high}
                — no detail-level
                — [no] mac ieee-address
                — mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}
                — no mode
                — [no] sap sap-id
                — [no] sdp sdp-id:vc-id
            — [no] mld-snooping
                — detail-level {low | medium | high}
                — no detail-level
                — [no] mac ieee-address
                — mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}
                — no mode
                — [no] sap sap-id
                — [no] sdp sdp-id:vc-id
            — [no] event-type {config-change | svc-oper-status-change | sap-oper-status-change |
              sdpbind-oper-status-change}
            — [no] sap sap-id
            — stp
                — all-events
                — [no] bpdu
                — [no] core-connectivity
                — [no] exception
                — [no] fsm-state-changes
                — [no] fsm-timers
                — [no] port-role
                — [no] port-state
                — [no] sap sap-id
```

Refer to the 7705 SAR OS OAM and Diagnostics Guide for information about CLI commands and syntax for OAM and diagnostics commands.

Command Descriptions

- VPLS Service Configuration Commands
- Routed VPLS Configuration Commands
- VPLS Show Commands
- VPLS Clear Commands
- VPLS Debug Commands

VPLS Service Configuration Commands

- Generic Commands
- VPLS Service Commands
- VPLS SAP Commands
- VPLS SAP ATM Commands
- VPLS STP Commands
- VPLS Filter and QoS Policy Commands
- Service Billing Commands
- VPLS SAP DHCP Commands
- VPLS SDP Commands

Generic Commands

shutdown

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

Context config>service>vpls
config>service>vpls>mac-move
config>service>vpls>split-horizon-group
config>service>vpls>sap
config>service>vpls>sap>dhcp
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>vpls>spoke-sdp>stp
config>service>vpls>stp
config>service>vpls>sap>stp
config>service>vpls>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping

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 can 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 are described below in Special Cases.

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

Special Cases **Service Admin State —** bindings to an SDP within the service will be put into the out-of-service state when the service is shut down. While the service is shut down, all customer packets are dropped and counted as discards for billing and debugging purposes.

Service Operational State — a service is regarded as operational providing that two SAPs or one SDP are operational

SDP (global) — when an SDP is shut down at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would normally be transmitted using this SDP binding will be discarded and counted as dropped packets.

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

SDP Keepalives — enables SDP connectivity monitoring keepalive messages for the SDP ID. The default state is disabled (shutdown), in which case the operational state of the SDP-ID is not affected by the keepalive message state.

VPLS SAPs and SDPs — SAPs are created in a VPLS and SDPs are bound to a VPLS in the administratively up default state. The created SAP will attempt to enter the operationally up state. An SDP will attempt to go into the in-service state once bound to the VPLS.

description

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

Context config>service>vpls
config>service>vpls>endpoint
config>service>vpls>sap
config>service>vpls>sap>dhcp
config>service>vpls>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 n/a

Parameters *description-string* — the description character string.

Values any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

VPLS Service Commands

vpls

Syntax **vpls** *service-id* [**customer** *customer-id*] [**m-vpls**] [**create**]
no vpls *service-id*

Context config>service

Description This command creates or edits a Virtual Private LAN Services (VPLS) 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.

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 and the **customer** keyword and *customer-id* must be specified in order to associate the service with a customer. The *customer-id* must already exist (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. To edit the customer association, the service must be deleted and recreated with a new customer association.

To create a management VPLS, include the **m-vpls** keyword when creating the VPLS. Associate a range of VLANs with the mVPLS by using the **managed-vlan-list** command.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

More than one VPLS service may be created for a single customer ID.

By default, no VPLS instances exist until they are explicitly created.

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs and SDPs defined within the service ID have been shut down and deleted, and the service has been shut down.

Parameters *service-id* — the unique service identification number or name that identifies 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 or name used for every 7705 SAR on which this service is defined.

Values 1 to 2147483647 or *service-name*

customer *customer-id* — specifies the customer ID number to be associated with the service. This parameter is required on service creation and is optional for service editing or deleting. The *customer-id* is not used with routed VPLS.

Values 1 to 2147483647

m-vpls — this keyword specifies that the VPLS is a management VPLS

create — this keyword is mandatory when creating a VPLS service

block-on-mesh-failure

Syntax [**no**] **block-on-mesh-failure**

Context config>service>vpls>spoke-sdp
config>service>vpls>endpoint

Description This command enables blocking (brings the entity to an operationally down state) after all configured SDPs or endpoints are in operationally down state. This event is signaled to a corresponding T-LDP peer by withdrawing the service label (status-bit-signaling non-capable peer) or by setting the "PW not forwarding" status bit in the T-LDP message (status-bit-signaling capable peer).

Default disabled

disable-aging

Syntax [**no**] **disable-aging**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command disables MAC address aging across a VPLS service or on a VPLS service SAP or spoke SDP.

As is the case for a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the VPLS forwarding database (FDB). The **disable-aging** command turns off aging for local and remote learned MAC addresses.

When **no disable-aging** is specified for a VPLS, it is possible to disable aging for specific SAPs and/or spoke SDPs by entering the **disable-aging** command at the appropriate level.

When the **disable-aging** command is entered at the VPLS level, the **disable-aging** state of individual SAPs or SDPs will be ignored.

The **no** form of this command enables aging on the VPLS service.

Default no disable-aging

disable-learning

Syntax [**no**] **disable-learning**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command disables learning of new MAC addresses in the VPLS forwarding database (FDB) for the service instance, SAP instance, or spoke SDP instance.

When **disable-learning** is enabled, new source MAC addresses will not be entered in the VPLS service forwarding database. This is true for both local and remote MAC addresses.

When **disable-learning** is disabled, new source MAC addresses will be learned and entered into the VPLS forwarding database.

This parameter is mainly used in conjunction with the **discard-unknown** command.

The **no** form of this command enables learning of MAC addresses.

Default no disable-learning (Normal MAC learning is enabled)

discard-unknown

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

Context config>service>vpls

Description By default, packets with unknown destination MAC addresses are flooded. If **discard-unknown** is enabled at the VPLS level, packets with an unknown destination MAC address will be dropped instead of being flooded (even when configured FDB size limits for VPLS or SAPs are not yet reached).

The **no** form of this command allows flooding of packets with unknown destination MAC addresses in the VPLS.

Default no discard-unknown—packets with unknown destination MAC addresses are flooded

endpoint

Syntax **endpoint** *endpoint-name* [**create**]
no endpoint

Context config>service>vpls

Description This command configures a service endpoint.

Parameters *endpoint-name* — specifies an endpoint name

Values 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 (“ ”).

create — this keyword is mandatory when creating a service endpoint

ignore-standby-signaling

Syntax [**no**] **ignore-standby-signaling**

Context config>service>vpls>endpoint
config>service>vpls>spoke-sdp

Description When this command is enabled, the node will ignores the standby bit received from T-LDP peers for the given spoke SDP and performs internal tasks without taking the standby bit into account—traffic can egress out to the spoke SDP.

This command is present at the endpoint level as well as at the 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 that is a part of the given explicit endpoint will inherit this setting from the endpoint configuration.

Default enabled

revert-time

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

Context config>service>vpls>endpoint

Description This command configures the time to wait before reverting to the primary spoke SDP.

For a regular endpoint, the *revert-time* setting affects only the pseudowire defined as “primary” (precedence 0). If the primary pseudowire fails and is then restored, the revert timer is started. After the revert timer expires, the primary pseudowire takes the active role in the endpoint. This behavior does not apply if both pseudowires are defined as “secondary”. For example, if the active secondary pseudowire fails and is restored, it will stay in standby until a configuration change or a force command occurs.

Parameters *revert-time* — specifies the time to wait, in seconds, before reverting back to the primary spoke SDP defined on this service endpoint, after having failed to move over to a backup spoke SDP

Values 0 to 600

infinite — specifying this keyword makes the endpoint non-revertive

static-mac

Syntax **static-mac** *ieee-address* [**create**]
no static-mac

Context config>service>vpls>endpoint

Description This command assigns a static MAC address to the endpoint. In the FDB, the static MAC address is then associated with the active spoke SDP.

Default n/a

Parameters *ieee-address* — specifies the static MAC address assigned to the endpoint

Values 6-byte MAC address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx) - cannot be all zeros

create — this keyword is mandatory when creating a static MAC address

suppress-standby-signaling

Syntax [**no**] **suppress-standby-signaling**

Context config>service>vpls>endpoint

Description When this command is enabled, the pseudowire standby bit (value 0x00000020) will not be sent to the T-LDP peer when the given spoke SDP is selected as a standby. This allows faster switchover because the traffic will be sent over this SDP and discarded at the blocking side of the connection. This is particularly applicable to multicast traffic.

Default enabled

pppoe-circuit-id

Syntax [**no**] **pppoe-circuit-id**

Context config>service>vpls
config>service>vpls>sap

Description When enabled, this command appends Agent-Circuit-Id information to PADI and PADR packets received from an ATM SAP (the subscriber) that is bound to a VPLS instance. The Agent-Circuit-Id information is compliant with RFC 4679 section-3.3.1; Agent-Circuit-Id. The ATM SAP must be configured for bridged llc-snap encapsulation.

The **pppoe-circuit-id** command can be enabled or disabled for a VPLS instance or an individual ATM SAP. When applied to a VPLS instance, **pppoe-circuit-id** appends the Agent-Circuit-Id to all ATM SAPs bound to that VPLS instance. Furthermore, **pppoe-circuit-id** can be applied to individual SAPs bound to that VPLS instance in order to override the VPLS setting. If there is a mix of enabled and disabled SAPs bound to the VPLS instance, applying the command to the VPLS will override the mix, enabling (or disabling) **pppoe-circuit-id** on all the SAPs.

In addition, any newly created SAPs bound to the VPLS will default to match the VPLS setting.

As per the DSL Forum TR-101 April'06 specification, section 3.9.2, any PPPoE vendor-specific tag that may already be present in the received frame is replaced by the 7705 SAR client-id tag.

The **no** version of this command disables appending the Agent-Circuit-Id information.

Default disabled

propagate-mac-flush

Syntax [**no**] **propagate-mac-flush**

Context config>service>vpls

Description This command enables the propagation of mac-flush messages received from the given T-LDP on all spoke and mesh SDPs within the context of the VPLS service. The propagation conforms to split-horizon principles and any datapath blocking in order to avoid looping of these messages.

Default disabled

fdb-table-high-wmark

Syntax [**no**] **fdb-table-high-wmark** *high-water-mark*

Context config>service>vpls

Description This command specifies the upper threshold value for FDB entries. The *high-water-mark* is configured as a percentage of the FDB. When the number of FDB entries exceeds the *high-water-mark*, the system raises a log event.

Parameters *high-water-mark* — specifies the upper threshold for FDB entries, which when exceeded, causes the system to raise a log event

Values 1 to 100

Default 95%

fdb-table-low-wmark

Syntax [**no**] **fdb-table-low-wmark** *low-water-mark*

Context config>service>vpls

Description This command specifies the lower threshold value for FDB entries. The *low-water-mark* is configured as a percentage of the FDB. When the number of FDB entries drops below the *low-water-mark*, the system raises a log event.

Parameters *low-water-mark* — specifies the lower threshold for FDB entries, which when dropped below, causes the system to raise a log event

Values 1 to 100

Default 90%

fdb-table-size

Syntax **fdb-table-size** *table-size*
no fdb-table-size [*table-size*]

Context config>service>vpls

Description This command specifies the maximum number of MAC entries in the FDB for the VPLS instance on this node.

The **fdb-table-size** specifies the maximum number of FDB entries for both learned and static MAC addresses for the VPLS instance.

The **no** form of this command returns the maximum FDB table size to the default.

Default 250

Parameters *table-size* — the maximum number of MAC entries in the FDB for the VPLS instance on the node

Values 1 to 16383

local-age

Syntax **local-age** *aging-timer*
no local-age

Context config>service>vpls

Description This command specifies the aging time for locally learned MAC addresses in the FDB for the VPLS instance. In a VPLS service, MAC addresses are associated with a SAP or SDP. MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

As is the case for 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 locally learned MAC and remotely learned MAC entries in the FDB. The **local-age** timer specifies the aging time for locally learned MAC addresses.

The **no** form of this command returns the local aging timer to the default value.

Default 300

Parameters *aging-timer* — the aging time for local MACs expressed in seconds

Values 60 to 86400

mac-move

Syntax [**no**] **mac-move**

Context config>service>vpls

Description This command enables the context to configure MAC move attributes. A sustained, high relearn rate can be a sign of a loop somewhere in the VPLS topology. Typically, the spanning tree protocol (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 the network against loops.

When enabled in a VPLS, **mac-move** monitors the relearn 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 increases linearly with the number of times the given SAP was disabled. A SAP can be configures as non-blockable with the limit-mac-move command. This means that when the relearn rate has exceeded the limit, another (blockable) SAP will be disabled instead. By default, all SAPs and spoke SDPs are configured as blockable when MAC-move is enabled.

When MAC-move is enabled and the relearn rate exceeds the maximum limit, the 7705 SAR sends a "Mac move rate for MAC ... exceeded" alarm. This alarm is raised for both blockable and non-blockable SAPs and spoke SDPs. The alarm frequency for non-blockable SAPs and spoke SDPs decreases if the MAC-move condition persists.

The **mac-move** command enables the feature at the service level for SAPs and spoke SDPs, as only those objects can be blocked by this feature. Mesh SDPs are never blocked, but their relearn rates (SAP-to-mesh/spoke-to-mesh or vice versa) are still measured.

The operation of this feature is the same on the SAP and spoke SDP. For example, if a MAC address moves from SAP to SAP, from SAP to spoke SDP, or between spoke SDPs, one will be blocked to prevent thrashing. If the MAC address moves between a SAP and mesh SDP or spoke SDP and mesh SDP combinations, the respective SAP or spoke SDP will be blocked.

The **mac-move** command 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-s interval. For example, when the move frequency is configured to 1 (1 relearn per second), **mac-move** will disable one of the VPLS ports when 5 relearns were detected during the 5-s interval because the average move frequency of 1 relearn per second has been reached. This can also occur in the first second if the relearn rate is 5 relearns per second or higher.

The **no** form of this command disables MAC move.

mac-subnet-length

Syntax **mac-subnet-length** *subnet-length*
no mac-subnet-length

Context config>service>vpls

Description This command specifies the number of bits to be considered when performing MAC learning (MAC source) and MAC switching (MAC destination). Specifically, this value identifies how many bits are used, starting from the beginning of the MAC address. For example, if the mask value of 28 is used, MAC learning will only do a lookup for the first 28 bits of the source MAC address when comparing it with existing FDB entries. Then, it will install the first 28 bits in the FDB while zeroing out the last 20 bits of the MAC address. When performing switching in the reverse direction, only the first 28 bits of the destination MAC address will be used to perform an FDB lookup to determine the next hop.

The **no** form of this command switches back to full MAC lookup.

Parameters *subnet-length* — specifies the number of bits to be considered when performing MAC learning or MAC switching

Values 24 to 48

move-frequency

Syntax **move-frequency** *frequency*
no move-frequency

Context config>service>vpls>mac-move

Description This command indicates the maximum rate at which MACs can be relearned in the VPLS service before the SAP where the moving MAC was last seen is automatically disabled in order to protect the system against undetected loops or duplicate MACs.

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

Default 2 (when **mac-move** is enabled); for example, 10 relearns in a 5-s period.

Parameters *frequency* — specifies the rate, in 5-s intervals, for the maximum number of relearns

Values 1 to 10

number-retries

Syntax **number-retries** *number-retries*
no number-retries

Context config>service>vpls>mac-move

Description This command configures the number of times that retries are performed for re-enabling the SAP or SDP bindings.

Parameters *number-retries* — specifies the number of retries for re-enabling the SAP/SDP. A zero (0) value indicates an unlimited number of retries.

Values 0 to 255

primary-ports

Syntax **primary-ports**

Context config>service>vpls>mac-move

Description This command enables the context to define primary VPLS ports. VPLS ports that were declared as secondary prior to the execution of this command will be moved from secondary port level to primary port level. Changing a port to the tertiary level (default) can only be done by first removing it from the secondary port level.

cumulative-factor

Syntax **cumulative-factor** *cumulative-factor*
no cumulative-factor

Context config>service>vpls>mac-move>primary-ports
config>service>vpls>mac-move>secondary-ports

Description This command configures a factor for the primary or secondary ports that defines how many MAC relearn periods should be used to measure the MAC relearn rate. The rate must be exceeded during consecutive periods before the corresponding ports (SAP and/or spoke SDP) are blocked by the **mac-move** feature.

Parameters *cumulative-factor* — specifies a MAC relearn period to be used for the MAC relearn rate

Values 3 to 10

sap

Syntax [**no**] **sap** *sap-id*

Context config>service>vpls>mac-move>primary-ports
config>service>vpls>mac-move>secondary-ports

Description This command configures the specified SAP to be a primary or secondary VPLS port.

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

spoke-sdp

Syntax [**no**] **spoke-sdp** *spoke-id*

Context config>service>vpls>mac-move>primary-ports
config>service>vpls>mac-move>secondary-ports

Description This command configures the specified spoke SDP to be a primary or secondary VPLS port.

Parameters *spoke-id* — specifies the spoke SDP to be used as the primary or secondary VPLS port

Values

spoke-id:	*sdp-id*:*vc-id*	
sdp-id:	1 to 17407	
vc-id:	1 to 4294967295	

secondary-ports

Syntax **secondary-ports**

Context config>service>vpls>mac-move

Description This command enables the context to define secondary VPLS ports. VPLS ports that were declared as primary prior to the execution of this command will be moved from the primary port level to the secondary port level. Changing a port to the tertiary level (default) can only be done by first removing it from the primary port level.

retry-timeout

Syntax **retry-timeout** *timeout*
no retry-timeout

Context config>service>vpls>mac-move

Description This command indicates the time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled.

It is recommended that the *timeout* value be equal to or larger than 5 s × cumulative factor of the highest-priority port so that the sequential order of port blocking will not be disturbed by reinitializing lower-priority ports.

A zero value indicates that the SAP will not automatically be 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.

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 re-enabled

Values 0 to 600

remote-age

Syntax **remote-age** *aging-timer*
no remote-age

Context config>service>vpls

Description This command specifies the aging time for remotely learned MAC addresses in the FDB for the VPLS instance. In a VPLS service, MAC addresses are associated with a SAP or an SDP. MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

As is the case for 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 locally learned MAC and remotely learned MAC entries in the FDB. The **remote-age** timer specifies the aging time for remotely learned MAC addresses. To reduce the amount of signaling required between switches, configure this timer to be larger than the **local-age** timer.

The **no** form of this command returns the remote aging timer to the default value.

Default 900

Parameters *seconds* — the aging time for remote MACs expressed in seconds

Values 60 to 86400

send-flush-on-failure

Syntax [**no**] **send-flush-on-failure**

Context config>service>vpls

Description This command enables sending out "flush-all-from-ME" messages to all LDP peers included in the 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 Rapid Spanning Tree Protocol (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.

Default no send-flush-on-failure

service-mtu

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

Context config>service>vpls

Description This command configures the service payload Maximum Transmission Unit (MTU), in bytes, for the service. This MTU value overrides the service-type default MTU. The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP bindings' operational state within the service.

The service MTU and a SAP's service delineation encapsulation overhead (that is, 4 bytes for a dot1q tag or 8 bytes for a qinq tag) are 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, the SAP will be placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP will be able to transition to the operative state.

When binding an SDP to a service, the service MTU is compared to the path MTU associated with the SDP. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced due to the MTU capabilities discovered by the tunneling mechanism of the SDP or the egress interface MTU capabilities based on the next hop in the tunnel path. If the service MTU is larger than the path MTU, the SDP binding for the service will be placed in an inoperative state. If the service MTU is equal to or less than the path MTU, then the SDP binding will be placed in an operational state.

If 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 VPLS: 1514

Table 59 shows MTU values for specific VC types.

Table 59: Show Service Service-MTU Output Fields

VC-Type Example	Service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
Ethernet (with preserved qinq)	1522	1508
VPLS	1514	1500
VPLS (with preserved dot1q)	1518	1504
VPLS (with preserved qinq)	1522	1508
VLAN (dot1p transparent to MTU value)	1514	1500

Parameters *octets* — the size of the MTU, in octets, expressed as a decimal integer

Values 1 to 9670

service-name

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

Context config>service>vpls

Description This command configures a service name that can be used in other configuration commands and show commands that reference the service.

Parameters *service-name* — up to 64 characters

split-horizon-group

Syntax [**no**] **split-horizon-group** *group-name* [**residential-group**]

Context config>service>vpls

Description This command creates a new split horizon group (SHG) 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. If the **residential-group** keyword is included, the split horizon group is a residential SHG.

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 reused in different VPLS instances.

An ATM SAP must be in a residential SHGs. If an Ethernet SAP is in a SHG, then that SHG cannot be a residential SHG.

Up to 30 split horizon groups can be defined per VPLS instance.

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

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

residential-group — defines a split horizon group as a residential split horizon group (RSHG). Doing so ensures that:

- SAPs that are members of this RSHG have:
  - → MAC pinning enabled per default (can be disabled)
  - → broadcast and multicast packets are discarded at the SAP egress point, thus blocking unknown flooded traffic
  - → downstream multicast packets are allowed when IGMP snooping is enabled
- Spoke SDPs that are members of this RSHG have:
  - → broadcast and multicast packets are NOT discarded at the spoke SDP egress point, thus allowing the unknown flooded traffic
  - → MAC pinning enabled per default (can be disabled)

Default By default, a split horizon group is not created as a residential group

VPLS SAP Commands

sap

Syntax **sap** *sap-id* [**split-horizon-group** *group-name*] [**create**]
no sap *sap-id*

Context config>service>vpls

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters that identify the service access point on the interface and within the 7705 SAR. 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. Channelized TDM ports are always access ports.

If a port is shut down, all SAPs on that port become operationally down. When a service is shut down, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP will also be deleted. For Internet Enhanced Service (IES), the IP interface must be shut down before the SAP on that interface may be removed.

Default n/a

Special Cases **VPLS SAP —** a VPLS SAP can be defined on SONET/SDH channels, on Ethernet ports, or on a GPON module port and on DSL module ports when the module is installed in the 7705 SAR-M (variants with module slots). Split horizon groups (SHGs) and residential SHGs can only be created in the scope of a VPLS service.

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

create — keyword used to create a SAP instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context

group-name — specifies the name of the split horizon group to which the SAP belongs

discard-unknown-source

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

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command specifies that packets received on a SAP or a spoke SDP with an unknown source MAC address will be dropped only if the maximum number of MAC addresses for that SAP or spoke SDP (see max-nbr-mac-addr) has been reached. If **max-nbr-mac-addr** has not been set for the SAP or spoke SDP, enabling **discard-unknown-source** has no effect.

When disabled, the packets are forwarded based on the destination MAC addresses.

The **no** form of this command causes packets with an unknown source MAC addresses to be forwarded by destination MAC addresses in VPLS.

Default no discard-unknown-source

limit-mac-move

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

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command indicates whether or not the MAC move agent, when enabled using mac-move, will limit the MAC relearn (move) rate on this SAP.

Default blockable

Parameters **blockable** — when the relearn rate is exceeded, the SAP or spoke SDP is blocked and the 7705 SAR raises a a "Mac move rate for MAC ... exceeded" alarm.

non-blockable — when the relearn rate is exceeded, this SAP or spoke SDP is not blocked, another blockable SAP or spoke SDP is blocked instead, and the 7705 SAR raises a a "Mac move rate for MAC ... exceeded" alarm. The alarm frequency decreases if the MAC-move condition persists.

mac-pinning

Syntax [**no**] **mac-pinning**

Context config>service>vpls>endpoint
config>service>vpls>sap
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command disables relearning of MAC addresses on other SAPs or SDPs within the VPLS. The MAC address will remain attached to a given SAP or SDP for the duration of its age timer.

The age of the MAC address entry in the FDB is set by the age timer. If **mac-aging** is disabled on a given VPLS service, any MAC address learned on a SAP or SDP with **mac-pinning** enabled will remain in the FDB on this SAP or SDP forever. Every event that would otherwise result in relearning will be logged (MAC address; original SAP; new SAP).

MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC pinning for such addresses is implicit.

Default disabled, except enabled at the creation of the SAP or spoke SDP that is part of a residential split horizon group (RSHG)

max-nbr-mac-addr

Syntax **max-nbr-mac-addr** *table-size*
no max-nbr-mac-addr

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>endpoint

Description This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this SAP, spoke SDP, or endpoint.

When the configured limit has been reached, and **discard-unknown-source** has been enabled for this SAP or spoke SDP (see discard-unknown-source), packets with unknown source MAC addresses will be discarded.

The **no** form of the command restores the global MAC learning limitations for the SAP or spoke SDP.

Default no max-nbr-mac-addr

Parameters *table-size* — specifies the maximum number of learned and static entries allowed in the FDB of this service

Values 1 to 16383

static-mac

Syntax [**no**] **static-mac** *ieee-address* [**create**]

Context config>service>vpls>sap

Description This command creates a local static MAC entry in the VPLS FDB associated with the SAP.

In a VPLS service, MAC addresses are associated with a SAP or an 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 FDB for the VPLS instance so that the MAC address will not be learned on the edge device.

Static MAC definitions on one edge device are not propagated to other edge devices participating in the VPLS instance; that is, each edge device has an independent FDB 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 FDB.

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 (cannot be all zeros). 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

VPLS SAP ATM Commands

atm

Syntax **atm**

Context config>service>vpls>sap

Description This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:

- configuring an ATM port or ATM port-related functionality on adapter cards supporting ATM functionality
- configuring ATM-related configuration for ATM-based SAPs that exist on adapter cards supporting ATM functionality

If ATM functionality is not supported for a given context, the command returns an error.

egress

Syntax **egress**

Context config>service>vpls>sap>atm

Description This command enables the context to configure egress ATM attributes for the SAP.

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
no traffic-desc

Context config>service>vpls>sap>atm>egress

Description This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP).

When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backwards direction.

The **no** form of the command reverts the traffic descriptor to the default traffic descriptor profile.

Default The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs

Parameters *traffic-desc-profile-id* — specifies a defined traffic descriptor profile (refer to "ATM QoS Traffic Descriptor Profiles" in the 7705 SAR OS Quality of Service Guide for information on the **atm-td-profile** t*raffic-desc-profile-id* command)

Values 1 to 1000

encapsulation

Syntax **encapsulation** *atm-encap-type*

Context config>service>vpls>sap>atm

Description This command specifies the data encapsulation for an ATM PVCC-delimited SAP. The definition references RFC 2684, *Multiprotocol Encapsulation over ATM AAL5,* and the ATM Forum LAN Emulation specification.

Ingress traffic that does not match the configured encapsulation will be dropped.

Default The encapsulation is driven by the service for which the SAP is configured.

For VPLS SAPs, the default and only option is **aal5snap-bridged**.

Parameters *atm-encap-type* — specifies the encapsulation type

Values **aal5snap-bridged:** bridged encapsulation for LLC encapsulated circuit (LLC/ SNAP precedes protocol datagram) as defined in RFC 2684

oam

Syntax **oam**

Context config>service>vpls>sap>atm

Description This command enables the context to configure OAM functionality for a PVCC delimiting a SAP. The ATM-capable adapter cards support the following F5 end-to-end OAM functionality (AIS, RDI, loopback):

- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
- GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3 June 1996
- GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

Syntax [**no**] **alarm-cells**

Context config>service>vpls>sap>atm>oam

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of its connection by propagating fault information through the network and by driving a PVCC operational status.

When **alarm-cells** functionality is enabled, a PVCC operational status is affected when a PVCC goes into an AIS or RDI state because of AIS/RDI processing. This assumes that nothing else affects the PVCC operational status; for example, the PVCC goes down, or enters a fault state and comes back up, or exits that fault state. RDI cells are generated when a PVCC is operationally down. No OAM-specific SNMP trap is raised whenever an endpoint enters or exits an AIS or RDI state; however, if an OAM state change results in a change to the operational status of the PVCC, then a trap is expected from an entity that the PVCC is associated with (for example, a SAP).

The **no** form of the command disables **alarm-cells** functionality for a PVCC. When **alarm-cells** functionality is disabled, the PVCC operational status is no longer affected by the PVCC OAM state changes due to AIS/RDI processing. When **alarm-cells** is disabled, a PVCC will change operational status to up from down due to **alarm-cell** processing. RDI cells are not generated as result of PVCC going into an AIS or RDI state; however, the PVCC OAM status will record OAM faults as described above.

Default enabled for PVCCs delimiting VPLS SAPs

subscriber-vlan

Syntax **subscriber-vlan** [*vlan-id*]
no subscriber-vlan

Context config>service>vpls>sap>atm

Description This command enables the push operation of a configured VLAN at ingress and a pop operation at egress on a per-ATM SAP basis. After AAL5 termination at ATM access ingress as per the configured encapsulation type, the configured VLAN tag is pushed to the received subscriber frame. The type of Ethernet frame is set to 0x8100 in order to designate the existence of the VLAN header, and the original Ethertype is shifted by 4 bytes (dot1q) or 8 bytes (qinq), enlarging the resulting subscriber frame by 4 or 8 bytes.

Using the **subscriber-vlan** command necessitates the use of the tagged (dot1q or qinq) uplink. In the uplink ingress direction (from the network to the 7705 SAR), the 7705 SAR is programmed to pop the VLAN tags. The first pop operation is mandatory, but if the frame is a single-tagged frame and there are no other VLAN tags, then the resulting untagged frame is forwarded to the subscriber interface without any errors.

Default no subscriber-vlan

Parameters *vlan-id* — specifies the VLAN ID for the subscriber VLAN

Values *, 0 to 4094 (* represents VLAN 4095 internally)

VPLS IGMP and MLD Snooping Commands

igmp-snooping

Syntax **igmp-snooping**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command enables the Internet Group Management Protocol (IGMP) snooping context.

mld-snooping

Syntax **mld-snooping**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command enables the Multicast Listener Discovery (MLD) snooping context.

disable-router-alert-check

Syntax [**no**] **disable-router-alert-check**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command enables or disables router alert checking for IGMP or MLD messages received on this interface. The router alert field in the IP header is used to inform the router to extract the packet. By default, router alert checks are enabled for IGMP and MLD packets (that is, no disable).

The **no** form of the command enables the router alert check.

Default no disable-router-alert-check

fast-leave

Syntax **no fast-leave**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command enables fast leave. When IGMP or MLD fast-leave processing is enabled, the 7705 SAR immediately removes a SAP or SDP from the multicast group when it detects an IGMP or MLD leave message on that SAP or SDP. Fast-leave processing allows the 7705 SAR to remove from the forwarding table a SAP or SDP that has sent a leave message without first sending group-specific queries to the SAP or SDP. This speeds up the process of changing channels (called zapping).

Fast leave should only be enabled when there is a single receiver present on the SAP or SDP. When fast leave is enabled, the configured **last-member-query-interval** value is ignored.

Default no fast-leave

import

Syntax **import** *policy-name*
no import

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command specifies the import routing policy used to filter IGMP or MLD packets on this SAP or SDP. Only a single policy can be imported on a single SAP or SDP at any time.

The **no** form of the command removes the policy association from the SAP or SDP.

Default no import

Parameters *policy-name* — the import policy name. Values can be a 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** *interval*
no last-member-query-interval

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command configures the maximum response time used in group-specific queries that are sent by the router acting as the querier in response to leave messages, and is also the amount of time between two consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. The configured **last-member-query-interval** is ignored when **fast-leave** is enabled on the SAP or SDP.

Default 10

Parameters *interval* — the frequency, in tenths of seconds, at which query messages are sent

Values 1 to 50

max-num-groups

Syntax **max-num-groups** *max-num-groups*
no max-num-groups

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command defines the maximum number of multicast groups that can be joined on this SAP or SDP. If the node receives an IGMP or MLD join message that would exceed the configured number of groups, the request is ignored.

The **no** form of this command disables the check.

Default no max-num-groups

Parameters *max-num-groups* — the maximum number of groups that can be joined on this SAP or SDP

Values 1 to 16000

max-num-grp-sources

Syntax **max-num-grp-sources** [*1..512*]
no max-num-grp-sources

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping

Description This command defines the maximum number of multicast (S,G)s that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of (S,G)s, the request is ignored.

The **no** form of this command disables the check.

Default no max-num-grp-sources

Parameters *1..512* — the maximum number of multicast sources allowed to be tracked per group

Values 1 to 512

max-num-sources

Syntax **max-num-sources** *max-num-sources*
no max-num-sources

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping

Description This command defines the maximum number of multicast sources that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of sources, the request is ignored.

The **no** form of this command disables the check.

Default no max-num-sources

Parameters *max-num-sources* — the maximum number of multicast sources allowed per group

Values 1 to 1000

mrouter-port

Syntax [**no**] **mrouter-port**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command specifies whether a multicast router is attached behind this SAP or SDP.

Configuring a SAP as an **mrouter-port** has two effects. First, all multicast traffic received on another SAP or SDP is copied to this SAP or SDP. Second, IGMP or MLD reports generated by the system as a result of a host joining or leaving a multicast group are sent to this SAP or SDP.

If two multicast routers exist in the network, one of them becomes the active querier and the other one (the non-querier) stops sending IGMP queries. The non-querier router continues to receive reports in order to keep its multicast trees up to date. To support this, enable **mrouter-port** on all SAPs or SDPs connected to a multicast router.

The version used for the reports (IGMPv1, v2, or v3, or MLDv1 or v2) can only be determined after an initial query has been received. Until the version is determined, no reports are sent on the SAP or spoke SDP, even if **mrouter-port** is enabled.

If the **send-queries** command is enabled on this SAP or spoke SDP, the **mrouter-port** command cannot be enabled. The **mrouter-port** and **send-queries** commands are mutually exclusive.

Default no mrouter-port

query-interval

Syntax **query-interval** *seconds*
no query-interval

Context config>service>vpls>igmp-snooping
config>service>vpls>mld-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command configures the IGMP or MLD query interval. If the **send-queries** command is enabled, this command specifies the interval between two consecutive general queries sent by the router acting as the querier 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, at which the router transmits general host-query messages

Values

1 to 65535	(at the VPLS level)
2 to 1024	(at the VPLS SAP or SDP level)

query-response-interval

Syntax **query-response-interval** *seconds*
no query-response-interval

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command configures the IGMP or MLD query response interval. If the **send-queries** command is enabled, this command specifies how long the router acting as the querier waits to receive a response from the host. 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* — the length of time to wait to receive a response to the host-query message from the host

Values 1 to 1023

query-src-ip

Syntax **query-src-ip** *ip-address*
query-src-ip *ipv6-address*
no query-src-ip

Context config>service>vpls>igmp-snooping
config>service>vpls>mld-snooping

Description This command configures the IP source address used in IGMP or MLD queries.

Parameters *ip-address* — the IPv4 address used for the IGMP query

Values a.b.c.d

ipv6-address — the IPv6 address used for the MLD query

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

report-src-ip

Syntax **report-src-ip** *ip-address*
report-src-ip *ipv6-address*
no report-src-ip

Context config>service>vpls>igmp-snooping
config>service>vpls>mld-snooping

Description This command specifies the source IP address used when generating IGMP or MLD reports. According to the IGMPv3 standard, a zero source address is allowed when 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 IPv4 source address in transmitted IGMP reports

Values a.b.c.d

ipv6-address — the source IPv6 source address in transmitted MLD reports

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

robust-count

Syntax **robust-count** *robust-count*
no robust-count

Context config>service>vpls>igmp-snooping
config>service>vpls>mld-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command allows tuning for the expected packet loss on a SAP or an SDP, and is comparable to a retry count. For a SAP or SDP, this command is functional when the **send-queries** command is enabled. If the **send-queries** command is not enabled, **robust-count** is ignored. If the SAP or SDP is expected to experience packet loss, the *robust-count* parameter may be increased. IGMP or MLD snooping on this SAP or SDP is robust up to the *robust-count* minus 1 packet.

Default 2

Parameters **robust-count** — specifies the robust count for the service, SAP, or SDP

Values

For VPLS:	1 to 255
For SAP or SDP:	2 to 7

send-queries

Syntax [**no**] **send-queries**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command specifies whether to send IGMP or MLD general query messages on the SAP or SDP.

When **send-queries** is configured, all query reports generated locally are of the type belonging to the configured **version**. If a report from a version higher than the configured version is received, the report is dropped and a "wrong version" counter is incremented. If **send-queries** is not configured, the **version** command has no effect and the version used is the version of the querier. This implies that, for example, when there is a version 2 querier, a specific version 3 group or group-source query is never sent when a host wants to leave a group.

Default no send-queries

static

Syntax **static**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command enables access to the context to configure static group addresses. Static group addresses can be configured on a SAP or SDP. When present as either a (*,G) or a (S,G) entry, multicast packets matching the configuration are forwarded even if no join message was registered for the specific group.

Default n/a

group

Syntax [**no**] **group** *grp-ip-address*
[**no**] **group** *grp-ipv6-address*

Context config>service>vpls>sap>igmp-snooping>static
config>service>vpls>sap>mld-snooping>static
config>service>vpls>mesh-sdp>igmp-snooping>static
config>service>vpls>mesh-sdp>mld-snooping>static
config>service>vpls>spoke-sdp>igmp-snooping>static
config>service>vpls>spoke-sdp>mld-snooping>static

Description This command adds a static multicast group either as a (*,G) or as one or more (S,G) records. When a static IGMP or MLD 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.

The group is not created until the **source** or **starg** is specified.

Default n/a

Parameters *grp-ip-address* — specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

Values a.b.c.d

grp-ipv6-address — specifies an MLD multicast group address that receives data on an interface. The IPv6 address must be unique for each static group.

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

source

Syntax [**no**] **source** *src-ip-address*
[**no**] **source** *src-ipv6-address*

Context config>service>vpls>sap>igmp-snooping>static>group
config>service>vpls>sap>mld-snooping>static>group
config>service>vpls>mesh-sdp>igmp-snooping>static>group
config>service>vpls>mesh-sdp>mld-snooping>static>group
config>service>vpls>spoke-sdp>igmp-snooping>static>group
config>service>vpls>spoke-sdp>mld-snooping>static>group

Description This command adds a static (S,G) entry, to allow multicast traffic for a multicast group from a specified source. For a multicast group, more than one source address can be specified. Static (S,G) entries cannot be added if a **starg** has been created previously.

The **no** form of the command removes the source from the configuration.

Default n/a

Parameters *src-ip-address* — specifies the IPv4 unicast address

Values a.b.c.d

src-ipv6-address — specifies the IPv6 unicast address

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

starg

Syntax [**no**] **starg**

Context config>service>vpls>sap>igmp-snooping>static>group
config>service>vpls>sap>mld-snooping>static>group
config>service>vpls>mesh-sdp>igmp-snooping>static>group
config>service>vpls>mesh-sdp>mld-snooping>static>group
config>service>vpls>spoke-sdp>igmp-snooping>static>group
config>service>vpls>spoke-sdp>mld-snooping>static>group

Description This command adds a static (*,G) entry to allow multicast traffic for the corresponding multicast group from any source. This command can only be enabled if there are no existing source addresses for this group.

The **no** form of the command removes the **starg** entry from the configuration.

Default no starg

version

Syntax **version** *version*
no version

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mesh-sdp>mld-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>spoke-sdp>mld-snooping

Description This command specifies the version of IGMP or MLD that is running on this SAP or SDP. This command can be used to configure a router capable of running IGMP version 1, 2, or 3, or MLD version 1 or 2. For IGMP or MLD to function correctly, all routers on a LAN must be configured to run the same version of IGMP or MLD on that LAN.

When the **send-queries** command is configured, all query reports generated locally are of the type belonging to the configured **version**. If a report from a version higher than the configured version is received, the report gets dropped and a "wrong version" counter is incremented.

If the **send-queries** command is not configured, the **version** command has no effect. The version used on that SAP or SDP is the version of the querier. This implies that, for example, when there is an IGMP version 2 querier, a specific version 3 group or group-source query is never sent when a host wants to leave a group.

The **no** form of the command returns the version to the default value.

Default 3 (IGMP)
2 (MLD)

Parameters *version* — specifies the IGMP or MLD version

Values IGMP: 1, 2, or 3
MLD: 1 or 2

VPLS STP Commands

stp

Syntax **stp**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command enables the context to configure the Spanning Tree Protocol (STP) parameters. The 7705 SAR runs the RSTP version of STP. The Alcatel-Lucent implementation of 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 service routers should not be blocked, the root path is calculated from the core perspective.

hello-time

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

Context config>service>vpls>stp

Description This command configures the Spanning Tree Protocol (STP) hello time for the Virtual Private LAN Service (VPLS) STP instance.

The **hello-time** command 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.

Since VPLS on the 7705 SAR runs in RSTP mode, the hello time is always taken from the locally configured parameter, unless RSTP fails and the SAP falls back to legacy STP operation, in which case the hello time for the spanning tree is determined by the root bridge.

The **no** form of this command returns the hello time to the default value.

Default 2 s

Parameters *hello-time* — the hello time for the STP instance, in seconds

Values 1 to 30

hold-count

Syntax **hold-count** *BDPU tx hold count*
no hold-count

Context config>service>vpls>stp

Description This command configures the peak number of BPDUs that can be transmitted in a period of 1 s.

The **no** form of this command returns the hold count to the default value.

Default 6

Parameters *BDPU tx hold count* — the hold count for the STP instance, in seconds

Values 1 to 10

mode

Syntax **mode rstp**
no mode

Context config>service>vpls>stp

Description This command specifies the version of Spanning Tree Protocol the bridge is currently running. VPLS on the 7705 SAR runs only in RSTP mode.

The **no** form of this command returns the STP variant to the default mode.

Default rstp

Parameters **rstp —** corresponds to the Rapid Spanning Tree Protocol, specified in IEEE 802.1D/D4-2003

priority

Syntax **priority** *bridge-priority*
no priority

Context config>service>vpls>stp

Description The **priority** command is used to populate the priority portion of the bridge ID field within outbound BPDUs (the most significant 4 bits of the bridge ID). It is also used as part of the decision process when determining the best BPDU between messages received and sent. All values will be truncated to multiples of 4096, conforming with IEEE 802.1t and 802.1D-2004.

The **no** form of this command returns the bridge priority to the default value.

Default By default, the bridge priority is configured to 4096, which is the highest priority

Parameters *bridge-priority* — the bridge priority for the STP instance

Values Allowed values are integers in the range of 4096 to 65535 (4096 is the highest priority).

The actual *bridge-priority* value stored and used is the allowed value (above) with its lowest 12 bits masked off, which means the actual range of values is 4096 to 61440 in increments of 4096.

managed-vlan-list

Syntax **managed-vlan-list**

Context config>service>vpls>sap

Description This command enables the context to configure VLAN ranges to be managed by a management VPLS. The list indicates, for each SAP, the ranges of associated VLANs that will be affected when the SAP changes state.

This command is only valid when the VPLS in which it is entered was created as a management VPLS.

range

Syntax [**no**] **range** *vlan-range*

Context config>service>vpls>sap>managed-vlan-list

Description This command configures a range of VLANs on an access port that are to be managed by an existing management VPLS.

This command is only valid when the VPLS in which it is entered was created as a management VPLS, and when the SAP in which it was entered was created on an Ethernet port with encapsulation type of dot1q or qinq.

To modify the range of VLANs, first the new range should be entered and then the old range removed. See Modifying VPLS Service Parameters.

Default n/a

Parameters *vlan-range* — specifies the VLAN start value and VLAN end value. The *end-vlan* value must be greater than *start-vlan value*. The format is *start-vlan-end-vlan*.

Values *start-vlan*: 0 to 4094

end-vlan: 0 to 4094

auto-edge

Syntax [**no**] **auto-edge**

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures automatic detection of the edge port characteristics of the SAP or spoke SDP.

If **auto-edge** is enabled and STP concludes there is no bridge behind the SAP or spoke SDP, the OPER_EDGE variable will dynamically be set to true. If **auto-edge** is enabled and a BPDU is received, the OPER_EDGE variable will dynamically be set to false (see edge-port).

The **no** form of this command removes automatic edge detection.

Default auto-edge

edge-port

Syntax [**no**] **edge-port**

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures the SAP or spoke SDP as an edge or non-edge port. If **auto-edge** is enabled for the SAP or spoke SDP, this value of the edge/non-edge setting will be used only as the initial value.

The function of the **edge-port** command is to tell STP that it is on the edge of the network (for example, there are no other bridges connected to that port), and, as a consequence, it can immediately transition to a forwarding state if the port becomes available.

STP, however, can detect that the actual situation is different from what **edge-port** may indicate.

Initially, the value of the SAP or spoke SDP parameter is set to **edge-port**. This value will change if:

- a BPDU is received on that port, meaning that there is another bridge connected to this port. In this case, **edge-port** becomes disabled.
- **auto-edge** is configured and no BPDU is received within a certain period of time, in which case, RSTP concludes that it is on an edge. In this case, **edge-port** becomes enabled.

The **no** form of this command returns the edge port setting to the default value.

Default no edge-port

link-type

Syntax **link-type** {**pt-pt** | **shared**}
no link-type

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command instructs STP on the maximum number of bridges behind the SAP or spoke SDP. If there is only a single bridge, transitioning to the forwarding state will be based on handshaking (fast transitions). If two or more bridges are connected via a shared media, their SAPs or spoke SDPs should 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

Parameters **pt-pt —** indicates a maximum of one bridge can exist behind this SAP or spoke SDP

shared — indicates that two or more bridges can exist behind this SAP or spoke SDP

path-cost

Syntax **path-cost** *path-cost*
no path-cost

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures the STP path cost for the SAP or spoke SDP.

The path cost is used by STP to calculate the path cost to the root bridge. The path cost in BPDUs received on the root port is incremented with the configured path cost for that SAP or spoke SDP. When BPDUs are sent out other egress SAP or spoke SDP, 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 and spoke SDPs are controlled by complex queuing dynamics, in the 7705 SAR the STP path cost is a purely static configuration.

The **no** form of this command returns the path cost to the default value.

Default 10

Parameters *path-cost* — the path cost for the SAP or spoke SDP

Values 1 to 200000000 (1 is the lowest cost)

port-num

Syntax [**no**] **port-num** *virtual-port-number*

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures the virtual port number that uniquely identifies a SAP or spoke SDP within configuration BPDUs. The internal representation of a SAP or spoke SDP is unique to a system and has a reference space much larger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a SAP or spoke SDP and identifies it with its own virtual port number, which is unique to any other SAP or spoke SDP defined on the VPLS. The virtual port number is assigned at the time that the SAP or spoke SDP is added to the VPLS. Since the order in which SAPs or spoke SDPs are added to the VPLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance.

Parameters *virtual-port-number* — the virtual port number for the SAP or spoke SDP

Values 1 to 2047

priority

Syntax **priority** *stp-priority*
no priority

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command configures the STP priority for the SAP or spoke SDP.

When configuration BPDUs are received, the STP priority is used in some circumstances as a tie-breaking mechanism to determine whether the SAP or spoke SDP will be designated or blocked. In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255) to create a 16-bit value. In the latest STP standard (802.1D-2004), only the upper 4 bits of the port priority field are used to encode the SAP or spoke SDP priority. The remaining 4 bits are used to extend the port ID field into a 12-bit virtual port number field. The virtual port number uniquely references a SAP or spoke SDP within the STP instance.

STP computes the actual priority by taking the input value and masking out the lower 4 bits.The result is the value that is stored in the SAP or spoke SDP priority parameter. For instance, if a value of 0 is entered, masking out the lower 4 bits results in a parameter value of 0. If a value of 255 is entered, the result is 240.

The **no** form of this command returns the STP priority to the default value.

Default 128

Parameters *stp-priority* — the STP priority value for the SAP or spoke SDP.

Values 0 to 255 (with masking, actual values are 0 to 240 in increments of 16)

Default 128

root-guard

Syntax [**no**] **root-guard**

Context config>service>vpls>sap>stp
config>service>vpls>spoke-sdp>stp

Description This command specifies whether this port is allowed to become an STP root port. It corresponds to the restrictedRole parameter in 802.1Q. If set, it can cause lack of spanning tree connectivity.

Default no root-guard

VPLS Filter and QoS Policy Commands

egress

Syntax **egress**

Context config>service>vpls>sap

Description This command enables the context to configure egress SAP 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>vpls>sap

Description This command enables the context to configure ingress SAP QoS policies and filter policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

agg-rate-limit

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

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description This command sets the aggregate rate limits (PIR and CIR) for the SAP. The *agg-rate* sets the PIR value. The *cir-rate* sets the CIR value. When aggregate rate limits are configured on a second-generation (Gen-2) adapter card, the scheduler mode must be set to 16-priority. On the third-generation (Gen-3) Ethernet adapter card, only the PIR component (*agg-rate*) of the **agg-rate-limit** can be set and the scheduler mode is always 4-priority-hqos. For information on adapter card generations, refer to the "Evolution of Ethernet Adapter Cards, Modules, and Platforms" section in the 7705 SAR OS Interface Configuration Guide.

Configuring the *cir-rate* is optional. If a *cir-rate* is not entered, then the *cir-rate* is set to its default value (0 kb/s). If a *cir-rate* has been set and the *agg-rate* is changed without re-entering the *cir-rate*, the *cir-rate* automatically resets to 0 kb/s. For example, to change the *agg-rate* from 2000 to 1500 while maintaining a *cir-rate* of 500, use the command **agg-rate-limit** 1500 **cir** 500.

The *cir-rate* for this command is blocked on Gen-3 Ethernet adapter cards and Gen-3 platforms with Ethernet ports, such as the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X. Gen-3 Ethernet adapter cards and Gen-3 platforms with Ethernet ports support 4-priority-hqos scheduling mode. Gen-3 platforms with a TDM MDA support 4-priority scheduling mode.

If the specified SAP is a LAG SAP, then *agg-rate* can be configured regardless of the scheduler mode setting, and *cir-rate* can be configured but is not applied to Gen-3 Ethernet adapter cards or platforms. If the active port is on a Gen-3 card or platform, then only agg-rate is applicable. If the active port is on a Gen-2 card, then agg-rate and cir-rate apply when the scheduler mode is set to 16-priority. If the active port is on a Gen-1 card, then *agg-rate* and *cir-rate* are not applicable.

The **no** form of the command sets the *agg-rate* to the maximum and the *cir-rate* to 0 kb/s.

Default no agg-rate-limit

Parameters *agg-rate* — sets the PIR for the aggregate of all the queues on the SAP. The **max** keyword applies the maximum physical port rate possible.

Values 1 to 10000000 kb/s, or **max**

Default max

cir-rate — sets the CIR for the aggregate of all the queues on the SAP. The **max** keyword applies the maximum physical port rate possible.

Values 0 to 10000000 kb/s, or **max**

Default 0 kb/s

filter

Syntax **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*]

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress
config>service>vpls>mesh-sdp>ingress
config>service>vpls>spoke-sdp>ingress

Description This command associates a filter policy with a SAP or SDP (mesh or spoke).

The **filter mac** command applies only to the **ingress** SAP and SDP contexts.

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 at a time. The filter ID must be defined before the **filter** command is executed. If the filter policy does not exist, the operation will fail and an error message will be returned.

Only one filter ID can be assigned to an interface unless the interface is dual-stack (supports both IPv4 and IPv6). A dual-stack interface can have one IPv4 and one IPv6 filter ID assigned to it.

In general, filters applied to SAPs or SDPs apply to all packets on the SAP or SDP. One exception is that 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 SDP. The filter ID itself is not removed from the system unless the **scope** of the created filter is set to **exclusive**. To avoid deletion of the filter ID and only break the association with the service object, use the **scope** command within the filter definition to set the **scope** to **template**. The default **scope** of a filter is **exclusive**.

Parameters *ip-filter-id* — specifies the IP filter policy. The filter ID or filter name must already exist within the created IP filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

ipv6-filter-id — specifies the IPv6 filter policy. The filter ID or filter name must already exist within the created IPv6 filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

mac-filter-id — specifies the MAC filter policy. The filter ID or filter name must already exist within the created MAC filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

match-qinq-dot1p

Syntax **match-qinq-dot1p** {**top** | **bottom**}
no match-qinq-dot1p

Context config>service>vpls>sap>ingress

Description This command specifies which dot1q tag position (top or bottom) in a qinq-encapsulated packet should be used when QoS evaluates dot1p classification.

The **no** form of the command restores the default dot1p evaluation behavior for the SAP, which means that the inner (bottom) tag (second tag) dot1p bits are used for classification.

By default, the dot1p bits from the inner tag service-delineating dot1q tag are used.

Table 60 shows which set of dot1p bits are used for QoS purposes when **match-qinq-dot1p** is configured. To use the table, find the row that represents the settings for **Port/SAP Type** and **Match-QinQ-Dot1q Setting**. Use the **Existing Packet Tags** column to identify which dot1q tags are available in the packet. Then use the **P-bits Used for Match** column to identify which dot1q tag contains the dot1p bits that are used for QoS dot1p classification.

Default no match-qinq-dot1p

Parameters **top —** the **top** parameter and **bottom** parameter are mutually exclusive. When the **top** parameter is specified, the outer tag's dot1p bits (topmost P-bits) are used (if existing) to match any **dot1p** *dot1p-value* entries.

bottom — the **bottom** parameter and **top** parameter are mutually exclusive. When the **bottom** parameter is specified, the bottommost P-bits (second tag's P-bits) are used (if existing) to match any **dot1p** *dot1p-value* entries.

Table 60: Match-QinQ-Dot1p Matching Behavior

Port/ SAP Type	Match-QinQ-Dot1p Setting [1]	Existing Packet Tags	P-bits Used for Match
Null	n/a	None	None
Null	n/a	Dot1p (VLAN ID 0)	None [2]
Null	n/a	Dot1q	None [2]
Null	n/a	TopQ BottomQ	None [2]
Dot1Q	n/a	None	None
Dot1Q	n/a	Dot1p (default SAP VLAN ID 0)	Dot1p P-bits
Dot1Q	n/a	Dot1q	Dot1q P-bits
QinQ/ X.Y	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.Y	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ X.0	Top	TopQ	TopQ P-bits
QinQ/ X.0	Default or Bottom	TopQ	TopQ P-bits
QinQ/ X.0	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.0	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ X.*	Top	TopQ	TopQ P-bits
QinQ/ X.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ X.*	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ 0.*	Top	None	None
QinQ/ 0.*	Default or Bottom	None	None
QinQ/ 0.*	Top	TopQ	TopQ P-bits
QinQ/ 0.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ 0.*	Top	TopQ BottomQ	TopQ P-bits

Table 60: Match-QinQ-Dot1p Matching Behavior (Continued)

Port/ SAP Type	Match-QinQ-Dot1p Setting [1]	Existing Packet Tags	P-bits Used for Match
QinQ/ 0.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ *.*	Top	None	None
QinQ/ *.*	Default or Bottom	None	None
QinQ/ *.*	Top	TopQ	TopQ P-bits
QinQ/ *.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ *.*	Top	TopQ BottomQ	TopQ P-bits
QinQ/ *.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits

Notes:

1. "Default" in this column refers to the **no** form of **match-qinq-dot1p** command.
2. For null encapsulation, the 7705 SAR does not process dot1p bits.

qinq-mark-top-only

Syntax **[no] qinq-mark-top-only**

Context config>service>vpls>sap>egress

Description When enabled, the **qinq-mark-top-only** command specifies which P-bits to mark during packet egress. When disabled, both sets of P-bits are marked. When enabled, only the P-bits in the top Q-tag are marked. The **no** form of the command is the default state (disabled).

Table 61 shows the dot1p remarking behavior for different egress port type/SAP type combinations and **qinq-mark-top-only** state, where "False" represents the default (disabled) state.

If a new tag is pushed, the dot1p bits of the new tag will be zero (unless the new tag is re-marked by the egress policy. The dot1p bits are configured using the **dot1p** parameter under the **config**>**qos** context.

Table 61: Dot1P Re-marking Behavior for the QinQ-mark-top-only Command

Egress Port Type/SAP Type	QinQ-mark-top-only State	Egress P-Bits Marked or Re-marked
Null [1]	n/a	None
Dot1q/ X [1]	n/a	Outer tag
Dot1q/ * [2]	n/a	None

Table 61: Dot1P Re-marking Behavior for the QinQ-mark-top-only Command (Continued)

Egress Port Type/SAP Type	QinQ-mark-top-only State	Egress P-Bits Marked or Re-marked
QinQ/ X.Y [1]	False	Two outer tags [3]
	True	Outer tag [3]
QinQ/ X.* [4]	True or False	Outer tag
QinQ/ X.0 [1]	True or False	Outer tag
QinQ/ 0.* [1]	True or False	None
QinQ/ *.* [2]	True or False	None

Notes:

1. This port type/SAP type is supported by the following services: Epipe, Ipipe, VPLS, IES, and VPRN.
2. This port type/SAP type is supported by the following services: Epipe and VPLS.
3. Normally, when a new tag is pushed, the dot1p bits of the new tag will be zero, unless the P-bits are remarked by the egress policy. However, an exception to this occurs when the egress SAP type is X.Y and only one new outer tag must be pushed. In this case, the new outer tag will have its dot1p bits set to the inner tag's dot1p bits.
4. This port type/SAP type is supported by the following services: Epipe, Ipipe, and VPLS.

Default no qinq-mark-top-only (disabled)

qos

Syntax **qos** *policy-id*
no qos

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description This command associates a QoS policy with an ingress or egress SAP.

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate both ingress and egress QoS policies. The **qos** command only allows ingress policies to be associated with ingress SAP, and egress policies with egress SAP. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP at one time.

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

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

Parameters *policy-id* — specifies the ingress or egress policy ID to associate with a SAP on ingress or egress. The policy ID or name must already exist.

Values 1 to 65535, or *policy-name* (up to 64 characters)

scheduler-mode

Syntax **scheduler-mode** {**4-priority** | **16-priority**}

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description This command sets the scheduler mode for the SAP and is part of the hierarchical QoS (H-QoS) feature on the 7705 SAR.

If the mode is 4-priority, then the SAP is considered an unshaped 4-priority SAP and the agg-rate-limit cannot be changed from its default values.

If the mode is 16-priority and the **agg-rate limit** parameters are configured to be non-default values, then the SAP is considered a shaped SAP. If the **agg-rate limit** parameters are left in their default settings, the SAP is considered an unshaped, 16-priority SAP.

This command is blocked on third-generation (Gen-3) Ethernet adapter cards and platforms, such as the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X with Ethernet ports, which use only 4-priority-hqos scheduling mode. It is also blocked on the 7705 SAR-X with a TDM MDA, which uses only 4-priority scheduling mode.

If the specified SAP is a LAG SAP, **scheduler-mode** can be configured but is not applied to Gen-3 adapter cards and platforms. If one of the ports in the LAG is on a Gen-1 adapter card, then **scheduler-mode** cannot be configured.

Default 4-priority

Parameters **4-priority —** sets the scheduler mode for the SAP to be 4-priority mode

16-priority — sets the scheduler mode for the SAP to be 16-priority mode

shaper-group

Syntax [**no**] **shaper-group** *shaper-group-name*

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description This command applies a shaper group to a SAP. The shaper group must already be created and must be within the shaper policy assigned to the Ethernet MDA (for ingress) or port (for egress). A shaper group is a dual-rate aggregate shaper used to shape aggregate access ingress or egress SAPs at a shaper group rate. Multiple aggregate shaper groups ensure fair sharing of available bandwidth among different aggregate shapers.

The default shaper group cannot be deleted.

The **no** form of this command removes the configured **shaper-group**.

Default shaper-group "default"

Parameters *shaper-group-name* — the name of the shaper group. To access the default shaper group, enter "default".

create — keyword used to create a shaper group

force-c-vlan-forwarding

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

Context config>service>vpls>sap

Description This command preserves the VLAN tag at the ingress SAP. The default (disabled) behavior is to strip off the VLAN tag at the ingress SAP and push a new VLAN tag at the egress SAP.

The **force-c-vlan-forwarding** command is only available on VPLS dot1q and qinq SAPs.

When the ingress and egress port **encap-type** are both dot1q, **force-c-vlan-forwarding** has the following behavior:

- if **force-c-vlan-forwarding** is enabled only at the ingress SAP

  The VLAN tag is preserved at the ingress SAP and a new VLAN tag is pushed at the egress SAP. The net effect at the egress SAP is that the packet contains two VLAN tags, where the inner tag is the preserved ingress tag.
- if **force-c-vlan-forwarding** is enabled only at the egress SAP

  A VLAN tag is not pushed at the egress SAP. The effect at the egress SAP is that the VLAN tag received at the ingress SAP is stripped off.
- if **force-c-vlan-forwarding** is enabled at both the ingress and egress SAPs

  The VLAN tag is preserved at the ingress SAP and there is no tag pushed at the egress SAP. The net effect is that the packet transmitted at the egress SAP contains the same VLAN tag that was received at the ingress SAP. Dot1p re-marking may occur at the egress SAP.

When the ingress and egress port **encap-type** are both qinq, **force-c-vlan-forwarding** has the following behavior:

- if **force-c-vlan-forwarding** is enabled only at the ingress SAP

  The inner VLAN tag is preserved at the ingress SAP and two new VLAN tags are pushed at the egress SAP. The net effect at the egress SAP is that the packet contains three VLAN tags, where the innermost tag is the preserved ingress tag.
- if **force-c-vlan-forwarding** is enabled only at the egress SAP

  Only the outer VLAN tag of the egress SAP is pushed at the egress SAP. The effect at the egress SAP is that the VLAN tag received at the ingress SAP is swapped with a new tag.
- if **force-c-vlan-forwarding** is enabled at both the ingress and egress SAPs

  The VLAN tag is preserved at the ingress SAP and the outer VLAN tag of the egress SAP is pushed. The net effect is that the packet transmitted at the egress SAP contains the same inner VLAN tag that was received at the ingress SAP. Dot1p re-marking may occur at the egress SAP.

The **no** version of this command sets the default behavior.

Default disabled

Service Billing Commands

accounting-policy

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

Context config>service>vpls>sap

Description This command creates the accounting policy context that can be applied to a SAP.

An accounting policy must be defined before it can be associated with a SAP. If the *acct-policy-id* does not exist, an error message is generated.

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

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

Default accounting-policy

Parameters *acct-policy-id* — the accounting policy ID as configured in the **config>log>accounting-policy** context

Values 1 to 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>vpls>sap

Description This command enables accounting and statistical data collection for a SAP, a network port, or an 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 CSM. 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 collect-stats

VPLS SAP DHCP Commands

dhcp

Syntax **dhcp**

Context config>service>vpls>sap

Description This command enables the context to configure DHCP parameters.

option

Syntax [**no**] **option**

Context config>service>vpls>sap>dhcp

Description This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 suboptions.

The **no** form of this command returns the system to the default.

Default no option

action

Syntax **action** [*dhcp-action*]
no action

Context config>service>vpls>sap>dhcp>option

Description This command configures the Relay Agent Information Option (Option 82) processing.

The **no** form of this command returns the system to the default value.

Default keep

Parameters *dhcp-action* — specifies the DHCP option action (**replace**, **drop**, or **keep**), as follows:

replace — in the upstream direction (from the client), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (towards the client), 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 the "Client Packets Dropped" 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 the case of Vendor Specific Options (VSOs). When the **keep** parameter is specified, the router will insert its own VSO into the Option 82 field (as per RFC 4243). This will only be done when the incoming message already has an Option 82 field. However, if adding the VSO causes the Option 82 field to exceed the maximum allowable length (255 octets), the packet is dropped.

If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

circuit-id

Syntax **circuit-id** [**ascii-tuple** | **vlan-ascii-tuple**] | **no circuit-id**

Context config>service>vpls>sap>dhcp>option

Description This command causes the router to send an ASCII-encoded "tuple" in the **circuit-id** suboption of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by "|". If no keyword is configured, then the **circuit-id** suboption will not be part of the information option (Option 82).

If disabled, the **circuit-id** suboption of the DHCP packet will be left empty.

Default no circuit-id

Parameters **ascii-tuple** — specifies that the included ASCII-encoded concatenated "tuple" consists of the access-node-identifier, service-id, and interface-name

vlan-ascii-tuple — specifies that the format will include VLAN-id and dot1p bits in addition to the **ascii-tuple** information. The format is supported on dot1q and qinq encapsulated ports only. Thus, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

remote-id

Syntax [**no**] **remote-id** [**mac** | **string** *string*]

Context config>service>vpls>sap>dhcp>option

Description This command specifies what information goes into the **remote-id** suboption in the DHCP Relay packet.

If disabled, the **remote-id** suboption 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** — specifies that the MAC address of the remote end is encoded in the suboption

string — specifies that *string* is encoded in the suboption

Values 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 (“ ”).

vendor-specific-option

Syntax [**no**] **vendor-specific-option**

Context config>service>vpls>sap>dhcp>option

Description This command configures the vendor-specific suboption within the Option 82 field of the DHCP relay packet.

client-mac-address

Syntax [**no**] **client-mac-address**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the MAC address in the vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the vendor-specific suboption of the DHCP relay packet.

sap-id

Syntax [**no**] **sap-id**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the SAP ID in the vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the vendor-specific suboption of the DHCP relay packet.

service-id

Syntax [**no**] **service-id**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command enables the sending of the service ID in the vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the vendor specific suboption of the DHCP relay packet.

string

Syntax **string** *text*
no string

Context config>service>vpls>sap>dhcp>option>vendor

Description This command specifies the string in the vendor specific suboption of the DHCP relay packet.

The **no** form of the command returns the default value.

Default no string

Parameters *text* — the string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

system-id

Syntax [**no**] **system-id**

Context config>service>vpls>sap>dhcp>option>vendor

Description This command specifies whether the system ID is encoded in the vendor specific suboption of Option 82.

snoop

Syntax [**no**] **snoop**

Context config>service>vpls>sap>dhcp

Description This command enables DHCP snooping of DHCP messages on the SAP. Enabling DHCP snooping on VPLS interfaces (SAPs) is required where vendor-specific information (as per RFC 4243) is to be inserted into the Option 82 field of the DHCP messages.

Use the **no** form of the command to disable DHCP snooping on the specified VPLS SAP binding.

Default no snoop

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 This command binds a VPLS service to an existing Service Distribution Point (SDP). Mesh SDPs bound to a service are logically treated like a single bridge "port" for flooded traffic, where flooded traffic received on any mesh SDP on the service is replicated to other "ports" (spoke SDPs and SAPs) and not transmitted on any mesh SDPs.

This command creates a binding between a service and an SDP. The SDP has an operational state 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 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 7705 SAR devices can participate in the service.

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

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

Special Cases **VPLS —** several SDPs can be bound to a VPLS. Each SDP must be destined for a different router. If two *sdp-id* bindings terminate on the same 7705 SAR, an error occurs and the second SDP binding is rejected.

Parameters *sdp-id* — the SDP identifier

Values 1 to 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 to 4294967295

vc-type — this option 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 that 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** option can still be used to define the dot1q value expected by the far-end provider equipment. A change of the 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*, as follows:

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

Values ether, vlan

ether — defines the VC type as Ethernet. The **ether** and **vlan** keywords are mutually exclusive. When the VC type is not defined, 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 **ether** and **vlan** keywords are mutually exclusive. When the VC type is not defined, the default is Ethernet for spoke SDP bindings.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] [**no-endpoint**]
spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] **endpoint** *endpoint-name*
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>vpls

Description This command binds a service to an existing SDP. A spoke SDP is treated like the equivalent of a traditional bridge "port", where flooded traffic received on the spoke SDP is replicated on all other "ports" (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state 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 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 to 17407

vc-id — the virtual circuit identifier

Values 1 to 4294967295

vc-type — this option 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 that 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** option can still be used to define the dot1q value expected by the far-end provider equipment. A change of the 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*, as follows:

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

Values ether, vlan

ether — defines the VC type as Ethernet. The **ether** and **vlan** keywords are mutually exclusive. When the VC type is not defined, 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 **ether** and **vlan** keywords are mutually exclusive. When the VC type is not defined, 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.

group-name **—** specifies the name of the split horizon group to which the SDP belongs

endpoint-name **—** specifies the service endpoint to which this SDP binding 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

control-word

Syntax [**no**] **control-word**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command enables the use of the control word on pseudowire packets in VPLS and enables the use of the control word individually on each mesh SDP or spoke SDP. By default, the control word is disabled.

When the control word is enabled, all VPLS packets 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 This command configures the egress SDP context.

ingress

Syntax **ingress**

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command configures the ingress SDP context.

vc-label

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

Context config>service>vpls>mesh-sdp>egress
config>service>vpls>spoke-sdp>egress

Description This command configures the egress VC label.

Parameters *egress-vc-label* — specifies an egress VC value that indicates a specific connection

Values 16 to 1048575

vc-label

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

Context config>service>vpls>mesh-sdp>ingress
config>service>vpls>spoke-sdp>ingress

Description This command configures the ingress VC label.

Parameters *ingress-vc-label* — specifies an ingress VC value that indicates a specific connection

Values 2048 to 18431

precedence

Syntax **precedence** [*precedence-value* | **primary**]
no precedence

Context config>service>vpls>spoke-sdp

Description This command configures the precedence of this SDP binding when there are multiple SDP bindings attached to one service endpoint. When an SDP binding goes down, the next highest precedence SDP binding begins forwarding traffic.

Default no precedence

Parameters *precedence-value* — specifies the precedence of this SDP binding

Values 1 to 4

primary — assigns this SDP as the primary spoke SDP

pw-status-signaling

Syntax [**no**] **pw-status-signaling**

Context config>service>vpls>spoke-sdp

Description This command enables pseudowire status signaling for spoke SDPs. The **no** form of this command disables pseudowire status signaling. When pseudowire status signaling is disabled, a 7705 SAR does not include the PW status TLV in the initial label mapping message of the pseudowire that is used for a spoke SDP. This forces both 7705 SAR PEs to use the pseudowire label withdrawal method for signaling pseudowire status. If the remote endpoint has **standby-signaling-master** enabled and it determines that a particular PW should be standby, based on the precedence of the PWs, it will withdraw the PW label. If the label is withdrawn for all PWs on a VPLS, the VPLS will go operationally down.

If pseudowire status signaling is enabled, the 7705 SAR includes the pseudowire status TLV in the initial label mapping message for the pseudowire.

Default pw-status-signaling

static-mac

Syntax [**no**] **static-mac** *ieee-address*

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command creates a remote static MAC entry in the VPLS FDB associated with the SDP.

In a VPLS service, MAC addresses are associated with a SAP or an SDP. MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Remote static MAC entries create a permanent MAC address to SDP association in the FDB for the VPLS instance so that the MAC address will not be learned on the edge device.

Static MAC definitions on one edge device are not propagated to other edge devices participating in the VPLS instance; that is, each edge device has an independent FDB for the VPLS.

Only one static MAC entry (local or remote) can be defined per MAC address per VPLS instance.

The **no** form of this command deletes the static MAC entry with the specified MAC address associated with the SDP from the VPLS FDB.

Default n/a

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.

vlan-vc-tag

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

Context config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command specifies an explicit dot1q value that is 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 0 is stored as the administrative dot1q value. Setting the value to 0 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

Routed VPLS Configuration Commands

allow-ip-int-binding

Syntax **allow-ip-int-binding**
no allow-ip-int-binding

Context config>service>vpls

Description This command sets a flag on the VPLS service that allows an IES or VPRN IP interface to be attached to the VPLS service in order to make the VPLS service routable. If the **allow-ip-int-binding** command is not enabled, the VPLS service cannot be attached to an IP interface.

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. The following features are disabled when the allow-ip-int-binding flag is set under VPLS:

- MAC filters
- residential SHG
- DHCP
- mVPLS
- mac-subnet-length
- GRE SDP (cannot be bound to the VPLS)



Note: The DHCP relay functionality under **service>ies>if>dhcp** can be used to dynamically assign IP addresses to the clients connected to routed VPLS SAPs.

If 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 fails. As soon as the allow-ip-int-binding flag is enabled on the VPLS, the corresponding IP interface will be attached and become operationally up. There is no need to toggle the **shutdown**/**no shutdown** command.

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 from an IES or VPRN service attached, the **no allow-ip-int-binding** command fails. Once the allow-ip-int-binding flag is reset on the VPLS service, the configuration restrictions associated with setting the flag are removed.

service-name

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

Context config>service>vpls

Description This command configures a service name, up to 64 characters in length, which adds a name identifier to a given service to use as display in show commands throughout the system.

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).

VPLS Show Commands

Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

egress-label

Syntax **egress-label** *start-label* [*end-label*]

Context show>service

Description This command displays service information 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 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 *start-label* is specified, services only using *start-label* are displayed.

Values 0, 2048 to 131071

end-label — the ending egress label value for which to display services using the label range

Values 2049 to 131071

Default the *start-label* value

Output The following output is an example of service egress-label information, and Table 62 describes the fields.

Sample Output

```
*A:ALU-48>show>service# egress-label 0 100000
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Binding        Type  I.Lbl                E.Lbl
-------------------------------------------------------------------------------
3          15:15              Spok  0                    0
5          5:5                Spok  0                    0
6          5:6                Spok  0                    0
5000       15:5000            Mesh  0                    0
5000       15:5001            Spok  0                    0
5001       5001:100           Spok  0                    0
-------------------------------------------------------------------------------
Number of Bindings Found : 6
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48>show>service#
```

Table 62: Show Service Egress Labels Output Fields

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

fdb-info

Syntax **fdb-info**

Context show>service

Description This command displays global forwarding database (FDB) usage information.

Output The following output is an example of global FDB usage information, and Table 63 describes the fields.

Sample Output

```
*A:ALU-48>show>service# fdb-info
===============================================================================
Forwarding Database(FDB) Information
===============================================================================
Service Id        : 5000               Mac Move          : Disabled
Primary Factor    : 3                  Secondary Factor  : 2
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                Total Count       : 0
Learned Count     : 0                  Static Count      : 0
OAM-learned Count : 0                  DHCP-learned Count: 0
Host-learned Count: 0
Remote Age        : 900                Local Age         : 300
High Watermark    : 95%                Low Watermark     : 90%
Mac Learning      : Enabled            Discard Unknown   : Disabled
Mac Aging         : Enabled            Relearn Only      : False
Mac Subnet Len    : 48
Incl PPP Circ-Id  : no
Service Id        : 5001               Mac Move          : Disabled
Primary Factor    : 3                  Secondary Factor  : 2
Mac Move Rate     : 2                  Mac Move Timeout  : 10
```

```
Mac Move Retries  : 3
Table Size        : 250                 Total Count       : 0
Learned Count     : 0                   Static Count      : 0
OAM-learned Count : 0                   DHCP-learned Count: 0
Host-learned Count: 0
Remote Age        : 900                 Local Age         : 300
High Watermark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabled             Discard Unknown   : Disabled
Mac Aging         : Enabled             Relearn Only      : False
Mac Subnet Len    : 48
Incl PPP Circ-Id  : no

-------------------------------------------------------------------------------
Total Service FDBs : 2
Total FDB Configured Size : 500
Total FDB Entries In Use  : 0
PBB MAC Address Indices In Use  : 0
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48>show>service#
```

Table 63: Show Service FDB-Info Output Fields

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 MACs can be relearned 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 MACs. The rate is computed as the maximum number of relearns allowed in a 5-s interval. The default rate of 10 relearns per second corresponds to 50 relearns in a 5-s period.
Mac Move Timeout	Indicates the time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled. A value of 0 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
OAM-learned Count	The current number of OAM-learned entries in the FDB of this service

Table 63: Show Service FDB-Info Output Fields (Continued)

Label	Description
DHCP-learned Count	The current number of DHCP-learned entries in the FDB of this service
Host-learned Count	The current number of host-learned 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 usage of the FDB table of this service at which a "table full" alarm is raised by the agent
Low Watermark	The usage 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 MAC aging is enabled
MAC Pinning	Specifies whether MAC pinning is enabled
Relearn Only	When enabled, indicates that either the FDB table of this service is full or that the maximum system-wide number of MACs supported by the agent has been reached, and thus MAC learning is temporarily disabled, and only MAC relearns 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
PBB MAC Address Indices in Use	Not applicable

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 The following output is an example of FDB information for a specific MAC address, and Table 64 describes the fields.

Sample Output

```
*A:ALU-48>show>service# fdb-mac
===============================================================================
Service Forwarding Database
===============================================================================
ServId    MAC               Source-Identifier    Type/Age  Last Change
-------------------------------------------------------------------------------
1         00:99:00:00:00:00  sap:1/2/7:0          Static    02/17/2007 13:58:41
-------------------------------------------------------------------------------
No. of Entries: 1
-------------------------------------------------------------------------------
Legend: L=Learned; P=MAC is protected
===============================================================================
*A:ALU-48>show>service#
```

Table 64: Show Service FDB-MAC Output Fields

Label	Description
ServId	The service ID number
MAC	The specified MAC address
Source-Identifier	The location where the MAC is defined

Table 64: Show Service FDB-MAC Output Fields (Continued)

Label	Description
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
Last Change	The date and time of the last change

id

Syntax **id** *service-id*

Context show>service

Description This command enables the context to display information for a specific *service-id*. The output display can be filtered by using one of the command filters in the Parameters list.

Parameters *service-id* — the unique service identification number or name that identifies the service in the service domain

all — display detailed information about the service (all)

arp — display ARP entries for the service (all)

base — display basic service information (base)

dhcp — display DHCP information (dhcp)

endpoint — display service endpoint information (endpoint)

fdb — display FDB entries (fdb)

labels — display labels being used by this service (labels)

mac-move — display MAC-move related information about the service (mac-move)

sap — display SAPs associated with the service (sap)

sdp — display SDPs associated with the service (pppoe-circuit-id)

split-horizon-group — display split horizon group information (split-horizon-group)

stp — display STP information (stp)

all

Syntax **all**

Context show>service>id

Description This command displays detailed information for all aspects of a service.

Output The following output is an example of detailed information about a service, and Table 65 describes the fields.

Sample Output

```
*A:ALU-48>show>service>id# all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 5001
Service Type      : VPLS
Name              : VPLS5001
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 10/26/2010 20:13:08
Last Mgmt Change  : 10/26/2010 20:13:09
Admin State       : Down                Oper State        : Down
MTU               : 1514                Def. Mesh VC Id   : 5001
SAP Count         : 1                   SDP Bind Count    : 1
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
Propagate MacFlush: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None

-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Split Horizon Group : shg5001
-------------------------------------------------------------------------------
Description       : (Not Specified)
Instance Id       : 1                   Last Change       : 10/26/2010 20:13:09

-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 5001:100  -(10.10.10.10)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 5001:100                Type              : Spoke
Split Horiz Grp   : shg5001
VC Type           : Ether                   VC Tag            : n/a
Admin Path MTU    : 0                       Oper Path MTU     : 0
Far End           : 10.10.10.10             Delivery          : MPLS

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

```
Ing mac Fltr       : n/a                       Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                       Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                       Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred             Oper ControlWord  : False
Last Status Change : 10/26/2010 20:13:08       Signaling         : TLDP
Last Mgmt Change   : 10/26/2010 20:13:09       Force Vlan-Vc     : Disabled
Endpoint           : endpoint5000              Precedence        : 4
PW Status Sig      : Enabled
Class Fwding State : Down
Flags              : SvcAdminDown SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
Time to RetryReset : never                     Retries Left      : 3
Mac Move           : Blockable                 Blockable Level   : Tertiary
Peer Pw Bits       : None
Peer Fault Ip      : None
Max Nbr of MAC Addr: No Limit                  Total MAC Addr    : 0

Learned MAC Addr   : 0                         Static MAC Addr   : 0

MAC Learning       : Enabled                   Discard Unkwn Srce: Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                     Block On Mesh Fail: False

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

Statistics           :
I. Fwd. Pkts.      : 0                         I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                         I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                         E. Fwd. Octets    : 0
E. Dro. Pkts.      : 0
Grp Enc Stats        :
I. Fwd. Pkts.      : 0                         I. Fwd. Octs.     : 0
I. Dro. Inv. Spi.  : 0                         I. Dro. OthEncPkt*: 0
E. Fwd. Pkts.      : 0                         E. Fwd. Octs.     : 0
E. Dro. Enc. Pkts. : 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

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

-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Down                      Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                       Port State        : Discarding
```

```
Port Number        : 0                        Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                      Designated Port Id: 0

Fwd Transitions    : 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/2/4:1/100
-------------------------------------------------------------------------------
Service Id         : 5001
SAP                : 1/2/4:1/100              Encap             : atm
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown L2OperDown
Multi Svc Site     : None
Last Status Change : 10/26/2010 20:13:08
Last Mgmt Change   : 10/26/2010 20:13:09
Sub Type           : regular
Split Horizon Group: shg5001

Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
Admin MTU          : 1524                     Oper MTU          : 1524
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                     qinq-pbit-marking : n/a
Ing Agg Rate Limit : max                      Egr Agg Rate Limit: max
ARP Reply Agent    : Disabled                 Host Conn Verify  : Disabled
Mac Learning       : Enabled                  Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                  Mac Pinning       : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Acct. Pol          : None                     Collect Stats     : Disabled

Anti Spoofing      : None                     Avl Static Hosts  : 0
                                              Tot Static Hosts  : 0

Calling-Station-Id : n/a
Application Profile: None

MCAC Policy Name   :                          MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                 MCAC Max Mand BW  : no limit
```

```
MCAC In use Mand BW: 0                        MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                        MCAC Avail Opnl BW: unlimited
Restr MacProt Src  : Disabled                 Restr MacUnpr Dst : Disabled
Time to RetryReset : never                    Retries Left      : 3
Mac Move           : Blockable                Blockable Level   : Tertiary
Egr MCast Grp      :
Auth Policy        : none

PPPoE Circuit-Id     : none

-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Down                     Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Unknown
Port Number        : N/A                      Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                      Designated Port   : N/A

Forward transitions: 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
Description        : (Not Specified)
Admin State        : Down                     Lease Populate    : 0
DHCP Snooping      : Down                     Action            : Keep

Proxy Admin State  : Down
Proxy Lease Time   : N/A
Emul. Server Addr  : Not Configured
-------------------------------------------------------------------------------
Subscriber Management
-------------------------------------------------------------------------------
Admin State        : Down                     MAC DA Hashing    : False
Def Sub-Id         : None
Def Sub-Profile    : None
Def SLA-Profile    : None
Sub-Ident-Policy   : None

Subscriber Limit   : 1
Single-Sub-Parameters
 Prof Traffic Only : False
```

```
 Non-Sub-Traffic   : N/A
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       n/a
Off. HiPrio           : 0                       n/a
Off. LowPrio          : n/a                     n/a

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

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

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

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : n/a                     n/a
Dro. InProf           : 0                       n/a
Dro. OutProf          : n/a                     n/a
-------------------------------------------------------------------------------
ATM SAP Configuration Information
-------------------------------------------------------------------------------
Ingress TD Profile : 1                        Egress TD Profile : 1
Alarm Cell Handling: Enabled                  AAL-5 Encap       : aal5snap-brid*
OAM Termination    : Enabled                  Periodic Loopback : Disabled

-------------------------------------------------------------------------------
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.a4:58:ff:00:00:00  Bridge Instance Id: 0
Bridge Priority    : 32768                    Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00              Bridge Max Age    : 20
```

```
Top. Change Count  : 0                          Bridge Fwd Delay  : 15
MST region revision: 0                          Bridge max hops   : 20
MST region name    :

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                          Root Forward Delay: 15
Rcvd Hello Time    : 2                          Root Max Age      : 20
Root Priority      : 32768                      Root Port         : N/A

-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 5001                Mac Move          : Disabled
Primary Factor    : 3                   Secondary Factor  : 2
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                 Total Count       : 0
Learned Count     : 0                   Static Count      : 0
OAM-learned Count : 0                   DHCP-learned Count: 0
Host-learned Count: 0
Remote Age        : 900                 Local Age         : 300
High Watermark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabled             Discard Unknown   : Disabled
Mac Aging         : Enabled             Relearn Only      : False
Mac Subnet Len    : 48
Incl PPP Circ-Id  : no

-------------------------------------------------------------------------------
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp               Oper    MRtr Pim  Send    Max  Max  MVR       Num
Id                    State   Port Port Queries Grps Srcs From-VPLS Grps
-------------------------------------------------------------------------------
sap:1/2/4:1/100       Down    No   No   No      None None Local     0
sdp:5001:100          Down    No   No   No      None None N/A       0

-------------------------------------------------------------------------------
MLD Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                 Oper    MRtr  Send      Max Num   MVR         Num
Id                      State   Port  Queries   Groups    From-VPLS   Groups
-------------------------------------------------------------------------------
sap:1/2/4:1/100         Down    No    Disabled  No Limit  Local       0
sdp:5001:100            Down    No    Disabled  No Limit  N/A         0
-------------------------------------------------------------------------------
DHCP Summary, service 5001
-------------------------------------------------------------------------------
Sap/Sdp               Snoop  Used/       Arp Reply   Info     Admin
                             Provided    Agent       Option   State
-------------------------------------------------------------------------------
sap:1/2/4:1/100       No     0/0         No          Keep     Down
```

```
sdp:5001:100           No     N/A         N/A         N/A      N/A
-------------------------------------------------------------------------------
Number of Entries : 2
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
ARP host Summary, service 5001
-------------------------------------------------------------------------------
Sap                    Used        Provided    Admin State
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No Entries found
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
Endpoint name               : endpoint5000
Description                 : (Not Specified)
Revert time                 : 0
Act Hold Delay              : 0
Ignore Standby Signaling    : false
Suppress Standby Signaling  : true
Block On Mesh Fail          : false
Tx Active                   : none
Tx Active Up Time           : 0d 00:00:00
Revert Time Count Down      : N/A
Tx Active Change Count      : 0
Last Tx Active Change       : 10/26/2010 20:13:08
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp: 5001:100 Prec:4                        Oper Status: Down

===============================================================================
VPLS Sites
===============================================================================
Site                Site-Id   Dest                Mesh-SDP  Admin   Oper  Fwdr
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
===============================================================================
*A:ALU-48>show>service>id#
```

Table 65: Show Service Service-ID (All) Output Fields

Label	Description
Service Id	The service identifier
Service Type	The type of service
Name	The service name
Description	Generic information about the service
Customer Id	The customer identifier
Last Status Change	The date and time of the most recent status change to this customer
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer
Admin State	The administrative state of the service
Oper State	The operational state of the service
MTU	The largest frame size (in octets) that the service can handle
Def. Mesh VC Id	This object is only valid in services that accept mesh SDP bindings. It is used to validate the VC ID portion of each mesh SDP binding defined in the service.
SAP Count	The number of SAPs defined on the service
SDP Bind Count	The number of SDPs bound to the service
Snd Flush on Fail	Specifies the state of sending a MAC flush on failure (enabled or disabled)
Host Conn Verify	Specifies the state of verifying host connectivity (enabled or disabled)
Propagate MacFlush	Specifies the state of propagating a MAC flush (enabled or disabled)
Def. Gateway IP	Specifies the defined gateway IP address for the service
Def. Gateway MAC	Specifies the defined gateway MAC address for the service
Split Horizon Group	
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
Instance Id	The identifier for this split horizon group instance

Table 65: Show Service Service-ID (All) Output Fields (Continued)

Label	Description
Service Distribution Points (SDPs) (see Table 81)	
Service Access Points (see Table 80)	
IGMP Snooping Base info (see Table 72)	
MLD Snooping Base info (see Table 72)	
VPLS Spanning Tree Information (see Table 83)	
DHCP Summary (see Table 68)	
ARP host Summary	
Sap	The SAP identifier
Used	The number of lease-states that are currently in use on a specific interface; that is, the number of clients on that interface who received an IP address via ARP. This value is always less than or equal to the “Provided” field.
Provided	The lease-populate value that is configured for a specific interface
Admin State	The administrative state of the service
Service Endpoints (see Table 69)	

base

Syntax **base**

Context show>service>id

Description This command displays basic information about the service ID, including service type, description, SAPs, and SDPs.

Output The following output is an example of basic information about a service, and Table 66 describes the fields.

Sample Output

```
*A:ALU-48>show>service>id# base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 5001
Service Type      : VPLS
Name              : VPLS5001
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 10/26/2010 20:13:08
Last Mgmt Change  : 10/26/2010 20:13:09
Admin State       : Down                Oper State        : Down
MTU               : 1514                Def. Mesh VC Id   : 5001
SAP Count         : 1                   SDP Bind Count    : 1
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
Propagate MacFlush: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm   Opr
-------------------------------------------------------------------------------
sap:1/2/4:1/100                         atm          1524    1524    Up    Down
sdp:5001:100 S(10.10.10.10)             n/a          0       0       Up    Down
===============================================================================
*A:ALU-48>show>service>id#
```

Table 66: Show Service Service-ID (Base) Output Fields

Label	Description
Service Id	The service identifier
Service Type	The type of service
Name	The service name
Description	Generic information about the service
Customer Id	The customer identifier
Last Status Change	The date and time of the most recent status change to this customer
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer
Admin State	The administrative state of the service

Table 66: Show Service Service-ID (Base) Output Fields (Continued)

Label	Description
Oper State	The operational state of the service
MTU	The largest frame size (in octets) that the service can handle
Def. Mesh VC Id	This object is only valid in services that accept mesh SDP bindings. It is used to validate the VC ID portion of each mesh SDP binding defined in the service.
SAP Count	The number of SAPs defined on the service
SDP Bind Count	The number of SDPs bound to the service
Snd Flush on Fail	Specifies the state of sending a MAC flush on failure (enabled or disabled)
Host Conn Verify	Specifies the state of verifying host connectivity (enabled or disabled)
Propagate MacFlush	Specifies the state of propagating a MAC flush (enabled or disabled)
Def. Gateway IP	Specifies the defined gateway IP address for the service
Def. Gateway MAC	Specifies the defined gateway MAC address for the service
Service Access and Destination Points	
Identifier	Specifies the service access (SAP) and destination (SDP) points
Type	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP
AdmMTU	Specifies the largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented
OprMTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this service to the far-end router, without requiring the packet to be fragmented
Adm	The administrative state of the service
Opr	The operational state of the service

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** *ip-int-name* | *ip-address*]

Context show>service>id>dhcp

Description This command displays DHCP statistics information.

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

sdp-id — the SDP identifier

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID for which to display information

Values 1 to 4294967295

ip-int-name — the IP interface name for which to display information

ip-address — the IP address of the interface for which to display information

Values a.b.c.d (host bits must be 0)

Output The following output is an example of DHCP statistics information for a SAP, an SDP, and an interface, and Table 67 describes the fields.

Sample Output

```
*A:ALU-48>show>service>id>dhcp# statistics
=====================================================
DHCP Statistics, service 5001
=====================================================
Client Packets Snooped               : 0
Client Packets Forwarded             : 0
Client Packets Dropped               : 0
Client Packets Proxied (RADIUS)      : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Snooped               : 0
Server Packets Forwarded             : 0
Server Packets Dropped               : 0
DHCP RELEASEs Spoofed                : 0
DHCP FORCERENEWs Spoofed             : 0
=====================================================
```

Table 67: Show Service Service-ID (DHCP Statistics) Output Fields

Label	Description
Client Packets Snooped	The number of client packets snooped
Client Packets Forwarded	The number of client packets forwarded
Client Packets Dropped	The number of client packets dropped
Client Packets Proxied (RADIUS)	The number of client packets proxied (RADIUS)
Client Packets Proxied (Lease-Split)	The number of client packets proxied (lease-split)
Server Packets Snooped	The number of server packets snooped
Server Packets Forwarded	The number of server packets forwarded
Server Packets Dropped	The number of server packets dropped
DHCP RELEASEs Spoofed	The number of DHCP releases spoofed
DHCP FORCERENEWs Spoofed	The number of DHCP forced renewals spoofed

summary

Syntax **summary**

Context show>service>id>dhcp

Description This command displays DHCP configuration summary information.

Output The following output is an example of DHCP summary information, and Table 68 describes the fields.

Sample Output

```
*A:ALU-48>show>service>id>dhcp# summary
===============================================================================
DHCP Summary, service 5001
===============================================================================
Sap/Sdp                 Snoop  Used/       Arp Reply   Info     Admin
                               Provided    Agent       Option   State
-------------------------------------------------------------------------------
sap:1/2/4:1/100         No     0/0         No          Keep     Down
sdp:5001:100            No     N/A         N/A         N/A      N/A
-------------------------------------------------------------------------------
Number of Entries : 2
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48>show>service>id>dhcp#
```

Table 68: Show Service Service-ID (DHCP Summary) Output Fields

Label	Description
Sap/Sdp	The SAP or SDP identifier of the router interface
Snoop	Specifies whether or not DHCP snooping is enabled
Used/Provided	Used: the number of lease-states that are currently in use on a specific interface; that is, the number of clients on that interface who received an IP address via DHCP. This value is always less than or equal to the "Provided" field. Provided: the lease-populate value that is configured for a specific interface
Arp Reply Agent	Specifies whether or not the ARP reply agent is enabled
Info Option	Indicates whether Option 82 processing is enabled on the interface
Admin State	Indicates the administrative state

endpoint

Syntax **endpoint** [*endpoint-name*]

Context show>service>id

Description This command displays service endpoint information.

Parameters *endpoint-name* — specifies a name for the endpoint

Values 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 (" ").

Output The following output is an example of endpoint information, and Table 69 describes the fields.

Sample Output

```
*A:ALU>show>service>id# endpoint
===============================================================================
Service 5001 endpoints
===============================================================================
Endpoint name                : endpoint5000
Description                  : (Not Specified)
Revert time                  : 0
Act Hold Delay               : 0
Ignore Standby Signaling     : false
Suppress Standby Signaling   : true
Block On Mesh Fail           : false
Tx Active                    : none
Tx Active Up Time            : 0d 00:00:00
```

```
Revert Time Count Down       : N/A
Tx Active Change Count       : 0
Last Tx Active Change        : 10/26/2010 20:13:57
-------------------------------------------------------------------------------
Members
-------------------------------------------------------------------------------
Spoke-sdp: 5001:100       Prec:4                          Oper Status: Down
===============================================================================
*A:ALU>show>service>id#
```

Table 69: Show Service Service-ID (Endpoint) Output Fields

Label	Description
Endpoint name	The name of the endpoint
Description	A description of the endpoint
Revert time	The programmable time delay to switch back to the primary spoke SDP
Act Hold Delay	Not applicable
Ignore Standby Signaling	Specifies whether or not ignore standby signaling is configured True: standby signaling is ignored False: standby signaling is not ignored
Suppress Standby Signaling	Specifies whether or not suppress standby signaling is configured True: standby signaling is suppressed False: standby signaling is not suppressed
Block On Mesh Fail	Specifies whether or not to take down the spoke SDP when the mesh SDP is down True: the spoke SDP is not taken down False: the spoke SDP is taken down
Tx Active	The identifier of the active spoke SDP
Tx Active Up Time	The total amount of time that a spoke SDP remains the active spoke SDP
Revert Time Count Down	The amount of time remaining before active transmission reverts to the primary spoke SDP
Tx Active Change Count	The number of times that the active spoke SDP has changed
Last Tx Active Change	The timestamp of the last active spoke SDP change
Members	
Spoke-sdp	Identifies the spoke SDP

Table 69: Show Service Service-ID (Endpoint) Output Fields (Continued)

Label	Description
Prec	Specifies the precedence of this SDP binding when there are multiple SDP bindings attached to one service endpoint
Oper Status	Indicates the operational status of the endpoint

fdb

Syntax **fdb** [**sap** *sap-id* | **sdp** *sdp-id* | **mac** *ieee-address* | **endpoint** *endpoint* | **detail**] [**expiry**]

Context show>service>id

Description This command displays FDB entries for a specified entity associated with the service.

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

sdp-id — specifies the SDP ID for the display

Values 1 to 17407

ieee-address — : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

endpoint — specifies an endpoint name

Values 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 (“ ”).

detail — displays detailed information

expiry — displays the time until entity is aged out

Output The following output is an example of FDB entries for a specific MAC address, and Table 70 describes the fields.

Sample Output

```
*A:ALU-48>show>service>id# fdb
===============================================================================
Forwarding Database, Service 5001
===============================================================================
Service Id        : 5001               Mac Move          : Disabled
Primary Factor    : 3                  Secondary Factor  : 2
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                Total Count       : 0
Learned Count     : 0                  Static Count      : 0
OAM-learned Count : 0                  DHCP-learned Count: 0
Host-learned Count: 0
Remote Age        : 900                Local Age         : 300
High Watermark    : 95%                Low Watermark     : 90%
```

```
Mac Learning      : Enabled             Discard Unknown   : Disabled
Mac Aging         : Enabled             Relearn Only      : False
Mac Subnet Len    : 48
Incl PPP Circ-Id  : no
===============================================================================
*A:ALU-48>show>service>id#

*A:ALU>show>service>id# fdb sap 1/2/4:1/100
===============================================================================
Forwarding Database, Service 5001
===============================================================================
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:ALU>show>service>id#

*A:ALU>show>service>id# fdb sdp 5001
===============================================================================
Forwarding Database, Service 5001
===============================================================================
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:ALU>show>service>id#

*A:ALU-48>show>service>id# fdb mac 34-34-34-34-34-34
===============================================================================
Forwarding Database, Service 5001
===============================================================================
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:ALU-48>show>service>id#

*A:ALU-48>show>service>id# fdb endpoint 10.10.10.10
===============================================================================
Forwarding Database, Service 5001
===============================================================================
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:ALU-48>show>service>id#

*A:ALU>show>service>id# fdb detail
===============================================================================
Forwarding Database, Service 5001
===============================================================================
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:ALU>show>service>id#
```

Table 70: Show Service Service-ID (FDB) Output Fields

Label	Description
ServID	The service ID
Mac Move	The administrative state of the MAC movement feature associated with this service
Primary Factor	A factor for the primary ports defining how many MAC relearn periods should be used to measure the MAC relearn rate
Secondary Factor	A factor for the secondary ports defining how many MAC relearn periods should be used to measure the MAC relearn rate
Mac Move Rate	The maximum rate at which MACs can be relearned 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 MACs The rate is computed as the maximum number of relearns allowed in a 5-s interval: for example, the default rate of 2 relearns per second corresponds to 10 relearns in a 5-s period
Mac Move Timeout	The time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled. A value of 0 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	The number of times retries are performed for re-enabling the SAP/SDP
Table Size	The maximum number of learned and static entries allowed in the FDB of this service
Total Count	The total number of learned entries 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
OAM-learned Count	The current number of OAM entries in the FDB of this service
DHCP-learned Count	The current number of DHCP-learned entries in the FDB of this service
Host-learned Count	The current number of host-learned 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.

Table 70: Show Service Service-ID (FDB) Output Fields (Continued)

Label	Description
Local Age	The number of seconds used to age out FDB entries learned on local SAPs
High Watermark	The usage of the FDB table of this service at which a table full alarm will be raised by the agent
Low Watermark	The usage of the FDB table of this service at which a table full alarm will be cleared by the agent
Mac Learning	Specifies whether MAC learning is enabled
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded
Mac Aging	Specifies whether MAC aging is enabled
Relearn Only	When enabled, displays that either the FDB table of this service is full or the maximum system-wide number of MACs supported by the agent has been reached, and thus MAC learning is temporarily disabled and only MAC relearns can take place
Mac Subnet Len	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: the number of seconds used to age out TLS FDB entries learned on local SAPs
	Age: the number of seconds used to age out TLS FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs.
	L: learned - dynamic entries created by the learning process OAM: entries created by the OAM process
	H: host, the entry added by the system for a static configured subscriber host
	D or DHCP: DHCP-installed MAC. Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease.
	P: indicates the MAC is protected by the MAC protection feature
	Static: statically configured
Last Change	Indicates the time of the most recent state changes

igmp-snooping

Syntax **igmp-snooping**

Context show>service>id

Description This command enables the context to display IGMP snooping information.

mld-snooping

Syntax **mld-snooping**

Context show>service>id

Description This command enables the context to display MLD snooping information.

all

Syntax **all**

Context show>service>id>igmp-snooping
show>service>id>mld-snooping

Description This command displays detailed information for all aspects of IGMP or MLD snooping on the VPLS service.

Output The following outputs are examples of IGMP snooping and MLD snooping displays for the specified VPLS service, and Table 71 describes the fields.

Sample Output (IGMP Snooping)

```
*A:7705:Dut-C# show service id 1 igmp-snooping all
===============================================================================
IGMP Snooping info for service 1
===============================================================================
-------------------------------------------------------------------------------
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Up
Querier     : 12.12.1.1 on SAP 1/9/7:34
-------------------------------------------------------------------------------
Sap/Sdp                   Oper MRtr Pim  Send Max   Max  Max   MVR       Num
Id                        Stat Port Port Qrys Grps  Srcs Grp   From-VPLS Grps
                                                       Srcs
-------------------------------------------------------------------------------
sap:1/9/7:12              Up   No   No   No   None  None None  Local     1
sap:1/9/7:34              Up   Yes  No   No   None  None None  Local     0
sdp:21:1                  Up   No   No   No   None  None None  N/A       1
sdp:23:1                  Up   No   No   No   None  None None  N/A       1
-------------------------------------------------------------------------------
```

```
IGMP Snooping Querier info
-------------------------------------------------------------------------------
Sap Id                  : 1/9/7:34
IP Address              : 12.12.1.1
Expires                 : 63487s
Up Time                 : 0d 00:00:38
Version                 : 3
General Query Interval  : 31744s
Query Response Interval : 1.0s
Robust Count            : 2
-------------------------------------------------------------------------------
IGMP Snooping Multicast Routers
-------------------------------------------------------------------------------
MRouter          Sap/Sdp Id                 Up Time        Expires   Version
-------------------------------------------------------------------------------
12.12.1.1        1/9/7:34                   0d 00:00:38    63487s    3
-------------------------------------------------------------------------------
Number of mrouters: 1
-------------------------------------------------------------------------------
IGMP Snooping Proxy-reporting DB
-------------------------------------------------------------------------------
Group Address    Mode     Up Time        Num Sources
-------------------------------------------------------------------------------
225.0.0.1        include  0d 00:00:37    1
-------------------------------------------------------------------------------
Number of groups: 1
-------------------------------------------------------------------------------
IGMP Snooping SAP 1/9/7:12 Port-DB
-------------------------------------------------------------------------------
Group Address   Mode    Type    From-VPLS  Up Time        Expires  Num   MC
                                                                    Src   Stdby
-------------------------------------------------------------------------------
225.0.0.1       include dynamic local      0d 00:00:34    0s       1
-------------------------------------------------------------------------------
Number of groups: 1
-------------------------------------------------------------------------------
IGMP Snooping SAP 1/9/7:34 Port-DB
-------------------------------------------------------------------------------
Group Address   Mode    Type    From-VPLS  Up Time        Expires  Num   MC
                                                                    Src   Stdby
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping SDP 21:1 Port-DB
-------------------------------------------------------------------------------
Group Address   Mode    Type    From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
225.0.0.1       include dynamic N/A        0d 00:00:37    0s       1
-------------------------------------------------------------------------------
Number of groups: 1
-------------------------------------------------------------------------------
IGMP Snooping SDP 23:1 Port-DB
-------------------------------------------------------------------------------
Group Address   Mode    Type    From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
225.0.0.1       include dynamic N/A        0d 00:00:34    0s       1
-------------------------------------------------------------------------------
Number of groups: 1
```

```
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups for SDP 5007:5007
-------------------------------------------------------------------------------
Source           Group
-------------------------------------------------------------------------------
12.12.0.12       225.0.0.1
12.12.0.13       225.0.0.1
12.12.0.14       225.0.0.1
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 3
-------------------------------------------------------------------------------
IGMP Snooping Statistics
-------------------------------------------------------------------------------
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         19            0             57
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V3 Reports              52            18            0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length               : 0
Bad IP Checksum          : 0
Bad IGMP Checksum        : 0
Bad Encoding             : 0
No Router Alert          : 0
Zero Source IP           : 0
Wrong Version            : 0
Lcl-Scope Packets        : 0
Rsvd-Scope Packets       : 0
Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
Exceeded Max Num Sources : 0
Exceeded Max Num Grp Srcs: 0
MCAC Policy Drops        : 0
MCS Failures             : 0
MVR From VPLS Cfg Drops  : 0
MVR To SAP Cfg Drops     : 0
-------------------------------------------------------------------------------
IGMP Snooping Multicast VPLS Registration info
-------------------------------------------------------------------------------
IGMP Snooping Admin State : Up
MVR Admin State           : Down
MVR Policy                : None
-------------------------------------------------------------------------------
```

```
Local SAPs/SDPs
------------------------------------------------------------------------------
Svc Id     Sap/Sdp                    Oper      From       Num Local
           Id                         State     VPLS       Groups
------------------------------------------------------------------------------
1          sap:1/9/7:12               Up        Local      1
1          sap:1/9/7:34               Up        Local      0
1          sdp:21:1                   Up        N/A        1
1          sdp:23:1                   Up        N/A        1
------------------------------------------------------------------------------
MVR SAPs (from-vpls=1)
------------------------------------------------------------------------------
Svc Id     Sap/Sdp                    Oper      From       Num MVR
           Id                         State     VPLS       Groups
------------------------------------------------------------------------------
No MVR SAPs found.
==============================================================================
*A:7705:Dut-C#
```

Sample Output (MLD Snooping)

```
*A:7705:Dut-C# show service id 1 mld-snooping all
==============================================================================
MLD Snooping info for service 1
==============================================================================
------------------------------------------------------------------------------
MLD Snooping Base info
------------------------------------------------------------------------------
Admin State : Up
Querier     : fe80::12 on SAP 1/9/7:34
------------------------------------------------------------------------------
Sap/Sdp                  Oper    MRtr  Send      Max Num   MVR        Num
Id                       State   Port  Queries   Groups    From-VPLS  Groups
------------------------------------------------------------------------------
sap:1/9/7:12             Up      No    Disabled  No Limit  Local      1
sap:1/9/7:34             Up      Yes   Disabled  No Limit  Local      0
sdp:21:1                 Up      No    Disabled  No Limit  N/A        1
sdp:23:1                 Up      No    Disabled  No Limit  N/A        1
------------------------------------------------------------------------------
MLD Snooping Querier info
------------------------------------------------------------------------------
Sap Id                  : 1/9/7:34
IP Address              : fe80::12
Expires                 : 125s
Up Time                 : 0d 00:00:38
Version                 : 2
General Query Interval  : 63s
Query Response Interval : 1.0s
Robust Count            : 2
------------------------------------------------------------------------------
MLD Snooping Multicast Routers
------------------------------------------------------------------------------
MRouter
                 Sap/Sdp Id             Up Time        Expires   Version
------------------------------------------------------------------------------
fe80::12
                 1/9/7:34               0d 00:00:38    125s      2
------------------------------------------------------------------------------
```

```
Number of mrouters: 1
-------------------------------------------------------------------------------
MLD Snooping Proxy-reporting DB
-------------------------------------------------------------------------------
Group Address
                 Mode     Up Time        Num Sources
-------------------------------------------------------------------------------
ff05::1
                 include  0d 00:00:38    1
-------------------------------------------------------------------------------
Number of groups: 1
-------------------------------------------------------------------------------
MLD Snooping SAP 1/9/7:12 Port-DB
-------------------------------------------------------------------------------
Group Address
                Mode    Type    From-VPLS  Up Time        Expires  Num   MC
                                                                    Src   Stdby
-------------------------------------------------------------------------------
ff05::1
                include dynamic local      0d 00:00:35    0s       1
-------------------------------------------------------------------------------
Number of groups: 1
-------------------------------------------------------------------------------
MLD Snooping SAP 1/9/7:34 Port-DB
-------------------------------------------------------------------------------
Group Address
                Mode    Type    From-VPLS  Up Time        Expires  Num   MC
                                                                    Src   Stdby
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
MLD Snooping SDP 21:1 Port-DB
-------------------------------------------------------------------------------
Group Address
                Mode    Type    From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
ff05::1
                include dynamic N/A        0d 00:00:38    0s       1
-------------------------------------------------------------------------------
Number of groups: 1
-------------------------------------------------------------------------------
MLD Snooping SDP 23:1 Port-DB
-------------------------------------------------------------------------------
Group Address
               Mode    Type    From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
ff05::1
                include dynamic N/A        0d 00:00:35    0s       1
-------------------------------------------------------------------------------
Number of groups: 1
-------------------------------------------------------------------------------
MLD Snooping Static Source Groups
-------------------------------------------------------------------------------
MLD Snooping Static Source Groups for SAP 1/10/6
-------------------------------------------------------------------------------
Source         Group
-------------------------------------------------------------------------------
2011::5        ff05::1
```

```
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 1
-------------------------------------------------------------------------------
MLD Snooping Statistics
-------------------------------------------------------------------------------
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         19            0             57
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              53            19            0
V1 Done                 0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad MLD Checksum        : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0
Wrong Version           : 0
Lcl-Scope Packets       : 0
Rsvd-Scope Packets      : 0
Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0
MCAC Policy Drops       : 0
MCS Failures            : 0
MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops    : 0
-------------------------------------------------------------------------------
MLD Snooping Multicast VPLS Registration info
-------------------------------------------------------------------------------
MLD Snooping Admin State  : Up
MVR Admin State           : Down
MVR Policy                : None
-------------------------------------------------------------------------------
Local SAPs/SDPs
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp               Oper      From       Num Local
           Id                    State     VPLS       Groups
-------------------------------------------------------------------------------
1          sap:1/9/7:12          Up        Local      1
1          sap:1/9/7:34          Up        Local      0
1          sdp:21:1              Up        N/A        1
1          sdp:23:1              Up        N/A        1
-------------------------------------------------------------------------------
MVR SAPs (from-vpls=1)
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp               Oper      From       Num MVR
           Id                    State     VPLS       Groups
-------------------------------------------------------------------------------
No MVR SAPs found.
===============================================================================
*A:7705:Dut-C#
```

Table 71: Show Service Service-ID (IGMP and MLD Snooping) Output Fields

Label	Description
IGMP Snooping info for service *id* MLD Snooping info for service *id*	Shows the service *id* for which snooping applies
IGMP Snooping Base info **MLD Snooping Base info**	
Admin State	The administrative state of the IGMP or MLD instance
Querier	The address of the IGMP or MLD querier on the IP subnet to which the interface is attached
Sap/Sdp Id	The SAP IDs and SDP IDs for the service ID
Oper Stat	The operational state of the SAP IDs and SDP IDs of the service ID
MRtr Port	Specifies whether the port is a multicast router port
Pim Port	(IGMP only) Specifies whether the port is a PIM port
Send Qrys Send Queries	Specifies whether the **send-queries** command is enabled or disabled
Max Grps	(IGMP only) The maximum number of multicast groups for the entry
Max Srcs	(IGMP only) The maximum number of multicast sources for the entry
Max Grp Srcs	(IGMP only) The maximum number of multicast group sources for the entry
Max Num Groups	(MLD only) The maximum number of multicast groups that can be joined on this SAP or SDP
MVR From-VPLS	Specifies multicast VPLS registration (MVR) from VPLS
Num Grps Num Groups	The actual number of multicast groups that can be joined on this SAP or SDP

Table 71: Show Service Service-ID (IGMP and MLD Snooping) Output Fields

Label	Description
IGMP Snooping Querier info **MLD Snooping Querier info**	
Sap Id	The SAP ID of the service
IP Address	The IP address of the querier
Expires	The length of time remaining before this query ages out, in seconds
Up Time	The length of time that the query has been enabled
Version	The configured version of IGMP or MLD
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	The number of IGMP or MLD message intervals that are used when there is expected packet loss on the SAP or SDP. Robust count is similar to a retry counter, where the number of retries is robust count minus 1.
IGMP Snooping Multicast Routers **MLD Snooping Multicast Routers**	
MRouter	The IP address of the multicast router
Sap/Sdp Id	The SAP ID or SDP ID connected to the multicast router
Up Time	The time since the multicast router was created
Expires	The length of time remaining before this entry ages out, in seconds
Version	The version of IGMP or MLD used by the multicast router
Number of mrouters	The number of multicast routers connected to this VPLS

Table 71: Show Service Service-ID (IGMP and MLD Snooping) Output Fields

Label	Description
IGMP Snooping Proxy-reporting DB **MLD Snooping Proxy-reporting DB**	
Group Address	The IP multicast group address for which this entry contains information
Mode	The type of membership reports received on the interface for the group: include or exclude Include — 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 or MLD membership report Exclude — reception of packets sent to the given multicast address is requested from all IP source addresses except those listed in the source-list parameter of the IGMP or MLD membership report
Up Time	The total operational time in seconds
Num Sources	The number of IGMP or MLD group-specific and source-specific queries received on this interface
Number of groups	The number of IGMP groups
IGMP Snooping SAP *id* **Port-DB** **MLD Snooping SAP** *id* **Port-DB**	
Group Address	The IP multicast group address for which this entry contains information
Mode	The type of membership reports received on the interface for the group: include or exclude Include — 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 or MLD membership report Exclude — reception of packets sent to the given multicast address is requested from all IP source addresses except those listed in the source-list parameter of the IGMP or MLD membership report
Type	Indicates how this group entry was learned: dynamic or static Dynamic — this group entry was learned by IGMP or MLD Static — this group entry was statically configured

Table 71: Show Service Service-ID (IGMP and MLD Snooping) Output Fields

Label	Description
From-VPLS	The VPLS from which the multicast streams corresponding to the groups learned via this SAP or SDP are copied. If "local", stream comes from its own VPLS.
Up Time	The time since the source group entry was created
Expires	The length of time remaining before this entry ages out, in seconds
Num Src	The number of IGMP or MLD group-specific and source-specific queries received on this SAP
MC Stdby	The multicast standby state
Number of groups	The number of groups configured for this SAP
IGMP Snooping SDP *id* **Port-DB** **MLD Snooping SDP** *id* **Port-DB**	
Group Address	The IP multicast group address for which this entry contains information
Mode	The type of membership reports received on the interface for the group: include or exclude Include — 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 or MLD membership report Exclude — reception of packets sent to the given multicast address is requested from all IP source addresses except those listed in the source-list parameter of the IGMP or MLD membership report
Type	Indicates how this group entry was learned: dynamic or static Dynamic — this group entry was learned by IGMP or MLD Static — this group entry was statically configured
From-VPLS	The VPLS from which the multicast streams corresponding to the groups learned via this SAP or SDP are copied. If "local", stream comes from its own VPLS.
Up Time	The time since the source group entry was created
Expires	The length of time remaining before this entry ages out, in seconds
Num Src	The number of IGMP or MLD group-specific and source-specific queries received on this SDP

Table 71: Show Service Service-ID (IGMP and MLD Snooping) Output Fields

Label	Description
Number of groups	The number of groups configured for this SDP
IGMP Snooping Static Source Groups **MLD Snooping Static Source Groups**	
IGMP Snooping Static Source Groups for SAP *id* **IGMP Snooping Static Source Groups for SDP** *id* **MLD Snooping Static Source Groups for SAP** *id* **MLD Snooping Static Source Groups for SDP** *id*	
Source	The source IP address of the static group
Group	The group IP address of the static (S,G) or (*,G)
Static (*,G)/(S,G) entries	The number of static (S,G) or (*, G) entries
IGMP Snooping Statistics **MLD Snooping Statistics**	
Message Type	The column heading for IGMP or MLD snooping messages
General Queries	The number of general query messages received, transmitted, and forwarded
Group Queries	The number of group query messages received, transmitted, and forwarded
Group-Source Queries	The number of group-source query messages received, transmitted, and forwarded
V1 Reports	The number of IGMPv1 or MLDv1 report messages received, transmitted, and forwarded
V2 Reports	The number of IGMPv2 or MLDv2 report messages received, transmitted, and forwarded
V3 Reports	(IGMP only) The number of IGMPv3 report messages received, transmitted, and forwarded
V2 Leaves	(IGMP only) The number of IGMPv2 leave messages received, transmitted, and forwarded
V1 Done	(MLD only) The number of MLDv1 done messages received, transmitted, and forwarded

Table 71: Show Service Service-ID (IGMP and MLD Snooping) Output Fields

Label	Description
Unknown Type	The number of unknown type messages received, transmitted, and forwarded
Drop Statistics	
Bad Length	The number of packets dropped due to bad length
Bad IP Checksum	(IGMP only) The number of packets dropped due to a bad IP checksum
Bad IGMP Checksum Bad MLD Checksum	The number of packets dropped due to a bad IGMP or MLD checksum
Bad Encoding	The number of packets dropped due to bad encoding
No Router Alert	The number of packets dropped because there was no router alert
Zero Source IP	The number of packets dropped due to a source IP address of 0.0.0.0 or 00:00:00:00:00:00:00:00
Wrong Version	The number of packets dropped due to a wrong version of IGMP or MLD
Lcl-Scope Packets	The number of local scope packets dropped
Rsvd-Scope Packets	The number of reserved scope packets dropped
Send Query Cfg Drops	The number of messages dropped because of send query configuration errors
Import Policy Drops	The number of messages dropped because of import policy
Exceeded Max Num Groups	The number of packets dropped because the maximum number of groups has been exceeded
Exceeded Max Num Sources	(IGMP only) The number of packets dropped because the maximum number of sources has been exceeded
Exceeded Max Num Grp Srcs	(IGMP only) The number of packets dropped because the maximum number of group sources has been exceeded
MCAC Policy Drops	The number of packets dropped due to multicast CAC
MCS Failures	The number of packets dropped due to multicast server (MCS) failures
MVR From VPLS Cfg Drops	The number of packets dropped due to VPLS configuration multicast VPLS registration (MVR)

Table 71: Show Service Service-ID (IGMP and MLD Snooping) Output Fields

Label	Description
MVR To SAP Cfg Drops	The number of packets dropped due to SAP configuration
IGMP Snooping Multicast VPLS Registration info **MLD Snooping Multicast VPLS Registration info**	
IGMP Snooping Admin State MLD Snooping Admin State	The administrative state of IGMP or MLD snooping
MVR Admin State	The administrative state of received MVR
MVR Policy	The MVR policy
Local SAPs/SDPs	
Svc Id	The service identifier for this SAP or SDP entry
Sap/Sdp Id	The local SAP or SDP identifier
Oper State	The operational state of this SAP or SDP
From VPLS	The VPLS from which the multicast streams corresponding to the groups learned via this SAP or SDP are copied. If "local", stream comes from its own VPLS.
Num Local Groups	The number of local groups associated with this SAP or SDP
MVR SAPs (from-vpls=*id*)	
Svc Id	The service identifier for this SAP or SDP entry
Sap/Sdp Id	The MVR SAP or SDP identifier
Oper State	The operational state of this SAP or SDP
From VPLS	The VPLS from which the multicast streams corresponding to the groups learned via this SAP or SDP are copied. If "local", stream comes from its own VPLS.
Num MVR Groups	The number of MVR groups associated with this SAP

base

Syntax **base**

Context show>service>id>igmp-snooping
show>service>id>mld-snooping

Description This command displays basic information about IGMP snooping or MLD snooping for the VPLS service.

Output The following outputs are examples of basic IGMP snooping and MLD snooping displays for the specified VPLS service, and Table 72 describes the fields.

Sample Output

```
*A:7705custDoc:Sar18>show>service>id# igmp-snooping base
===============================================================================
IGMP Snooping Base info for service 5007
===============================================================================
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                  Oper MRtr Pim  Send Max   Max  Max   MVR       Num
Id                       Stat Port Port Qrys Grps  Srcs Grp   From-VPLS Grps
                                                      Srcs
-------------------------------------------------------------------------------
sap:1/10/6               Down No   No   No   None  None None  Local     0
sdp:5007:5007            Down Yes  No   No   None  None None  N/A       0
===============================================================================
*A:7705custDoc:Sar18>show>service>id# igmp-snooping

*A:7705custDoc:Sar18>show>service>id# mld-snooping base
===============================================================================
MLD Snooping Base info for service 5001
===============================================================================
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                  Oper    MRtr  Send      Max Num   MVR        Num
Id                       State   Port  Queries   Groups    From-VPLS  Groups
-------------------------------------------------------------------------------
sap:1/6/4                Down    No    Disabled  No Limit  Local      0
sdp:5001:5001            Down    No    Disabled  No Limit  N/A        0
===============================================================================
*A:7705custDoc:Sar18>show>service>id#
```

Table 72: Show Service Service-ID IGMP and MLD Snooping (Base) Output Fields

Label	Description
Admin State	The administrative state of the IGMP or MLD instance
Querier	The address of the IGMP or MLD querier on the IP subnet to which the interface is attached
Sap/Sdp Id	The SAP IDs and SDP IDs for the service ID
Oper Stat	The operational state of the SAP IDs and SDP IDs of the service ID
MRtr Port	Specifies whether the port is a multicast router port
Pim Port	(IGMP only) Specifies whether the port is a PIM port
Send Qrys Send Queries	Specifies whether the **send-queries** command is enabled or disabled
Max Grps	(IGMP only) The maximum number of multicast groups for the entry
Max Srcs	(IGMP only) The maximum number of multicast sources for the entry
Max Grp Srcs	(IGMP only) The maximum number of multicast group sources for the entry
Max Num Groups	(MLD only) The maximum number of multicast groups that can be joined on this SAP or SDP
MVR From-VPLS	Specifies MVR from VPLS
Num Grps Num Groups	The actual number of multicast groups that can be joined on this SAP or SDP

port-db

Syntax **port-db sap** *sap-id* [**detail**]
port-db sap *sap-id* **group** *grp-ip-address*
port-db sap *sap-id* **group** *grp-ipv6-address*
port-db sdp *sdp-id:vc-id* [**detail**]
port-db sdp *sdp-id:vc-id* **group** *grp-ip-address*
port-db sdp *sdp-id:vc-id* **group** *grp-ipv6-address*

Context show>service>id>igmp-snooping
show>service>id>mld-snooping

Description This command displays information on the IGMP or MLD snooping port database for the VPLS service.

Parameters *grp-ip-address* — displays the IGMP snooping port database for the specified multicast group address

Values a.b.c.d

grp-ipv6-address — displays the MLD snooping port database for the specified multicast group address

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

sap-id — displays the IGMP or MLD snooping port database for the specified SAP. See Table 39 for a full list of SAP IDs.

sdp-id — displays the IGMP or MLD snooping port database for the specified SDP. For a spoke SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID for which to display information

Values 1 to 4294967295

Output The following outputs are examples of port database information for IGMP snooping and MLD snooping displays for the specified VPLS service, and Table 73 describes the fields.

Sample Output (IGMP Snooping)

```
*A:7705custDoc:Sar18>show>service>id>igmp-snooping# port-db sap 1/10/6
===============================================================================
IGMP Snooping SAP 1/10/6 Port-DB for service 5007
===============================================================================
Group Address   Mode    Type    From-VPLS  Up Time        Expires  Num   MC
                                                                    Src   Stdby
-------------------------------------------------------------------------------
225.0.0.2       exclude static  local      29d 01:28:14   never    0
225.0.0.3       include static  local      29d 01:27:34   never    2
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================

*A:7705custDoc:Sar18>show>service>id>igmp-snooping# port-db sap 1/10/6 detail
===============================================================================
IGMP Snooping SAP 1/10/6 Port-DB for service 5007
===============================================================================
-------------------------------------------------------------------------------
IGMP Group 225.0.0.2
-------------------------------------------------------------------------------
Mode             : exclude              Type             : static
Up Time          : 29d 01:28:18         Expires          : never
Compat Mode      : IGMP Version 3
V1 Host Expires  : 0s                   V2 Host Expires  : 0s
MVR From-VPLS    : local                MVR To-SAP       : local
MC Standby       : no
-------------------------------------------------------
Source Address  Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------------
No sources.
-------------------------------------------------------------------------------
IGMP Group 225.0.0.3
-------------------------------------------------------------------------------
Mode             : include              Type             : static
Up Time          : 29d 01:27:38         Expires          : never
Compat Mode      : IGMP Version 3
V1 Host Expires  : 0s                   V2 Host Expires  : 0s
MVR From-VPLS    : local                MVR To-SAP       : local
MC Standby       : no
-------------------------------------------------------
Source Address  Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------------
11.12.12.12     29d 01:27:38 never    static   Fwd
11.12.13.13     29d 01:21:48 never    static   Fwd
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================

*A:7705:Dut-C# show service id 1 igmp-snooping port-db sdp 23:1

===============================================================================
IGMP Snooping SDP 23:1 Port-DB for service 1
===============================================================================
Group Address   Mode    Type    From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
225.0.0.1       include dynamic N/A        0d 00:10:48    0s       2
```

```
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
*A:7705:Dut-C#
```

Sample Output (MLD Snooping)

```
*A:7705custDoc:Sar18>show>service>id>mld-snooping# port-db sap 1/10/6
===============================================================================
MLD Snooping SAP 1/10/6 Port-DB for service 5007
===============================================================================
Group Address
                Mode    Type    From-VPLS  Up Time        Expires  Num   MC
                                                                    Src   Stdby
-------------------------------------------------------------------------------
ff05::1
                include static  local      39d 06:42:40   never    1
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================

*A:7705custDoc:Sar18>show>service>id>mld-snooping# port-db  sap  1/10/6  detail
===============================================================================
MLD Snooping SAP 1/10/6 Port-DB for service 5007
===============================================================================
-------------------------------------------------------------------------------
MLD Group ff05::1
-------------------------------------------------------------------------------
Mode             : include              Type             : static
Up Time          : 38d 02:54:18         Expires          : never
Compat Mode      : MLD Version 2
V1 Host Expires  : 0s                                    :
MVR From-VPLS    : local                MVR To-SAP       : local
MC Standby       : no
-------------------------------------------------------
Source Address
                Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------------
2011::5
                38d 02:54:18 never    static   Fwd
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
*A:7705custDoc:Sar18>show>service>id>mld-snooping#

*A:7705custDoc:Sar18>show>service>id>mld-snooping# port-db sdp 5007:5007
===============================================================================
MLD Snooping SDP 5007:5007 Port-DB for service 5007
===============================================================================
Group Address
                Mode    Type    From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
ff05::1
                exclude static  N/A        39d 06:45:13   never    0
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
```

```
*A:7705custDoc:Sar18>show>service>id>mld-snooping# port-db sdp 5007:5007 detail
===============================================================================
MLD Snooping SDP 5007:5007 Port-DB for service 5007
===============================================================================
-------------------------------------------------------------------------------
MLD Group ff05::1
-------------------------------------------------------------------------------
Mode             : exclude              Type             : static
Up Time          : 39d 06:45:17         Expires          : never
Compat Mode      : MLD Version 2
V1 Host Expires  : 0s                                    :
MVR From-VPLS    : N/A                  MVR To-SAP       : N/A
-------------------------------------------------------
Source Address
                Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------------
No sources.
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
```

Table 73: Show Service Service-ID IGMP and MLD Snooping (Port-DB) Output Fields

Label	Description
Group Address	The IP multicast group address for which this entry contains information
Mode	The type of membership reports received on the interface for the group: include or exclude Include — 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 or MLD membership report Exclude — 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: dynamic or static Dynamic — this group entry was learned by IGMP or MLD Static — this group entry was statically configured
From-VPLS	The VPLS from which the multicast streams corresponding to the groups learned via this SAP or SDP are copied. If “local”, stream comes from its own VPLS.
Up Time	The time since the source group entry was created
Expires	The length of time remaining before this entry ages out, in seconds
Num Src	Indicates the number of IGMP or MLD group and source specific queries received on this SAP

Table 73: Show Service Service-ID IGMP and MLD Snooping (Port-DB) Output Fields

Label	Description
MC Stdby	(SAP only) Indicates the multicast standby state
Number of groups	Indicates the number of groups configured for this SAP or SDP
IGMP Group MLD Group	The IPv4 or IPv6 multicast group address for which this entry contains information
Compat Mode	The IGMP or MLD compatibility 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; MLD must operate in version 1 mode. IGMPv3 and MLDv2 hosts must track the compatibility mode of each attached network on a per-local-interface basis. A host's compatibility mode is determined from the host compatibility mode variable, which can be in one of three states for IGMP (IGMPv1, IGMPv2, or IGMPv3) or one of two states for MLD (MLDv1 or MLDv2). The compatibility mode variable is kept on a per-interface basis and is dependent on the version of general queries heard on that interface and the older version querier present timers for the interface.
V1 Host Expires	The time remaining until the local router assumes that there are no longer any IGMPv1 members on the IP subnet attached to this interface. Upon receiving an IGMPv1 membership report, this value is reset to the setting of the group membership timer. While the time remaining on the timer is non-zero, the local router ignores any IGMPv3 leave messages for this group that are received on this interface.
V2 Host Expires	The time remaining until the local router assumes that there are no longer any IGMPv2 members on the IP subnet attached to this interface. Upon receiving an IGMPv2 membership report, this value is reset to the setting of the group membership timer. While the time remaining on the timer is non-zero, the local router ignores any IGMPv3 leave messages for this group that are received on this interface.
MVR From-VPLS	The VPLS from which the multicast VPLS registration (MVR) streams corresponding to the groups learned via this SAP or SDP are copied. If “local”, it is from its own VPLS.
MVR To-SAP	The SAP to which the multicast VPLS registration (MVR) streams corresponding to the groups learned via this SAP or SDP are copied. If “local”, it is from its own VPLS.
Source Address	The source address for which this entry contains information

Table 73: Show Service Service-ID IGMP and MLD Snooping (Port-DB) Output Fields

Label	Description
Fwd/Blk	Indicates whether the corresponding multicast stream will be blocked or forwarded

proxy-db

Syntax **proxy-db** [**detail**]
proxy-db group *grp-ip-address*
proxy-db group *grp-ipv6-address*

Context show>service>id>igmp-snooping
show>service>id>mld-snooping

Description This command displays information on the IGMP or MLD snooping proxy reporting database for the VPLS service.

Parameters *grp-ip-address* — displays the IGMP snooping proxy reporting database for the specified multicast group address

Values a.b.c.d

grp-ipv6-address — displays the MLD snooping proxy reporting database for the specified multicast group address

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

Output The following outputs are examples of proxy reporting database information for IGMP snooping and MLD snooping displays for the specified VPLS service, and Table 74 describes the fields.

Sample Output

```
*A:7705:Dut-C# show service id 1 igmp-snooping proxy-db detail

===============================================================================
IGMP Snooping Proxy-reporting DB for service 1
===============================================================================
-------------------------------------------------------------------------------
IGMP Group 225.0.0.1
-------------------------------------------------------------------------------
Up Time : 0d 00:03:21                  Mode : exclude
------------------------------
Source Address  Up Time
------------------------------
No sources.
-------------------------------------------------------------------------------
Number of groups: 1
```

```
===============================================================================
*A:7705:Dut-C#

*A:7705:Sar18>show>service>id>igmp-snooping# proxy-db detail
===============================================================================
IGMP Snooping Proxy-reporting DB for service 5007
===============================================================================
-------------------------------------------------------------------------------
IGMP Group 225.0.0.2
-------------------------------------------------------------------------------
Up Time : 28d 01:13:36                  Mode : exclude
------------------------------
Source Address  Up Time
------------------------------
No sources.
-------------------------------------------------------------------------------
IGMP Group 225.0.0.3
-------------------------------------------------------------------------------
Up Time : 28d 01:12:56                  Mode : include
------------------------------
Source Address  Up Time
------------------------------
11.12.12.12     28d 01:12:56
11.12.13.13     28d 01:07:06
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================

*A:7705:Sar18>show>service>id>igmp-snooping# proxy-db group 225.0.0.3
===============================================================================
IGMP Snooping Proxy-reporting DB for service 5007
===============================================================================
-------------------------------------------------------------------------------
IGMP Group 225.0.0.3
-------------------------------------------------------------------------------
Up Time : 28d 01:25:31                  Mode : include
------------------------------
Source Address  Up Time
------------------------------
11.12.12.12     28d 01:25:31
11.12.13.13     28d 01:19:41
===============================================================================

*A:7705:Dut-C# show service id 1 mld-snooping proxy-db

===============================================================================
MLD Snooping Proxy-reporting DB for service 1
===============================================================================
Group Address
                 Mode     Up Time         Num Sources
-------------------------------------------------------------------------------
ff05::1
                 exclude  0d 00:02:18    0
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
*A:7705:Dut-C#
```

Table 74: Show Service Service-ID IGMP and MLD Snooping (Proxy-DB) Output Fields

Label	Description
IGMP Group Address	The IP multicast group address for which this entry contains information
Up Time	The time since the source group entry was created
Mode	The type of membership reports received on the interface for the group: include or exclude Include — 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 or MLD membership report Exclude — reception of packets sent to the given multicast address is requested from all IP source addresses except those listed in the source-list parameter
Source Address	The source address for which this entry contains information
Num Sources	Indicates the number of IGMP or MLD group and source specific queries received on this SAP
Number of groups	Indicates the number of groups configured for this SAP or SDP

querier

Syntax **querier**

Context show>service>id>igmp-snooping
show>service>id>mld-snooping

Description This command displays information on the current IGMP or MLD snooping querier for the VPLS service.

Output The following outputs are examples of querier information for IGMP snooping and MLD snooping displays for the specified VPLS service, and Table 75 describes the fields.

Sample Output

```
*A:ALA-1>show>service>id>igmp-snooping# querier
===============================================================================
IGMP Snooping Querier info for service 10
===============================================================================
Sap Id                  : 1/1/1
IP Address              : 10.10.10.1
Expires                 : 6s
Up Time                 : 0d 00:56:50
```

```
Version                  : 3
General Query Interval   : 5s
Query Response Interval  : 2.0s
Robust Count             : 2
===============================================================================
*A:ALA-1>show>service>id>snooping#
```

Table 75: Show Service Service-ID IGMP and MLD Snooping (Querier) Output Fields

Label	Description
SAP Id	The SAP ID of the service
IP Address	The IP address of the querier
Expires	The length of time remaining before this query ages out, in seconds
Up Time	The length of time that the query has been enabled
Version	The configured version of IGMP or MLD
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	The number of IGMP or MLD message intervals that are used when there is expected packet loss on the SAP or SDP. Robust count is similar to a retry counter, where the number of retries is robust count minus one packet.

static

Syntax **static** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>igmp-snooping
show>service>id>mld-snooping

Description This command displays information on static IGMP or MLD snooping source group membership for the VPLS service.

Parameters *sap-id* — displays static IGMP or MLD snooping source groups for the specified SAP. See Table 39 for a full list of SAP IDs.

sdp-id — displays the IGMP or MLD snooping source groups for the specified SDP

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID for which to display information

Values 1 to 4294967295

Output The following outputs are examples of static SAP or SDP multicast source and group information for IGMP snooping and MLD snooping displays for the specified VPLS service, and Table 76 describes the fields.

Sample Output

```
*A:7705custDoc:Sar18>show>service>id>igmp-snooping# static
===============================================================================
IGMP Snooping Static Source Groups for service 5007
===============================================================================
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups for SAP 1/10/6
-------------------------------------------------------------------------------
Source            Group
-------------------------------------------------------------------------------
*                 225.0.0.2
11.12.12.12       225.0.0.3
11.12.13.13       225.0.0.3
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 3
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups for SDP 5007:5007
-------------------------------------------------------------------------------
Source            Group
-------------------------------------------------------------------------------
12.12.0.12        225.0.0.1
12.12.0.13        225.0.0.1
12.12.0.14        225.0.0.1
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 3
===============================================================================
*A:7705custDoc:Sar18>show>service>id>igmp-snooping#

*A:7705custDoc:Sar18>show>service>id>igmp-snooping# static sap 1/10/6
===============================================================================
IGMP Snooping Static Source Groups for SAP 1/10/6 (service 5007)
===============================================================================
Source            Group
-------------------------------------------------------------------------------
*                 225.0.0.2
11.12.12.12       225.0.0.3
11.12.13.13       225.0.0.3
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 3
===============================================================================
*A:7705custDoc:Sar18>show>service>id>igmp-snooping#

*A:7705custDoc:Sar18>show>service>id>igmp-snooping# static sdp 5007:5007
===============================================================================
IGMP Snooping Static Source Groups for SDP 5007:5007 (service 5007)
===============================================================================
Source            Group
-------------------------------------------------------------------------------
12.12.0.12        225.0.0.1
12.12.0.13        225.0.0.1
12.12.0.14        225.0.0.1
-------------------------------------------------------------------------------
```

```
Static (*,G)/(S,G) entries: 3
===============================================================================


*A:7705custDoc:Sar18>show>service>id>mld-snooping# static sdp 5007:5007
===============================================================================
MLD Snooping Static Source Groups for SDP 5007:5007 (service 5007)
===============================================================================
Source
                Group
-------------------------------------------------------------------------------
*
                ff05::1
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 1
===============================================================================
```

Table 76: Show Service Service-ID IGMP and MLD Snooping (Static) Output Fields

Label	Description
Source	The source IP address of the static group
Group	The group IP address of the static (S,G) or (*/G)
Static (*,G)/(S,G) entries	The number of static (S,G) or (*, G) entries

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>igmp-snooping
show>service>id>mld-snooping

Description This command displays IGMP or MLD snooping statistics information for the specified VPLS SAP or SDP.

Parameters *sap-id* — displays IGMP or MLD snooping statistics for the specified SAP. See Table 39 for a full list of SAP IDs.

sdp-id — displays the IGMP or MLD snooping statistics for the specified SDP

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID for which to display information

Values 1 to 4294967295

Output The following outputs are examples of multicast statistics for SAPs and SDPs for IGMP snooping and MLD snooping displays for the specified VPLS service, and Table 77 describes the fields.

Sample Output

```
*A:7705custDoc:Sar18>show>service>id>igmp-snooping# statistics sap 1/10/6
===============================================================================
IGMP Snooping Statistics for SAP 1/10/6 (service 5007)
===============================================================================
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         0             0             0
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V3 Reports              0             0             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length               : 0
Bad IP Checksum          : 0
Bad IGMP Checksum        : 0
Bad Encoding             : 0
No Router Alert          : 0
Zero Source IP           : 0
Wrong Version            : 0
Lcl-Scope Packets        : 0
Rsvd-Scope Packets       : 0

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
Exceeded Max Num Sources : 0
Exceeded Max Num Grp Srcs: 0
MCAC Policy Drops        : 0
MCS Failures             : 0

MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops    : 0
===============================================================================
*A:7705custDoc:Sar18>show>service>id>igmp-snooping#

*A:7705custDoc:Sar18>show>service>id>igmp-snooping# statistics sdp 5007:5007
===============================================================================
IGMP Snooping Statistics for SDP 5007:5007 (service 5007)
===============================================================================
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         0             0             0
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V3 Reports              0             0             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
```

```
Drop Statistics
-------------------------------------------------------------------------------
Bad Length               : 0
Bad IP Checksum          : 0
Bad IGMP Checksum        : 0
Bad Encoding             : 0
No Router Alert          : 0
Zero Source IP           : 0
Wrong Version            : 0
Lcl-Scope Packets        : 0
Rsvd-Scope Packets       : 0

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
Exceeded Max Num Sources : 0
Exceeded Max Num Grp Srcs: 0
MCAC Policy Drops        : 0
===============================================================================
*A:7705custDoc:Sar18>show>service>id>igmp-snooping#

*A:7705custDoc:Sar18>show>service>id>mld-snooping# statistics sap 1/10/6
===============================================================================
MLD Snooping Statistics for SAP 1/10/6 (service 5007)
===============================================================================
Message Type             Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries          0             0             0
Group Queries            0             0             0
Group-Source Queries     0             0             0
V1 Reports               0             0             0
V2 Reports               0             0             0
V1 Done                  0             0             0
Unknown Type             0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length               : 0
Bad MLD Checksum         : 0
Bad Encoding             : 0
No Router Alert          : 0
Zero Source IP           : 0
Wrong Version            : 0
Lcl-Scope Packets        : 0
Rsvd-Scope Packets       : 0

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
MCS Failures             : 0

MVR From VPLS Cfg Drops  : 0
MVR To SAP Cfg Drops     : 0
===============================================================================
```

```
*A:7705custDoc:Sar18>show>service>id>mld-snooping# statistics sdp  5007:5007
===============================================================================
MLD Snooping Statistics for SDP 5007:5007 (service 5007)
===============================================================================
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         0             0             0
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad MLD Checksum        : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0
Wrong Version           : 0
Lcl-Scope Packets       : 0
Rsvd-Scope Packets      : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0
===============================================================================
```

Table 77: Show Service Service-ID IGMP and MLD Snooping (Statistics) Output Fields

Label	Description
Message Type	The column heading for IGMP or MLD snooping messages
General Queries	The number of general query messages received, transmitted, and forwarded
Group Queries	The number of group query messages received, transmitted, and forwarded
Group-Source Queries	The number of group-source query messages received, transmitted, and forwarded
V1 Reports	The number of IGMPv1 or MLDv1 report messages received, transmitted, and forwarded
V2 Reports	The number of IGMPv2 or MLDv2 report messages received, transmitted, and forwarded
V3 Reports	(IGMP only) The number of IGMPv3 report messages received, transmitted, and forwarded

Table 77: Show Service Service-ID IGMP and MLD Snooping (Statistics) Output Fields (Continued)

Label	Description
V2 Leaves	(IGMP only) The number of IGMP leave messages received, transmitted, and forwarded
V1 Done	(MLD only) The number of MLD done messages received, transmitted, and forwarded
Unknown Type	The number of unknown type messages received, transmitted, and forwarded
Drop Statistics	
Bad Length	The number of packets dropped due to bad length
Bad IP Checksum	(IGMP only) The number of packets dropped due to a bad IP checksum
Bad IGMP Checksum Bad MLD Checksum	The number of packets dropped due to a bad IGMP or MLD checksum
Bad Encoding	The number of packets dropped due to bad encoding
No Router Alert	The number of packets dropped because there was no router alert
Zero Source IP	The number of packets dropped due to a source IP address of 0.0.0.0 or 00:00:00:00:00:00:00:00
Wrong Version	The number of packets dropped due to a wrong version of IGMP or MLD
Lcl-Scope Packets	The number of local scope packets dropped
Rsvd-Scope Packets	The number of reserved scope packets dropped
Send Query Cfg Drops	The number of messages dropped because of send query configuration errors
Import Policy Drops	The number of messages dropped because of import policy
Exceeded Max Num Groups	The number of packets dropped because the maximum number of groups has been exceeded
Exceeded Max Num Sources	The number of packets dropped because the maximum number of sources has been exceeded
Exceeded Max Num Grp Srcs	The number of packets dropped because the maximum number of group sources has been exceeded

Table 77: Show Service Service-ID IGMP and MLD Snooping (Statistics) Output Fields (Continued)

Label	Description
MCAC Policy Drops	The number of packets dropped due to multicast CAC
MCS Failures	The number of packets dropped due to multicast server (MCS) failures
MVR From VPLS Cfg Drops	The number of packets dropped due to VPLS configuration multicast VPLS registration (MVR)
MVR To SAP Cfg Drops	The number of packets dropped due to SAP configuration

labels

Syntax **labels**

Context show>service>id

Description This command displays information about ingress and egress labels for the specified service.

Output The following output is an example of service label information, and Table 78 describes the fields.

Sample Output

```
*A:ALU-48>show>service>id# labels
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Binding        Type  I.Lbl                E.Lbl
-------------------------------------------------------------------------------
5001       5001:100           Spok  0                    0
-------------------------------------------------------------------------------
Number of Bound SDPs : 1
-------------------------------------------------------------------------------
```

Table 78: Show Service Service-ID (Labels) Output Fields

Label	Description
Svc Id	The service identifier
Sdp Binding	The SDP binding identifier
Type	Indicates whether the SDP is spoke or mesh
I.Lbl	The ingress label used by the far-end device to send packets to this device in this service by the SDP

Table 78: Show Service Service-ID (Labels) Output Fields (Continued)

Label	Description
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 Bound SDPs	The number of SDP bindings

mac-move

Syntax **mac-move**

Context show>service>id

Description This command displays information related to the **mac-move** feature for the specified service.

Output The following output is an example of service MAC move information, and Table 79 describes the fields.

Sample Output

```
*A:ALU-48>show>service>id# mac-move
===============================================================================
Service Mac Move Information
===============================================================================
Service Id        : 5001               Mac Move          : Disabled
Primary Factor    : 3                  Secondary Factor  : 2
Mac Move Rate     : 2                  Mac Move Timeout  : 10
Mac Move Retries  : 3
-------------------------------------------------------------------------------
SAP Mac Move Information: 1/2/4:1/100
-------------------------------------------------------------------------------
Admin State       : Up                 Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown L2OperDown
Time to RetryReset: never              Retries Left      : 3
Mac Move          : Blockable          Blockable Level   : Tertiary
-------------------------------------------------------------------------------
SDP Mac Move Information: 5001:100
-------------------------------------------------------------------------------
Admin State       : Up                 Oper State        : Down
Flags              : SvcAdminDown SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
Time to RetryReset: never              Retries Left      : 3
Mac Move          : Blockable          Blockable Level   : Tertiary
===============================================================================
*A:ALU-48>show>service>id#
```

Table 79: Show Service Service-ID (MAC Move) Output Fields

Label	Description
Service Id	The service identifier
Mac Move	The administrative state of the MAC movement feature associated with this service
Primary Factor	A factor for the primary ports defining how many MAC relearn periods should be used to measure the MAC relearn rate
Secondary Factor	A factor for the secondary ports defining how many MAC relearn periods should be used to measure the MAC relearn rate
Mac Move Rate	The maximum rate at which MACs can be relearned 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 MACs. The rate is computed as the maximum number of relearns allowed in a 5-s interval: for example, the default rate of 2 relearns per second corresponds to 10 relearns in a 5-s period.
Mac Move Timeout	The time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled. A value of 0 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	The number of times retries are performed for re-enabling the SAP or SDP
SAP Mac Move Information:	
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, PortOperDown, L2OperDown.

Table 79: Show Service Service-ID (MAC Move) Output Fields (Continued)

Label	Description
Time to RetryReset	The time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled. A value of 0 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.
Retries Left	The number of remaining attempts to re-enable the SAP
Mac Move	Specifies whether MAC move is configured as blockable or not blockable on the SAP
Blockable Level	Specifies the level at which MAC move is blockable on the SAP (primary, secondary, or tertiary)
SDP Mac Move Information:	
Admin State	The administrative state of the SDP
Oper State	The operational state of the SDP
Flags	Specifies the conditions that affect the operating status of this SDP. Display output includes: SvcAdminDown, SdpOperDown, NoIngVCLabel, NoEgrVCLabel, PathMTUTooSmall.
Time to RetryReset	The time, in seconds, to wait before a SDP that has been disabled after exceeding the maximum relearn rate is re-enabled. A value of 0 indicates that the SDP will not be automatically re-enabled after being disabled. If after the SDP is re-enabled it is disabled again, the effective retry timeout is doubled in order to avoid thrashing.
Retries Left	The number of remaining attempts to re-enable the SDP
Mac Move	Specifies whether MAC move is configured as blockable or not blockable on the SDP
Blockable Level	Specifies the level at which MAC move is blockable on the SDP (primary, secondary, or tertiary)

sap

Syntax **sap**
sap *sap-id* [**atm** | **base** | **detail** | **qos** | **sap-stats** | **stats**]

Context show>service>id

Description This command displays information about SAPs. When the **sap** command is used without specifying a *sap-id*, the display shows all the information for all SAPs in the service. Including the *sap-id* and a filtering keyword with the **sap** command displays information pertaining to the keyword.

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

atm — filters the command results to display ATM information for the specified SAP

base — filters the command results to display base information for the specified SAP

detail — displays detail information for the specified SAP

qos — filters the command results to display QoS information for the specified SAP

sap-stats — filters the command results to display SAP statistics information for the specified SAP

stats — filters the command results to display statistics information for the specified SAP

Output The following output is an example of service SAP information, and Table 80 describes the fields.

Sample Output

```
*A:ALU-48>show>service>id# sap
===============================================================================
SAP(Summary), Service 5001
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/2/4:1/100                     5001       1     none    1     none   Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48>show>service>id#
```

```
*A:ALU-48>show>service>id# sap 1/2/4:1/100 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 5001
SAP                : 1/2/4:1/100             Encap             : atm
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown L2OperDown
Multi Svc Site     : None
Last Status Change : 10/26/2010 20:13:11
Last Mgmt Change   : 10/26/2010 20:13:12
Sub Type           : regular
Split Horizon Group: shg5001

Max Nbr of MAC Addr: No Limit                Total MAC Addr    : 0
Learned MAC Addr   : 0                       Static MAC Addr   : 0
Admin MTU          : 1524                    Oper MTU          : 1524
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                    qinq-pbit-marking : n/a
Ing Scheduler Mode : 16-priority             Egr Scheduler Mode: 4-priority
Ing Agg Rate Limit : max                     Egr Agg Rate Limit: max
Ing Agg cir        : 10
Q Frame-Based Acct : Disabled
ARP Reply Agent    : Disabled                Host Conn Verify  : Disabled
Mac Learning       : Enabled                 Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                 Mac Pinning       : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None

Acct. Pol          : None                    Collect Stats     : Disabled

Anti Spoofing      : None                    Avl Static Hosts  : 0
                                             Tot Static Hosts  : 0
Calling-Station-Id : n/a
Application Profile: None

MCAC Policy Name   :                         MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                       MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                       MCAC Avail Opnl BW: unlimited
Restr MacProt Src  : Disabled                Restr MacUnpr Dst : Disabled
Time to RetryReset : never                   Retries Left      : 3
Mac Move           : Blockable               Blockable Level   : Tertiary
Egr MCast Grp      :
Auth Policy        : none

PPPoE Circuit-Id     : none

-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Down                    Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                     Port State        : Unknown
```

```
Port Number        : N/A                       Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : N/A
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                       Designated Port   : N/A

Forward transitions: 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                         MST BPDUs tx      : 0

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                         Egress qos-policy : 1
Shared Q plcy      : n/a                       Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
Description        : (Not Specified)
Admin State        : Down                      Lease Populate    : 0
DHCP Snooping      : Down                      Action            : Keep

Proxy Admin State  : Down
Proxy Lease Time   : N/A
Emul. Server Addr  : Not Configured
-------------------------------------------------------------------------------
Subscriber Management
-------------------------------------------------------------------------------
Admin State        : Down                      MAC DA Hashing    : False
Def Sub-Id         : None
Def Sub-Profile    : None
Def SLA-Profile    : None
Sub-Ident-Policy   : None

Subscriber Limit   : 1
Single-Sub-Parameters
 Prof Traffic Only : False
 Non-Sub-Traffic   : N/A
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                          Packets                  Octets
Forwarding Engine Stats
Dropped                : 0                        n/a
Off. HiPrio            : 0                        n/a
Off. LowPrio           : n/a                      n/a

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio            : 0                        n/a
Dro. LowPrio           : n/a                      n/a
For. InProf            : 0                        0
```

```
For. OutProf          : 0                       0

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

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

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : n/a                     n/a
Dro. InProf           : 0                       n/a
Dro. OutProf          : n/a                     n/a
-------------------------------------------------------------------------------
ATM SAP Configuration Information
-------------------------------------------------------------------------------
Ingress TD Profile : 1                        Egress TD Profile : 1
Alarm Cell Handling: Enabled                  AAL-5 Encap       : aal5snap-brid*
OAM Termination    : Enabled                  Periodic Loopback : Disabled
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:ALU-48>show>service>id#


*A:ALU-48>show>service>id# sap 1/2/4:1/100 qos
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 5001
SAP                : 1/2/4:1/100              Encap             : atm
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown L2OperDown
Multi Svc Site     : None
Last Status Change : 10/26/2010 20:13:11
Last Mgmt Change   : 10/26/2010 20:13:12

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
===============================================================================
*A:ALU-48>show>service>id#
```

Table 80: Show Service ID (SAP) Output Fields

Label	Description
Service Id	The service identifier
SAP	The SAP identifier
Encap	The encapsulation type of the SAP
Admin State	The administrative state of the SAP
Oper State	The operating state of the SAP
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
Last Status Change	The time of the most recent operating status change to this SAP
Last Mgmt Change	The time of the most recent management-initiated change to this SAP
Sub Type	The supported sub type: regular
Split Horizon Group	Indicates the split horizon group that this SAP is a member of
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented
Ingr IP Fltr-Id	The IPv4 ingress filter policy ID assigned to the SAP
Egr IP Fltr-Id	The IPv4 egress filter policy ID assigned to the SAP
Ingr Mac Fltr-Id	Not applicable
Egr Mac Fltr-Id	Not applicable

Table 80: Show Service ID (SAP) Output Fields (Continued)

Label	Description
Ingr IPv6 Fltr-Id	The IPv6 ingress filter policy ID assigned to the SAP
Egr IPv6 Fltr-Id	The IPv6 egress filter policy ID assigned to the SAP
tod-suite	Indicates whether a time-based policy is applied to a multiservice site
qinq-pbit-marking	Indicates the qinq P-bit marking for the service: both or top
Ing Scheduler Mode	The scheduler mode for the SAP in the access ingress direction: 4-priority or 16-priority
Egr Scheduler Mode	The scheduler mode for the SAP in the access egress direction: 4-priority or 16-priority
Ing Agg Rate Limit	The PIR rate limit in the access ingress direction for the aggregate of the SAP queues
Egr Agg Rate Limit	The PIR rate limit in the access egress direction for the aggregate of the SAP queues
Ing Agg cir	The CIR rate limit in the access ingress direction for the aggregate of the SAP queues
Egr Agg cir	The CIR rate limit in the access egress direction for the aggregate of the SAP queues
ARP Reply Agent	Specifies whether or not the ARP reply agent is enabled
Host Conn Verify	Specifies the state of verifying host connectivity (enabled or disabled)
Mac Learning	Specifies whether MAC learning is enabled
Discard Unkwn Srce	Specifies whether frames received with an unknown destination MAC are discarded
Mac Aging	Specifies whether MAC aging is enabled
Mac Pinning	Specifies whether MAC pinning is enabled
Acct. Pol	The accounting policy ID applied to the SAP
Collect Stats	Specifies whether accounting statistics are collected on the SAP

Table 80: Show Service ID (SAP) Output Fields (Continued)

Label	Description
Time to RetryReset	The time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled A value of 0 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.
Retries Left	The number of remaining attempts to re-enable the SAP
Mac Move	Indicates the administrative state of the MAC movement feature associated with the SAP
Blockable Level	Specifies the level at which MAC move is blockable on the SAP (primary, secondary, or tertiary)
PPPoE Circuit-Id	Specifies the agent-circuit-id, as specified in RFC 4679 (applies to ATM VPLS instances only)
Stp Service Access Point specifics	
The fields under STP SAP specifics do not apply to VPLS services on the 7705 SAR.	
QOS	
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP
Egress qos-policy	The egress QoS policy ID assigned to the SAP
Shared Q plcy	Not applicable
Multipoint shared	Not applicable
DHCP	
Admin State	Specifies whether DHCP Relay is enabled on this SAP
DHCP Snooping	The status of the DHCP snooping function (up or down)
Action	The DHCP action to be used for the Relay Agent Information Option (Option 82) processing (forward, drop, or keep)
Subscriber Management	
The fields under subscriber management do not apply to VPLS services on the 7705 SAR.	
Sap Statistics	
Last Cleared Time	The date and time that a clear command was issued on statistics

Table 80: Show Service ID (SAP) Output Fields (Continued)

Label	Description
Forwarding Engine Stats	
Dropped	The number of packets or octets dropped by the forwarding engine
Off. HiPrio	The number of high-priority packets or octets offered to the forwarding engine
Off. LowPrio	The number of low-priority packets offered to the forwarding engine
Queueing Stats (Ingress QoS Policy)	
Dro. HiPrio	The number of high-priority packets or octets discarded, as determined by the SAP ingress QoS policy
Dro. LowPrio	The number of low-priority packets discarded, as determined by the SAP ingress QoS policy
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP ingress QoS policy
For. OutProf	The number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP ingress QoS policy
Queueing Stats (Egress QoS Policy)	
Dro. InProf	The number of in-profile packets or octets discarded, as determined by the SAP egress QoS policy
Dro. OutProf	The number of out-of-profile packets or octets discarded, as determined by the SAP egress QoS policy
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP egress QoS policy
For. OutProf	The number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP egress QoS policy
Sap per Queue stats	
Ingress Queue *n*	The index of the ingress QoS queue of this SAP, where *n* is the index number
Off. HiPrio	The number of packets or octets of high-priority traffic for the SAP (offered)

Table 80: Show Service ID (SAP) Output Fields (Continued)

Label	Description
Off. LoPrio	The number of packets or octets of low-priority traffic for the SAP (offered)
Dro. HiPrio	The number of high-priority traffic packets or octets dropped
Dro. LoPrio	The number of low-priority traffic packets or octets dropped
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	The number of out-of-profile packets or octets (rate above CIR) forwarded
Egress Queue *n*	The index of the egress QoS queue of the SAP, where *n* is the index number
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	The number of out-of-profile packets or octets (rate above CIR) forwarded
Dro. InProf	The number of in-profile packets or octets dropped for the SAP
Dro. OutProf	The number of out-of-profile packets or octets discarded
ATM SAP Configuration Information	
Ingress TD Profile	The profile ID of the traffic descriptor applied to the ingress SAP
Egress TD Profile	The profile ID of the traffic descriptor applied to the egress SAP
Alarm Cell Handling	Indicates that OAM cells are being processed
OAM Termination	Indicates whether this SAP is an OAM termination point
AAL-5 Encap	Indicates the type of AAL5 encapsulation for this ATM SAP
OAM Termination	Indicates the state of the OAM termination for this ATM SAP (enabled or disabled)
Periodic Loopback	Indicates the state of the periodic loopback for this ATM SAP (enabled or disabled)

pppoe-circuit-id

Syntax **pppoe-circuit-id statistics**

Context show>service>id>sap

Description This command displays statistics related to the agent-circuit-id, as specified in RFC 4679. This command applies to ATM VPLS instances only.

sdp

Syntax **sdp**
sdp {*sdp-id*[**:***vc-id*] | **far-end** *ip-address*} [**detail**]

Context show>service>id

Description This command displays information about SDPs. When the **sdp** command is used without specifying a *sdp-id*, the display shows all the information for all SDPs in the service. Including the *sdp-id* and a filtering keyword with the **sdp** command displays information pertaining to the keyword.

Parameters *sdp-id* — the SDP identifier

Values 1 to 17407

vc-id — the VC identifier

Values 1 to 4294967295

ip-address — displays information for the SDP having this as far-end IP address

Values a.b.c.d

detail — adds details to SDP information

Output The following output is an example of service SDP information, and Table 81 describes the fields.

Sample Output

```
*A:ALU-48>show>service>id# sdp 5001
===============================================================================
Service Destination Point (Sdp Id : 5001)
===============================================================================
SdpId            Type IP address      Adm     Opr       I.Lbl       E.Lbl
-------------------------------------------------------------------------------
5001:100         Spok 10.10.10.10     Up      Down      0           0
===============================================================================
*A:ALU-48>show>service>id#
```

```
*A:ALU-48>show>service>id# sdp 5001 detail
===============================================================================
Service Destination Point (Sdp Id : 5001) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 5001:100  -(10.10.10.10)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id              : 5001:100                  Type              : Spoke
Split Horiz Grp     : shg5001
VC Type             : Ether                     VC Tag            : n/a
Admin Path MTU      : 0                         Oper Path MTU     : 0
Far End             : 10.10.10.10               Delivery          : MPLS

Admin State         : Up                        Oper State        : Down
Acct. Pol           : None                      Collect Stats     : Disabled
Ingress Label       : 0                         Egress Label      : 0
Ing mac Fltr        : n/a                       Egr mac Fltr      : n/a
Ing ip Fltr         : n/a                       Egr ip Fltr       : n/a
Ing ipv6 Fltr       : n/a                       Egr ipv6 Fltr     : n/a
Admin ControlWord   : Not Preferred             Oper ControlWord  : False
Last Status Change  : 10/26/2010 20:14:00       Signaling         : TLDP
Last Mgmt Change    : 10/26/2010 20:14:01       Force Vlan-Vc     : Disabled
Endpoint            : endpoint5000              Precedence        : 4
PW Status Sig       : Enabled
Class Fwding State  : Down
Flags               : SvcAdminDown SdpOperDown
                      NoIngVCLabel NoEgrVCLabel
                      PathMTUTooSmall
Time to RetryReset  : never                     Retries Left      : 3
Mac Move            : Blockable                 Blockable Level   : Tertiary
Peer Pw Bits        : None
Peer Fault Ip       : None
Max Nbr of MAC Addr: No Limit                   Total MAC Addr    : 0
Learned MAC Addr    : 0                         Static MAC Addr   : 0

MAC Learning        : Enabled                   Discard Unkwn Srce: Disabled
BPDU Translation    : Disabled
L2PT Termination    : Disabled
MAC Pinning         : Disabled
Ignore Standby Sig  : False                     Block On Mesh Fail: False

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

Statistics            :
I. Fwd. Pkts.       : 0                         I. Dro. Pkts.     : 0
I. Fwd. Octs.       : 0                         I. Dro. Octs.     : 0
E. Fwd. Pkts.       : 0                         E. Fwd. Octets    : 0
MCAC Policy Name    :
MCAC Max Unconst BW: no limit                   MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                          MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                          MCAC Avail Opnl BW: unlimited

Associated LSP LIST :
No LSPs Associated
```

```
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Down                     Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Discarding
Port Number        : 0                        Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                      Designated Port Id: 0

Fwd Transitions    : 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
===============================================================================

===============================================================================
*A:ALU-48>show>service>id#

*A:ALU-48>show>service>id# sdp far-end 10.10.10.10
===============================================================================
Service Destination Point(Far-End : 10.10.10.10)
===============================================================================
SdpId            Type IP address      Adm     Opr       I.Lbl       E.Lbl
-------------------------------------------------------------------------------
5001:100         Spok 10.10.10.10     Up      Down      0           0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================

*A:ALU-48>show>service>id# sdp far-end 10.10.10.10 detail
===============================================================================
Service Destination Point(Far-End : 10.10.10.10) Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 5001:100  -(10.10.10.10)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id             : 5001:100                 Type              : Spoke
Split Horiz Grp    : shg5001
VC Type            : Ether                    VC Tag            : n/a
Admin Path MTU     : 0                        Oper Path MTU     : 0
Far End            : 10.10.10.10              Delivery          : MPLS

Admin State        : Up                       Oper State        : Down
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 0                        Egress Label      : 0
Ing mac Fltr       : n/a                      Egr mac Fltr      : n/a
Ing ip Fltr        : n/a                      Egr ip Fltr       : n/a
Ing ipv6 Fltr      : n/a                      Egr ipv6 Fltr     : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
Last Status Change : 10/26/2010 20:14:00      Signaling         : TLDP
Last Mgmt Change   : 10/26/2010 20:14:01      Force Vlan-Vc     : Disabled
```

```
Endpoint           : endpoint5000             Precedence        : 4
Class Fwding State : Down
Flags              : SvcAdminDown SdpOperDown
                     NoIngVCLabel NoEgrVCLabel
                     PathMTUTooSmall
Time to RetryReset : never                    Retries Left      : 3
Mac Move           : Blockable                Blockable Level   : Tertiary
Peer Pw Bits       : None
Peer Fault Ip      : None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0

MAC Learning       : Enabled                  Discard Unkwn Srce: Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                    Block On Mesh Fail: False

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

Statistics           :
I. Fwd. Pkts.      : 0                        I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                        I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octets    : 0
E. Dro. Pkts.      : 0
Grp Enc Stats      :
I. Fwd. Pkts.      : 0                        I. Fwd. Octs.     : 0
I. Dro. Inv. Spi.  : 0                        I. Dro. OthEncPkt*: 0
E. Fwd. Pkts.      : 0                        E. Fwd. Octs.     : 0
E. Dro. Enc. Pkts. : 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 :
No LSPs Associated


-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Down                     Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                      Port State        : Discarding
Port Number        : 0                        Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                      Designated Port Id: 0

Fwd Transitions    : 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
```

```
TCN BPDUs rcvd     : 0                          TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                          RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48>show>service>id#
```

Table 81: Show Service ID (SDP) Output Fields

Label	Description
SDP Id	The SDP identifier
Type	Indicates whether this service SDP binding is a spoke or a mesh
Split Horizon Group	The name of the split horizon group
VC Type	The VC type: ether or vlan
VC Tag	The explicit dot1q value used when encapsulating to the SDP far end
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
Far End	Specifies the IP address of the remote end of the GRE, MPLS, or IP tunnel defined by this SDP
Delivery	Specifies the type of delivery used by the SDP: GRE, MPLS, or IP
Admin State	The administrative state of this SDP
Oper State	The operational state of this SDP
Acct. Pol	The accounting policy applied to the SDP
Collect Stats	Specifies whether accounting statistics are collected on the SDP
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP
Ing mac Fltr	Not applicable

Table 81: Show Service ID (SDP) Output Fields (Continued)

Label	Description
Egr mac Fltr	Not applicable
Ing ip Fltr	The SDP ingress filter policy ID for IPv4
Egr ip Fltr	The SDP egress filter policy ID for IPv4
Ing ipv6 Fltr	Not applicable
Egr ipv6 Fltr	Not applicable
Admin ControlWord	The administrative state of the control word: Preferred (control word enabled) or Not Preferred (control word disabled)
Oper ControlWord	The operational state of the control word: True (control word enabled) or False (control word disabled)
Last Status Change	The date and time of the most recent status change to this SDP
Signaling	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP
Last Mgmt Change	The date and time of the most recent management-initiated change to this SDP
Endpoint	The name of the service endpoint
Precedence	Specifies the precedence of this SDP binding when there are multiple SDP bindings attached to one service endpoint
PW Status Sig	Specifies whether pseudowire status signaling for spoke SDPs is enabled or disabled
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SvcAdminDown, SdpOperDown, NoIngVCLabel NoEgrVCLabel, PathMTUTooSmall.
Retries Left	The number of remaining attempts to re-enable the SDP
Mac Move	The administrative state of the MAC movement feature associated with the SDP
Blockable Level	Specifies the level at which MAC move is blockable on the SAP (primary, secondary, or tertiary)
MAC Learning	Specifies whether MAC learning is enabled

Table 81: Show Service ID (SDP) Output Fields (Continued)

Label	Description
MAC Pinning	Specifies whether MAC pinning is enabled in this SDP
Ignore Standby Sig	Specifies whether or not ignore standby signaling is configured True: standby signaling is ignored False: standby signaling is not ignored
Block On Mesh Fail	Specifies whether or not to take down the spoke SDP when the mesh SDP is down True: the spoke SDP is not taken down False: the spoke SDP is taken down
KeepAlive Information	
Admin State	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	The length of the SDP echo request messages transmitted on this SDP
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 time to wait before the keepalive operating status is eligible to enter the alive state
Statistics	
I. Fwd. Pkts.	The number of forwarded ingress packets
I. Dro. Pkts.	The number of dropped ingress packets
I. Fwd. Octs.	The number of forwarded ingress octets
I. Dro. Octs.	The number of dropped ingress octets
E. Fwd. Pkts.	The number of forwarded egress packets
E. Fwd. Octets	The number of forwarded egress octets
E. Dro. Pkts.	The number of dropped egress octets
Grp Enc Stats	
I. Fwd. Pkts.	The number of forwarded ingress packets

Table 81: Show Service ID (SDP) Output Fields (Continued)

Label	Description
I. Fwd. Octs.	The number of forwarded ingress octets
I. Dro. Inv. Spi.	The number of ingress packets dropped due to an invalid SPI
I. Dro. OthEncPkt*	The number of ingress packets dropped due to a packet encapsulation other than that configured
E. Fwd. Pkts.	The number of forwarded egress packets
E. Fwd. Octs.	The number of forwarded egress octets
E. Dro. Enc. Pkts	The number of dropped egress packets
Associated LSP LIST	If the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far-end field. If the SDP type is GRE, the following message displays: SDP delivery mechanism is not MPLS
Number of SDPs	The total number of SDPs applied to this service ID

split-horizon-group

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

Context show>service>id

Description This command displays information related to all split horizon groups in the service or the specified service split horizon group.

Parameters *group-name* — specifies a split horizon group name

Values 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 (“ ”).

Output The following output is an example of service split horizon group information, and Table 82 describes the fields.

Sample Output

```
*A:ALU-48>show>service>id# split-horizon-group
===============================================================================
Service: Split Horizon Group
===============================================================================
Name                                Description
-------------------------------------------------------------------------------
R   shg5001
-------------------------------------------------------------------------------
R = Residential Split Horizon Group
A = Auto Created Split Horizon Group
No. of Split Horizon Groups: 1

*A:ALU-48>show>service>id# split-horizon-group shg5001
===============================================================================
Service: Split Horizon Group
===============================================================================
Name                                Description
-------------------------------------------------------------------------------
R   shg5001
-------------------------------------------------------------------------------
Associations
-------------------------------------------------------------------------------
SAP                                 1/2/4:1/100
SDP                                 5001:100
-------------------------------------------------------------------------------
R = Residential Split Horizon Group
SAPs Associated : 1                 SDPs Associated : 1
===============================================================================
*A:ALU-48>show>service>id#
```

Table 82: Show Service ID (Split Horizon Group) Output Fields

Label	Description
Name	The name of the split horizon group. When preceded by "R", the group is a residential split horizon group.
Description	A description of the split horizon group as configured by the user
Associations	A list of SAPs and SDPs associated with the split horizon group

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 The following output is an example of information about service-id STP, and Table 83 describes the fields.

Sample Output

```
A:ALU-48#> show service id 5000 stp
===============================================================================
Stp info, Service 5000
===============================================================================
Bridge Id          : 80:00.a4:58:ff:00:00:00  Top. Change Count : 0
Root Bridge        : N/A                      Stp Oper State    : Down
Primary Bridge     : N/A                      Topology Change   : Inactive
Mode               : Rstp                     Last Top. Change  : 0d 00:00:00
Vcp Active Prot.   : N/A
Root Port          : N/A                      External RPC      : 0

===============================================================================
Stp port info
===============================================================================
Sap/Sdp/PIP Id     Oper-     Port-      Port-       Port-  Oper-  Link-  Active
                   State     Role       State       Num    Edge   Type   Prot.
-------------------------------------------------------------------------------
1/5/5              Down      N/A        Disabled    2048   N/A    Pt-pt  N/A
15:5001            Down      N/A        Discard     0      N/A    Pt-pt  N/A
===============================================================================
A:ALU-48#
```

```
A:ALU-48#> show service id 5000 stp detail
===============================================================================
Spanning Tree Information
===============================================================================
-------------------------------------------------------------------------------
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
VPLS oper state    : Down                     Core Connectivity : Down
Stp Admin State    : Down                     Stp Oper State    : Down
Mode               : Rstp                     Vcp Active Prot.  : N/A

Bridge Id          : 80:00.a4:58:ff:00:00:00  Bridge Instance Id: 0
Bridge Priority    : 32768                    Tx Hold Count     : 6
Topology Change    : Inactive                 Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00              Bridge Max Age    : 20
Top. Change Count  : 0                        Bridge Fwd Delay  : 15
MST region revision: 0                        Bridge max hops   : 20
MST region name    :

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                        Root Forward Delay: 15
Rcvd Hello Time    : 2                        Root Max Age      : 20
Root Priority      : 32768                    Root Port         : N/A

-------------------------------------------------------------------------------
Spanning Tree Sap/Spoke SDP Specifics
-------------------------------------------------------------------------------
SAP Identifier     : 1/5/5                    Stp Admin State   : Up
Port Role          : N/A                      Port State        : Unknown
Port Number        : 2048                     Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                      Designated Port   : N/A
Forward transitions: 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                        MST BPDUs tx      : 0

SDP Identifier     : 15:5001                  Stp Admin State   : Down
Port Role          : N/A                      Port State        : Discarding
Port Number        : 0                        Port Priority     : 128
Port Path Cost     : 10                       Auto Edge         : Enabled
Admin Edge         : Disabled                 Oper Edge         : N/A
Link Type          : Pt-pt                    BPDU Encap        : Dot1d
Root Guard         : Disabled                 Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                      Designated Port Id: 0
Fwd Transitions    : 0                        Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                        Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                        TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                        RST BPDUs tx      : 0

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

Table 83: Show Service ID (STP) Output Fields

Label	Description
Stp info, Service #	
Bridge Id	The MAC address used to identify this bridge in the network
Top Change Count	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	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.
Topology Change	Specifies whether a topology change is currently in progress
Mode	Always RSTP
Last Top. Change	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
Root Port	The port number of the port that offers the lowest cost path from this bridge to the root bridge
Stp port info	
Sap/Sdp/PIP Id	The ID of the access port where the SAP or spoke SDP is defined
Oper-State	Indicates the operational state of the Rapid Spanning Tree Protocol instance associated with this service. This field is applicable only when STP is enabled on the router.
Port-State	The port identifier of the port on the designated bridge for this port's segment
Port-Num	The value of the port number field that is contained in the least significant 12 bits of the 16-bit port ID associated with the SAP or spoke SDP

Table 83: Show Service ID (STP) Output Fields (Continued)

Label	Description
Oper-Edge	The state of the oper-edge variable: true or false
Link-Type	The link type (number of bridges that can exist behind the SAP or spoke SDP): pt-pt or shared
VPLS Spanning Tree Information	
VPLS oper state	The operational state of the Rapid Spanning Tree Protocol instance associated with this service. This field is applicable only when STP is enabled on the router.
Core Connectivity	The connectivity status to the core
Stp Admin State	The administrative state of the Rapid Spanning Tree Protocol instance associated with this service
Stp Oper State	The operational state of the Rapid Spanning Tree Protocol instance associated with this service. This field is applicable only when STP is enabled on the router.
Mode	Always RSTP
Bridge Id	The MAC address used to identify this bridge in the network
Bridge Priority	The priority of the Spanning Tree Protocol instance associated with this service
Tx Hold Count	The interval length during which no more than two Configuration BPDUs shall be transmitted by this bridge
Topology Change	Specifies whether a topology change is currently in progress
Bridge Hello Time	The amount of time between the transmission of Configuration BPDUs
Last Top. Change	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
Bridge Max Age	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.

Table 83: Show Service ID (STP) Output Fields (Continued)

Label	Description
Top. Change Count	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
Bridge Fwd Delay	Specifies how fast a bridge changes its state when moving toward the forwarding state
MST region revision	Not applicable
MST region name	Not applicable
Root Bridge	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	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	The amount of time between the transmission of configuration BPDUs
Root Max Age	The maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded
Root Priority	The priority of the bridge that is currently selected as root-bridge for the network
Root Port	The port number of the port that offers the lowest cost path from this bridge to the root bridge
Spanning Tree Sap/Spoke SDP Specifics	
SAP Identifier	The ID of the access port where this SAP is defined
SDP Identifier	The ID of the SDP and VC
Stp Admin State	The administrative state of the Rapid Spanning Tree Protocol instance associated with this service

Table 83: Show Service ID (STP) Output Fields (Continued)

Label	Description
Port State	The port identifier of the port on the designated bridge for this port's segment
Port number	The value of the port number field that is contained in the least significant 12 bits of the 16-bit port ID associated with the SAP or spoke SDP
Port Priority	The value of the port priority field that is contained in the most significant 4 bits of the 16-bit port ID associated with the SAP or spoke SDP
Port Path Cost	The contribution of this port to the path cost of paths towards the spanning tree root which include this port
Auto Edge	Specifies whether auto edge is enabled or disabled
Oper Edge	The state of the oper-edge variable: true or false
Link Type	The link type (number of bridges that can exist behind the SAP or spoke SDP): pt-pt or shared
BPDU Encap	The type of encapsulation used on BPDUs sent out and received on this SAP or spoke SDP
Root Guard	Specifies whether the port is allowed to become an STP root port
CIST Desig Bridge	The bridge identifier of the bridge which this port considers to be the designated bridge for this port's segment
Designated Port	The port identifier of the port on the designated bridge for this port's segment

ingress-label

Syntax **ingress-label** *start-label* [*end-label*]

Context show>service

Description This command displays services using a 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 are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *start-label* — the starting ingress label value for which to display services using the label range. If only *start-label* is specified, services only using *start-label* are displayed.

Values 0, 2048 to 131071

end-label — the ending ingress label value for which to display services using the label range

Values 2049 to 131071

Default the *start-label* value

Output The following output is an example of information about services using the specified range of ingress labels, and Table 84 describes the fields.

Sample Output

```
*A:ALU-48>show>service# ingress-label 0
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Binding        Type  I.Lbl                E.Lbl
-------------------------------------------------------------------------------
3          15:15              Spok  0                    0
5          5:5                Spok  0                    0
6          5:6                Spok  0                    0
5000       15:5000            Mesh  0                    0
5000       15:5001            Spok  0                    0
5001       5001:100           Spok  0                    0
-------------------------------------------------------------------------------
Number of Bindings Found : 6
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48>show>service#
```

Table 84: Show Service Ingress-Label Output Fields

Label	Description
Svc ID	The service identifier
SDP Binding	The SDP binding identifier
Type	Indicates whether the SDP is spoke or mesh
I.Lbl	The ingress label used by the far-end device to send packets to this device in this service by the SDP
E.Lbl	The egress label used by this device to send packets to the far-end device in this service by the SDP
Number of Bindings Found	The number of SDP bindings within the label range specified

sap-using

Syntax **sap-using** [**sap** *sap-id*]
sap-using interface [*ip-address* | *ip-int-name*]
sap-using description
sap-using authentication-policy *auth-plcy-name*
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**ingress** | **egress**] **atm-td-profile** *td-profile-id*
sap-using [**ingress** | **egress**] **qos-policy** [*qos-policy-id* | *qos-policy-name*]
sap-using [**ingress** | **egress**] **scheduler-mode** {**4-priority** | **16-priority**}

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 **description —** displays a SAP summary table with description information

ingress — specifies matching an ingress policy

egress — specifies matching an egress policy

qos-policy-id **—** the ingress or egress QoS policy ID for which to display matching SAPs

Values 1 to 65535

qos-policy-name **—** the ingress or egress QoS policy name for which to display matching SAPs

Values up to 64 characters

td-profile-id **—** displays SAPs using this traffic descriptor

filter-id — the filter policy for which to display matching SAPs specifies. The filter ID or filter name must already exist.

Values 1 to 65535 or *filter-name* (up to 64 characters)

auth-plcy-name — the session authentication policy for which to display matching SAPs

sap-id — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

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 to 223.255.255.255

ip-int-name — the IP interface name for which to display matching SAPs

scheduler-mode — specifies the scheduler mode for which to display the SAPs

Output The following output is an example of information about SAPs matching the specified properties, and Table 85 describes the fields.

Sample Output

```
*A:ALU-48>show>service# sap-using
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.   Ing.    Egr.   Egr.    Adm   Opr
                                           QoS    Fltr    QoS    Fltr
-------------------------------------------------------------------------------
1/5/1                           2          1      none    1      none    Up    Down
1/5/2                           2          1      none    1      none    Up    Down
1/5/5                           5000       1      mac     1      ip4     Up    Down
1/2/4:1/100                     5001       1      none    1      none    Up    Down
-------------------------------------------------------------------------------
Number of SAPs : 4
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
Sap Aggregation Groups
===============================================================================
GroupName                       SvcId      Ing.   Ing.    Egr.   Egr.    Adm   Opr
                                           QoS    Fltr    QoS    Fltr
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:ALU-48>show>service#
```

```
*A:ALU-48# show service sap-using description
===============================================================================
Service Access Points
===============================================================================
PortId                              SvcId        Adm  Opr  Description
-------------------------------------------------------------------------------
1/1/2                               1            Down Down (Not Specified)
1/2/1.1                             4            Up   Down (Not Specified)
1/10/4                              5            Up   Down (Not Specified)
-------------------------------------------------------------------------------
Number of SAPs : 3
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48#

*A:7705custDoc:Sar18>show>service# sap-using ingress scheduler-mode 4-priority
===========================================================================
Service Access Points Using Ingress 4-priority Scheduler Mode
===========================================================================
PortId                          SvcId      Scheduler Mode    Adm  Opr
---------------------------------------------------------------------------
1/12/6                          6000       4-priority        Up   Down
---------------------------------------------------------------------------
Number of SAPs : 1
---------------------------------------------------------------------------
===========================================================================
*A:7705custDoc:Sar18>show>service#
```

Table 85: Show Service SAP-Using Output Fields

Label	Description
Port ID	The ID of the access port where the SAP is defined
SvcID	The service identifier
Ing.QoS	The SAP ingress QoS policy number specified on the ingress SAP
Ing. Fltr	The filter policy ID applied to the ingress SAP
Egr.QoS	The SAP egress QoS policy number specified on the egress SAP
Egr. Fltr	The filter policy ID applied to the egress SAP
Scheduler Mode	The scheduler mode of the SAP: 4-priority or 16-priority
Adm	The administrative state of the SAP
Opr	The actual state of the SAP
Description	The description of the SAP
Number of SAPs	The total number of SAPs listed in the output
GroupName	The group name of any aggregation groups

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 1 to 2147483647 or *service-name*

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id*[*:vc-id*] | **interface** [*ip-address* | *ip-int-name*]]

Context clear>service>id>dhcp

Description This command clears DHCP statistics for this service.

Parameters *sap-id* — clears the specified SAP statistics. See Table 39 for a full list of SAP IDs.

sdp-id — the SDP ID to be cleared

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID to be cleared

Values 1 to 4294967295

ip-int-name — clears the statistics for the IP interface with the specified name

ip-addr — clears the statistics for the IP interface with the specified IP address

fdb

Syntax **fdb** {**all** | **mac** *ieee-address* | **sap** *sap-id* | **mesh-sdp** *sdp-id*[*:vc-id*] | **spoke-sdp** *sdp-id:vcid*}

Context clear>service>id

Description This command clears FDB entries for the service.

Parameters **all —** clears all FDB entries

ieee-address **—** clears only FDB entries in the FDB table with the specified 48-bit 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 Table 39 for a full list of SAP IDs.

mesh-sdp — clears only service FDB entries associated with the specified mesh SDP ID. For a mesh SDP, the VC ID is optional.

spoke-sdp — clears only service FDB entries associated with the specified spoke SDP ID. For a spoke SDP, the VC ID must be specified.

sdp-id — the SDP ID for which to clear associated FDB entries

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID for which to clear associated FDB entries

Values 1 to 4294967295

igmp-snooping

Syntax **igmp-snooping**

Context clear>service>id

Description This command enables the context to clear IGMP snooping-related data.

mld-snooping

Syntax **mld-snooping**

Context clear>service>id

Description This command enables the context to clear MLD snooping-related data.

port-db

Syntax **port-db sap** *sap-id* [**group** *grp-address* [**source** *src-ip-address*]]
port-db sap *sap-id* [**group** *grp-ipv6-address*]
port-db sap *sap-id* **group** *grp-ipv6-address* **source** *src-ipv6-address*
port-db sdp *sdp-id:vc-id* [**group** *grp-address* [**source** *src-ip-address*]]
port-db sdp *sdp-id:vc-id* [**group** *grp-ipv6-address*]
port-db sdp *sdp-id:vc-id* **group** *grp-ipv6-address* **source** *src-ipv6-address*

Context clear>service>id>igmp-snooping
clear>service>id>mld-snooping

Description This command clears the information on the IGMP or MLD snooping port database for the VPLS service.

Parameters *sap-id* — clears IGMP or MLD snooping statistics matching the specified SAP ID. See Table 39 for a full list of SAP IDs.

sdp-id — clears only IGMP or MLD snooping entries associated with the specified SDP. For a spoke SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID for which to clear information

Values 1 to 4294967295

Default For mesh SDPs only, all VC IDs

grp-address — clears IGMP snooping statistics matching the specified group address

Values a.b.c.d

grp-ipv6-address — clears MLD snooping statistics matching the specified group address

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

src-ip-address — clears IGMP snooping statistics matching the specified source address

Values a.b.c.d

src-ipv6-address — clears MLD snooping statistics matching the specified source address

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

querier

Syntax **querier**

Context clear>service>id>igmp-snooping
clear>service>id>mld-snooping

Description This command clears the information on the IGMP or MLD snooping queriers for the VPLS service.

statistics

Syntax **statistics** {**all** | **sap** *sap-id* | **sdp** *sdp-id:vc-id*}

Context clear>service>id>igmp-snooping
clear>service>id>mld-snooping

Description This command clears IGMP or MLD snooping statistics.

Parameters **all** — clears all statistics for the service ID

sap-id — clears statistics for the specified SAP. See Table 39 for a full list of SAP IDs.

sdp-id:vc-id — clears statistics for the specified SDP

Values *sdp-id* : 1 to 17407
vc-id : 1 to 4294967295

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[*:vc-id*] **ingress-vc-label**

Context clear>service>id

Description This command clears and resets the mesh SDP bindings for the service.

Parameters *sdp-id* — the mesh SDP ID to be reset

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

Default All VC IDs on the SDP ID

pppoe-circuit-id

Syntax **pppoe-circuit-id statistics**

Context clear>service>id>sap

Description This command clears and resets the SAP statistics for the service.

spoke-sdp

Syntax **spoke-sdp** *sdp-id:vc-id* **ingress-vc-label**

Context clear>service>id

Description This command clears and resets the spoke SDP bindings for the service.

Parameters *sdp-id* — the spoke SDP ID to be reset

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

cem

Syntax **cem**

Context clear>service>statistics>id

Description This command clears CEM statistics for this service.

counters

Syntax **counters**

Context clear>service>statistics>id

Description This command clears all traffic queue counters associated with the service ID.

spoke-sdp

Syntax **spoke-sdp** *sdp-id:vc-id* {**all** | **counters** | **l2pt** | **mrp**}

Context clear>service>statistics>id

Description This command clears statistics for the spoke SDP bound to the service.

Parameters *sdp-id* — the spoke SDP ID for which to clear statistics

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

all — clears all queue statistics associated with the SDP

counters — clears all queue counters associated with the SDP

l2pt — clears all L2PT statistics associated with the SDP

mrp — clears all MRP statistics associated with the SDP

stp

Syntax **stp**

Context clear>service>statistics>id

Description This command clears all spanning tree statistics for the service ID.

sap

Syntax **sap** *sap-id* {**all** | **cem** | **counters** | **l2pt** | **stp** | **mrp**}

Context clear>service>statistics

Description This command clears statistics for the SAP bound to the service.

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

all — clears all queue statistics associated with the SAP

cem — clears all CEM statistics associated with the SAP

counters — clears all queue counters associated with the SAP

l2pt — clears all L2PT statistics associated with the SAP

stp — clears all STP statistics associated with the SAP

mrp — clears all MRP statistics associated with the SAP

sdp

Syntax **sdp** *sdp-id* **keep-alive**

Context clear>service>statistics

Description This command clears keepalive statistics associated with the SDP ID.

Parameters *sdp-id* — the SDP ID for which to clear statistics

Values 1 to 17407

keep-alive — clears the keep alive history associated with this SDP ID

dhcp

Syntax **dhcp**

Context clear>router

Description This command enables the context to clear and reset DHCP entities.

statistics

Syntax **statistics** [**interface** *ip-int-name* | *ip-address*]

Context clear>router>dhcp

Description This command clears DHCP statistics.

Parameters *ip-int-name* — clears the statistics for the IP interface with the specified name

ip-addr — clears the statistics for the IP interface with the specified IP address

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

igmp-snooping

Syntax [**no**] **igmp-snooping**

Context debug>service>id

Description This command enables IGMP snooping debugging.

The **no** form of the command disables IGMP snooping debugging.

mld-snooping

Syntax [**no**] **mld-snooping**

Context debug>service>id

Description This command enables MLD snooping debugging.

The **no** form of the command disables MLD snooping debugging.

detail-level

Syntax **detail-level** {**low** | **medium** | **high**}
no detail-level

Context debug>service>id>igmp-snooping
debug>service>id>mld-snooping

Description This command enables debugging for the IGMP or MLD tracing detail level.

The **no** form of the command disables debugging for the IGMP or MLD tracing detail level.

mac

Syntax [**no**] **mac** *ieee-address*

Context debug>service>id>igmp-snooping
debug>service>id>mld-snooping

Description This command debugs IGMP or MLD packets for the specified MAC address.

The **no** form of the command disables the MAC debugging.

Parameters *ieee-address* — 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

mode

Syntax **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
no mode

Context debug>service>id>igmp-snooping
debug>service>id>mld-snooping

Description This command enables debugging for the IGMP or MLD tracing mode.

The **no** form of the command disables debugging for the IGMP or MLD tracing mode.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>igmp-snooping
debug>service>id>mld-snooping

Description This command enables debugging for IGMP or MLD packets for a specific SAP.

The **no** form of the command disables debugging for the SAP.

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>id>igmp-snooping
debug>service>id>mld-snooping

Description This command enables debugging for a specific SDP.

The **no** form of the command disables debugging for the SDP.

Parameters *sdp-id* — specifies the mesh SDP or spoke SDP. For an IGMP spoke SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID for which to display information

Values 1 to 4294967295

event-type

Syntax **[no] event-type {config-change | svc-oper-status-change | sap-oper-status-change | sdpbind-oper-status-change}**

Context debug>service>id

Description This command enables a particular debugging event type.

The **no** form of the command disables the event type debugging.

Parameters **config-change —** debugs configuration change events

svc-oper-status-change — debugs service operational status changes

sap-oper-status-change — debugs SAP operational status changes

sdpbind-oper-status-change — debugs SDP operational status changes

sap

Syntax **[no] sap** *sap-id*

Context debug>service>id

Description This command enables debugging for a particular SAP.

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

stp

Syntax **stp**

Context debug>service>id

Description This command enables the context for debugging STP.

all-events

Syntax **all-events**

Context debug>service>id>stp

Description This command enables STP debugging for all events.

bpdu

Syntax [**no**] **bpdu**

Context debug>service>id>stp

Description This command enables STP debugging for received and transmitted BPDUs.

core-connectivity

Syntax [**no**] **core-connectivity**

Context debug>service>id>stp

Description This command enables STP debugging for core connectivity.

exception

Syntax [**no**] **exception**

Context debug>service>id>stp

Description This command enables STP debugging for exceptions.

fsm-state-changes

Syntax [**no**] **fsm-state-changes**

Context debug>service>id>stp

Description This command enables STP debugging for FSM state changes.

fsm-timers

Syntax [**no**] **fsm-timers**

Context debug>service>id>stp

Description This command enables STP debugging for FSM timer changes.

port-role

Syntax [**no**] **port-role**

Context debug>service>id>stp

Description This command enables STP debugging for changes in port roles.

port-state

Syntax [**no**] **port-state**

Context debug>service>id>stp

Description This command enables STP debugging for port states.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>stp

Description This command enables STP debugging for a specific SAP.

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

Internet Enhanced Service

In This Chapter

This chapter provides information about Internet Enhanced Service (IES), used to provide IP routing services; that is, direct forwarding of IP traffic between CE devices, and also to facilitate the transport of in-band management datagrams of the 7705 SAR over ATM links.

Topics in this chapter include:

- IES for In-band Management
- IES for Customer Traffic
- Configuring IES with CLI
- IES Command Reference

Internet Enhanced Services can coexist with IES management SAP services on the same 7705 SAR node. IP over ATM is used exclusively for in-band management of the 7705 SAR. Up to two IPoATM SAPs can be bound to IES along with many other SAPs with other (non-ATM) supported SAP encapsulation types. Traffic from IPoATM SAPs is extracted to the CSM for further processing. Traffic received from other IES SAPs is forwarded as per the forwarding table (FIB).

IES for In-band Management

Topics in this section include:

- Setting Up Connections Between the 5620 SAM and the 7705 SAR
- Encapsulation
- Layer 2 and Layer 3 Traffic Management
- Troubleshooting and Fault Detection Services
- IP ECMP Load Balancing

In the HSDPA offload application (see HSDPA Offload), the main uplink out of a typical cell site is over the ATM network using leased lines. Mission-critical traffic such as voice, signaling, and synchronization traffic is carried over the ATM network.

Internet Enhanced Service (IES) provides a reliable means of diverting the node management IP packets from the DSL IP network to the more reliable Layer 2 ATM network. To do this, IES provides an IP address and interworking function between the Layer 3 IP network and the Layer 2 ATM network. Without this capability, the in-band IP management traffic for the 7705 SAR could only be connected to an IP network.

IES can be used for in-band management of the 7705 SAR over the ATM network. IP over an ATM SAP bound to IES is for in-band management purposes only, and IP traffic from the ATM SAP is only extracted to the CSM; it is not forwarded.

IES management service is supported on the following cards for the 7705 SAR-8 and 7705 SAR-18:

- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card

IES management service is also supported on the T1/E1 ports on the following:

- 7705 SAR-F
- 7705 SAR-M (variants with T1/E1 ports)
- 7705 SAR-A (variants with T1/E1 ports)
- 7705 SAR-X
- 4-port T1/E1 and RS-232 Combination module

The service can be created on an ATM port or on an IMA group.

In the 7705 SAR, all traffic received over IES management SAPs is extracted directly to the control plane (CSM) in the same way as management traffic received over the CSM console port or Ethernet management port, or management traffic destined for the 7705 SAR over an Ethernet or MLPPP encapsulated network port. With IES management, the traffic transported is always IP packets. At the termination point of the ATM link, the IP packets are extracted to the CSM for further processing.

Setting Up Connections Between the 5620 SAM and the 7705 SAR

IP over ATM is used for in-band management of the 7705 SAR. This requires the use of IP addresses so that the packets can be routed through the network using a routing table to indicate the next hop. Because Apipe interfaces (SAPs) do not have IP addresses, Apipes cannot be used to carry the management traffic.

With IES, the ATM SAP can be used for the forwarding of management IP packets. To set up a connection, IES is enabled on an interface on the 7705 SAR and the IP address for the interface is defined. A PVCC connection is then set up between the 7705 SAR and the remote router (SR) attached to the network manager (5620 SAM).

The IP datagrams are encapsulated into AAL5 for transport over the ATM network.

At the remote SR end, the SAP is bound to a VPRN instance to ensure that LDP signaling to the system IP address of the 7705 SAR flows through the IP/GRE link and not over the ATM link. Within the VPRN, an IP address is assigned at the termination SAP. The IP datagram is extracted from the ATM cell at this termination point and is routed to the 5620 SAM.

Alternatively, manually configured connections can be used instead of signaled pseudowires.



Note: The remote IP address must be manually configured and a static route must be set up between the two connections. This configuration is beyond the scope of this document; refer to the 7705 SAR OS Router Configuration Guide for information.

For redundancy, it is recommended that two VCs be configured per ATM port or IMA group. This requires the configuration of two static routes. ECMP must be enabled to allow duplicate routes in the routing table, and BFD can be enabled to trigger a faster handoff to the other route in case of route failure.

Encapsulation

To run IP traffic over ATM links, the system uses routed VC-mux encapsulation as specified in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. Since the only supported Layer 3 protocol over the management VC is IP, the VC mux encapsulation method is implemented to reduce complexity and overhead; likewise, routing mode is preferred over bridged mode.

The maximum MTU size supported is 2048 bytes.

Layer 2 and Layer 3 Traffic Management

ATM traffic descriptors can be applied at the ingress (policing) and egress (shaping and service category scheduling and prioritization) of the IES SAP in order to provide traffic management functions at Layer 2.

Management IP traffic that is destined for the CSM is classified at Layer 3 and is forwarded into the fabric from one of three of the adapter card control queues:

- high priority
- low priority
- FTP priority

The high-priority and low-priority queues are limited to 1 Mb/s and the FTP queue is rate-limited to 3 Mb/s ingress to the fabric toward the control plane.



Note: Proper configuration of the traffic descriptor profiles is essential for proper operation of the IES SAP. If no profile is assigned, the default UBR service category is assumed. All IES 7705 SAR traffic is scheduled; no shaping is supported in this mode. To ensure that IP traffic transported over the IES SAP is prioritized fairly, ATM layer traffic descriptors should be assigned. See IES Management SAP Commands in the IES Command Reference section for information.

Troubleshooting and Fault Detection Services

The IES in-band management service supports ATM OAM F4 (VP level) and F5 (VC level) cell generation and termination. For more information on OAM, refer to the 7705 SAR OS OAM and Diagnostics Guide, “OAM and SAA”.

Bidirectional forwarding detection (BFD) can also be configured on the IES interface. BFD is a simple protocol for detecting failures in a network. BFD uses a “hello” mechanism that sends control messages periodically to the far end and receives periodic control messages from the far end. BFD is implemented for static routes in asynchronous mode only, meaning that neither end responds to control messages; rather, the messages are sent in the time period configured at each end.

To support redundancy, ECMP must be enabled to allow duplicate routes in the routing table, and BFD must be enabled to trigger the handoff to the other route in case of failure.

Due to the lightweight nature of BFD, it can detect failures faster than other detection protocols, making it ideal for use in applications such as mobile transport.

If the configured number of consecutive missed BFD messages is reached, the static route to the peer is declared not active.



Note: Layer 2 AIS/RDI cells that are received on the IES SAP will disable the IP interface. Link failures detected by BFD will also disable the IP interface.

IP ECMP Load Balancing

IP ECMP allows the configuration of load balancing across all IP interfaces at the system level or interface level on the network side. You can include Layer 4 port attributes and/or the TEID attribute in the hashing algorithm with the `l4-load-balancing` and `teid-load-balancing` commands in the `config>service>ies>interface` context. Configuration at the interface level overrides the system-level settings for the specific interface.

For more information on IP ECMP, refer to the7705 SAR OS Router Configuration Guide, “Static Routes, Dynamic Routes, and ECMP”.

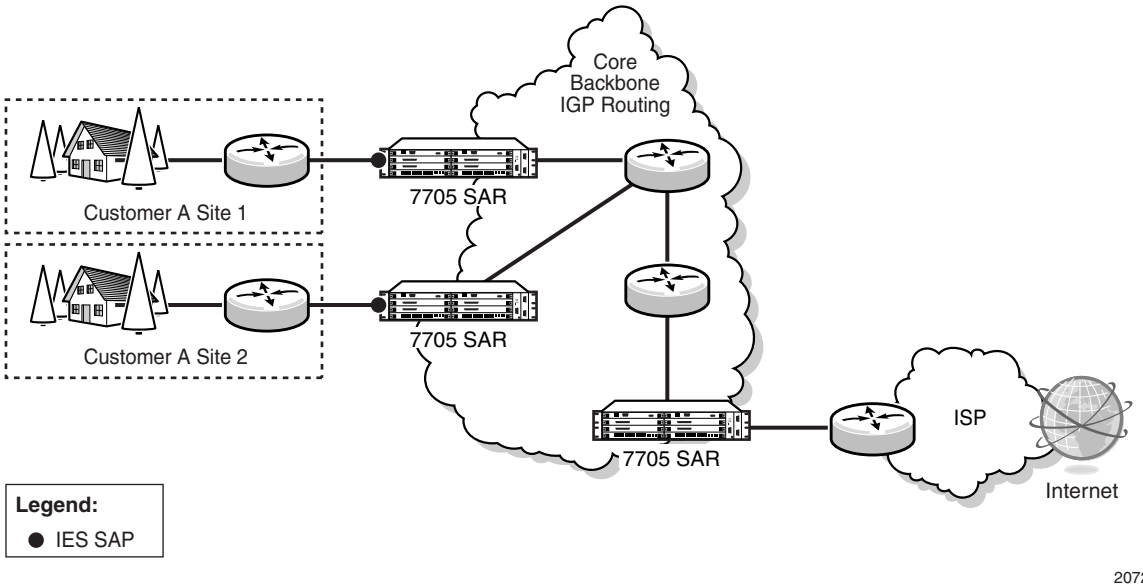
IES for Customer Traffic

Topics in this section include:

- DHCP Relay and DHCPv6 Relay
- IPCP
- IPSec Support
- Security Zones and IES
- Proxy ARP
- Configurable ARP Retry Timer
- Unnumbered Interfaces
- Troubleshooting and Fault Detection Services
- SAPs
- Spoke SDP Termination to IES

IES provides IP connectivity between customer access points. From the customer's perspective, IES provides a direct IP connection and can be used for Internet connectivity, as shown in Figure 89. The customer is assigned an IP interface and a SAP is associated with the IP interface to designate a customer access point to the service—one SAP per interface. SAPs can be MC-MLPPP, PPP/MLPPP, LAG, or null/dot1q/qinq Ethernet. SDPs are not required, because traffic is routed rather than being encapsulated in a tunnel.

Figure 89: IES for Customer Access to the Internet



IES supports static routes on customer IP interfaces (that is, SAPs). These routes are redistributed into the global routing table of the 7705 SAR.

IES is supported on the following:

- the 16-port T1/E1 ASAP Adapter card
- the 32-port T1/E1 ASAP Adapter card
- the Packet Microwave Adapter card
- any T1/E1 port of the 7705 SAR-F
- any T1/E1 port of the 7705 SAR-M (variants with T1/E1 ports)
- any T1/E1 port of the 7705 SAR-A (variants with T1/E1 ports)
- any T1/E1 port of the 4-port T1/E1 and RS-232 Combination module
- any port on the 8-port Ethernet Adapter card
- any port on the 6-port Ethernet 10Gbps Adapter card
- any port on the 8-port Gigabit Ethernet Adapter card
- any port on the 10-port 1GigE/1-port 10GigE X-Adapter card (10-port 1GigE mode)
- any port on the 4-port SAR-H Fast Ethernet module
- any port on the 6-port SAR-M Ethernet module
- any Ethernet port of the 7705 SAR-F
- any Ethernet port of the 7705 SAR-M (all variants)
- any Ethernet port of the 7705 SAR-A (both variants)
- any Ethernet port of the 7705 SAR-W
- any Ethernet port of the 7705 SAR-Wx (all variants)
- any Ethernet port of the 7705 SAR-H
- any Ethernet port of the 7705 SAR-Hc
- any Ethernet port of the 7705 SAR-X
- GPON module ports and DSL module ports when the module is installed in the 7705 SAR-M (variants with module slots)

Ports must be in access mode.

The encapsulation type for Ethernet ports must be null, dot1q, or qinq. The encapsulation type for DSL and GPON module ports on the 7705 SAR-M must be null or dot1q; however, the encapsulation type on a DSL port on the 7705 SAR-Wx can be null, dot1q, or qinq.

IES IPv6 SAPs are supported on the following cards, modules, and ports:

- the 8-port Ethernet Adapter card (version 2 only)
- the 6-port Ethernet 10Gbps Adapter card
- the 8-port Gigabit Ethernet Adapter card
- the 10-port 1GigE/1-port 10GigE X-Adapter card (10-port 1GigE mode)
- the Packet Microwave Adapter card
- the 4-port SAR-H Fast Ethernet module
- the 6-port SAR-M Ethernet module
- any Ethernet port on the 7705 SAR-F
- any Ethernet port on the 7705 SAR-M (all variants)
- any Ethernet port on the 7705 SAR-A (both variants)
- the 7705 SAR-W
- any Ethernet port on the 7705 SAR-Wx (all variants)
- the 7705 SAR-H
- any Ethernet port on the 7705 SAR-Hc
- any Ethernet port of the 7705 SAR-X
- GPON module ports and DSL module ports when the module is installed in the 7705 SAR-M (variants with module slots)

For more information on IPv6 addressing, refer to the 7705 SAR OS Router Configuration Guide, “Internet Protocol Versions”.

More than one Internet Enhanced Service can be created for a single customer ID, and more than one IP interface can be created within a single IES. All IP interfaces created within an IES belong to the same customer.

The service provider applies billing, ingress/egress shaping and policing to the customer.



Note:

- Internet Enhanced Services require that the fabric mode be set to aggregate mode rather than per-destination mode. IES is only supported with aggregate-mode fabric profiles. If the fabric mode is set to per-destination mode, creation of the Internet Enhanced Service is blocked through the CLI. The fabric mode must be changed to aggregate mode before IES can be configured. As well, if IES is configured, alteration of the fabric mode is blocked.
- For information on configuring fabric mode, refer to the 7705 SAR OS Quality of Service Guide, “Configurable Ingress Shaping to Fabric (Access and Network)”.

DHCP Relay and DHCPv6 Relay

The 7705 SAR provides DHCP/BOOTP Relay agent services and DHCPv6 Relay agent services for DHCP clients. DHCP is used for IPv4 network addresses and DHCPv6 is used for IPv6 network addresses. Both DHCP and DHCPv6 are known as stateful protocols because they use dedicated servers to maintain parameter information.

Unless stated otherwise, DHCP is equivalent to "DHCP for IPv4", or DHCPv4.

In the stateful autoconfiguration model, hosts obtain interface addresses and/or configuration information and parameters from a server. The server maintains a database that keeps track of which addresses have been assigned to which hosts.

The 7705 SAR supports DHCP Relay on access IP interfaces associated with IES and VPRN and on network interfaces. Each DHCP instance supports up to eight DHCP servers.

The 7705 SAR supports DHCPv6 Relay on access IP interfaces associated with IES and VPRN. Each DHCPv6 instance supports up to eight DHCPv6 servers.

Note:

- The 7705 SAR acts as a Relay agent for DHCP and DHCPv6 requests and responses, and can also be configured to function as a DHCP or DHCPv6 server. DHCPv6 functionality is only supported on network interfaces and on access IP interfaces associated with VPRN.
- When used as a CPE, the 7705 SAR can act as a DHCP client to learn the IP address of the network interface. Dynamic IP address allocation is supported on both network and system interfaces.
- For more information on DHCP and DHCPv6, refer to the 7705 SAR OS Router Configuration Guide, "DHCP and DHCPv6".

DHCP Relay

The 7705 SAR provides DHCP/BOOTP Relay agent services for DHCP clients. DHCP is a configuration protocol used to communicate network information and configuration parameters from a DHCP server to a DHCP-aware client. DHCP is based on the BOOTP protocol, with additional configuration options and the added capability of allocating dynamic network addresses. DHCP-capable devices are also capable of handling BOOTP messages.

A DHCP client is an IP-capable device (typically a computer or base station) that uses DHCP to obtain configuration parameters such as a network address. A DHCP server is an Internet host or router that returns configuration parameters to DHCP clients. A DHCP/BOOTP Relay agent is a host or router that passes DHCP messages between clients and servers.

Home computers in a residential high-speed Internet application typically use the DHCP protocol to have their IP address assigned by their Internet service provider.

The DHCP protocol requires the client to transmit a request packet with a destination broadcast address of 255.255.255.255 that is processed by the DHCP server. Since IP routers do not forward broadcast packets, this would suggest that the DHCP client and server must reside on the same network segment. However, for various reasons, it is sometimes impractical to have the server and client reside in the same IP network. When the 7705 SAR is acting as a DHCP Relay agent, it processes these DHCP broadcast packets and relays them to a preconfigured DHCP server. Therefore, DHCP clients and servers do not need to reside on the same network segment.

DHCP OFFER messages are not dropped if they contain a yiaddr that does not match the local configured subnets on the DHCP relay interface. This applies only to regular IES and VPRN interfaces with `no lease-populate` configured on the DHCP relay interface.

DHCP Options

DHCP options are codes that the 7705 SAR inserts in packets being forwarded from a DHCP client to a DHCP server. Some options have additional information stored in suboptions.

The 7705 SAR supports the Relay Agent Information Option 82 as specified in RFC 3046. The following suboptions are supported:

- circuit ID
- remote ID
- vendor-specific options

DHCPv6 Relay

DHCPv6 Relay operation is similar to DHCP in that servers send configuration parameters such as IPv6 network addresses to IPv6 nodes, but DHCPv6 Relay is not based on the DHCP or BOOTP protocol. DHCPv6 can be used instead of stateless autoconfiguration (refer to the 7705 SAR OS Router Configuration Guide, “Neighbor Discovery”) or in conjunction with it.

DHCPv6 is also oriented around IPv6 methods of addressing, especially the use of reserved, link-local scoped multicast addresses. DHCPv6 clients transmit messages to these reserved addresses, allowing messages to be sent without the client knowing the address of any DHCP server. This transmission allows efficient communication even before a client has been assigned an IP address. When a client has an address and knows the identity of a server, it can communicate with the server directly using unicast addressing.

The DHCPv6 protocol requires the client to transmit a request packet with a destination multicast address of ff02::1:2 (all DHCP servers and relay agents on the local network segment) that is processed by the DHCP server.

Similar to DHCP address allocation, if a client needs to obtain an IPv6 address and other configuration parameters, it sends a Solicit message to locate a DHCPv6 server, then requests an address assignment and other configuration information from the server. Any server that can meet the client's requirements responds with an Advertise message. The client chooses one of the servers and sends a Request message, and the server sends back a Reply message with the confirmed IPv6 address and configuration information.

If the client already has an IPv6 address, either assigned manually or obtained in some other way, it only needs to obtain configuration information. In this case, exchanges are done using a two-message process. The client sends an Information Request message, requesting only configuration information. A DHCPv6 server that has configuration information for the client sends back a Reply message with the information.

The 7705 SAR supports the DHCPv6 Relay Agent option in the same way that it supports the DHCP Relay Agent option. This means that when the 7705 SAR is acting as a DHCPv6 Relay Agent, it relays messages between clients and servers that are not connected to the same link.

DHCPv6 Options

DHCPv6 options are codes that the 7705 SAR inserts in packets being forwarded from a DHCPv6 client to a DHCPv6 server. DHCPv6 supports interface ID and remote ID options as defined in RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPV6)* and RFC 4649, *DHCPv6 Relay Agent Remote-ID Option*.

IPCP

Similar to DHCP over Ethernet interfaces, Internet Protocol Control Protocol (IPCP) extensions to push IP information over PPP/MLPPP IES SAPs are supported. Within this protocol, extensions can be configured to define the remote IP address and DNS IP address to be signaled via IPCP on the associated PPP interface. The IPCP-based IP and DNS assignment process is similar to DHCP behavior; IPCP-based IP/DNS assignment uses PPP/MLPPP IP layer protocol handshake procedures. PPP/MLPPP connected devices hooked up to IES can benefit from this feature for the assignment of IP and DNS to the associated interface.

IPSec Support

The 7705 SAR supports IPSec and IPSec tunnels, where IES is used as a public (untrusted) network-facing service and VPRN is used as a private (trusted) network-facing service. IES interfaces support the provisioning of tunnel SAPs as part of IPSec provisioning. The *sap-id* for a tunnel SAP is `tunnel-1.public:`*tag*.

For more information, see the IPSec chapter in this guide.

Security Zones and IES

The 7705 SAR supports a number of mechanisms for node security, including Access Control Lists (ACLs), Network Address Translation (NAT), and stateful, zone-based firewalls. For information about ACLs, NAT, and firewalls, refer to the 7705 SAR OS Router Configuration Guide, "Configuring Security Parameters".

NAT and firewall security configurations are both based on zones. Zones segment a network, making it easier to control and organize traffic. A zone consists of a group of Layer 3 interfaces with common criteria, bundled together. Security policies, which define a set of rules that determine how NAT or firewall should direct traffic, can be applied to the entire zone or to multiple zones. To enable NAT or firewall functionality, security policy and profile parameters must be configured under the `config>security` context in the CLI, and a security zone must be configured under one or more of the following contexts:

- `config>router>zone`
- `config>service>vprn>zone`
- `config>service>ies>zone`

A zone is created by adding at least one Layer 3 interface to the zone configuration. Multiple zones can be created within each service or within the router context. Layer 3 interfaces from different services cannot be grouped into a single common zone. Table 86 lists the supported interfaces that can be added to zones in each CLI context for NAT or firewall.

Table 86: Security Zone Interfaces per Context

CLI Context	Interface Type	NAT	Firewall
Router	Layer 3	✓	✓
VPRN	SAP	✓	✓
	Spoke-SDP termination	✓	✓
	IPSec private		
	Routed VPLS	✓	✓
IES	SAP	✓	✓
	Spoke-SDP termination	✓	✓
	IPSec public	✓	
	Routed VPLS	✓	✓

Proxy ARP

Proxy ARP is supported on IES interfaces.

Proxy ARP is a technique by which a router on one network responds to ARP requests intended for another node that is physically located on another network. The router effectively pretends to be the destination node by sending an ARP response to the originating node that associates the router's MAC address with the destination node's IP address (acts as a proxy for the destination node). The router then takes responsibility for routing traffic to the real destination.

For more information on proxy ARP, refer to the 7705 SAR OS Router Configuration Guide, "Proxy ARP".

Configurable ARP Retry Timer

A timer is available to configure a shorter retry interval when an ARP request fails. An ARP request may fail for a number of reasons, such as network connectivity issues. By default, the 7705 SAR waits 5000 ms before retrying an ARP request. The configurable retry timer makes it possible to shorten the retry interval to between 100 and 30 000 ms.



Note: The ARP retry default value of 5000 ms is intended to protect CPU cycles on the 7705 SAR, especially when it has a large number of interfaces. Configuring the ARP retry timer to a value shorter than the default should be done only on mission-critical links, such as uplinks or aggregate spoke SDPs transporting mobile traffic; otherwise, the retry interval should be left at the default value.

The configurable ARP retry timer is supported on VPRN and IES service interfaces, as well on the router interface.

Unnumbered Interfaces

Unnumbered interfaces are supported on IES and VPRN services for IPv4. Unnumbered interfaces are point-to-point interfaces that are not explicitly configured with a dedicated IP address and subnet; instead, they borrow (or link to) an IP address from another interface on the system (the system IP address, another loopback interface, or any other numbered interface) and use it as the source IP address for packets originating from the interface.

This feature is supported via both dynamic and static ARP for unnumbered interfaces to allow interworking with unnumbered interfaces that may not support dynamic ARP.

The use of unnumbered interfaces has no effect on IPv6 routes; however, the `unnumbered` command must only be used in cases where IPv4 is active (IPv4 only and mixed IPv4/IPv6 environments). When using an unnumbered interface for IPv4, the loopback address used for the unnumbered interface must have an IPv4 address. The interface type for the unnumbered interface is automatically point-to-point.

Troubleshooting and Fault Detection Services

Bidirectional forwarding detection (BFD) can be configured on the IES interface. BFD is a simple protocol for detecting failures in a network. BFD uses a “hello” mechanism that sends control messages periodically to the far end and expects to receive periodic control messages from the far end. On the 7705 SAR, BFD is implemented for static routes in asynchronous mode only, meaning that neither end responds to control messages; rather, the messages are sent periodically from each end.

To support redundancy with fast switchover, BFD must be enabled to trigger the handoff to the other route in case of failure.

Due to the lightweight nature of BFD, it can detect failures faster than other detection protocols, making it ideal for use in applications such as mobile transport.

If BFD packets are not received in the configured amount of time, the associated static route is declared “not active”, causing a reroute to an alternative path, if any.

Note: Link failures detected by BFD will disable the IP interface.

The 7705 SAR also supports Internet Control Message Protocol (ICMP and ICMPv6). ICMP is a message control and error reporting protocol that also provides information relevant to IP packet processing. For more information on ICMP and ICMPv6, refer to the 7705 SAR OS Router Configuration Guide, “ICMP and ICMPv6”.

SAPs

This section contains information on the following topics:

- Encapsulations
- Routing Protocols
- QoS Policies
- QinQ (IES)
- IP Filter Policies on an IES SAP

Encapsulations

The following SAP encapsulations are supported on the 7705 SAR Internet Enhanced Service:

- Ethernet null
- Ethernet dot1q
- Ethernet qinq
- PPP/MLPPP/MC-MLPPP

Routing Protocols

The OSPFv2 routing protocol is supported on IES SAPs (that is, access IP interfaces).

The SAP for the IES IP interface is created at the IES service level, but the routing protocol for the IES IP interface is configured at the routing protocol level for the main router instance in the global context. Refer to the chapter on OSPF in the 7705 SAR OS Routing Protocols Guide for information on configuring the OSPF routing protocol.

QoS Policies

When applied to an Internet Enhanced Service SAP, service ingress QoS policies only create the unicast queues defined in the policy.

Service egress QoS policies function in the same way as Ethernet and IP pseudowire services, where class-based queues are created based on the QoS policy. Multiple queues are supported. Refer to the 7705 SAR OS Quality of Service Guide, “Creating a Service Egress QoS Policy”.

Both Layer 2 and Layer 3 match criteria can be used in the QoS policies for traffic classification in an IES.

QinQ (IES)

IES supports QinQ functionality. For details, see QinQ Support.

IP Filter Policies on an IES SAP

IPv4 filter policies can be applied to ingress IES management SAPs.

IPv4 and IPv6 filter policies can be applied to both ingress and egress IES SAPs (null, dot1q, or qinq interfaces).

Configuration and assignment of IP filter policies is similar for all services. Refer to the 7705 SAR OS Router Configuration Guide, “Filter Policies”, for information on configuring IP filters.

Spoke SDP Termination to IES

This feature enables a customer to exchange traffic between a VLL or VPLS (Layer 2) service and an IES or VPRN (Layer 3) service. Customer premises traffic coming in from a VLL or VPLS service (SAP to spoke SDP) is forwarded over the IP/MPLS network to the IES or VPRN service, and vice versa. Network QoS policies can be applied to the spoke SDP to control traffic forwarding to the Layer 3 service.

In a Layer 3 spoke SDP termination to an IES or VPRN service, where the destination IP address resides within the IES or VPRN network, CE device-generated ARP frames must be processed by the Layer 3 interface. When an ARP frame is received over the spoke SDP at the Layer 3 interface endpoint, the 7705 SAR responds to the ARP frame with its own MAC address. When an ARP request is received from the routed network and the ARP entry for the CE device that is connected to the spoke SDP is not known, the 7705 SAR initiates an ARP frame to resolve the MAC address of the next hop or CE device.

Figure 90 shows traffic terminating on a specific IES or VPRN service that is identified by the SDP ID and VC label present in the service packet.

Figure 90: SDP ID and VC Label Service Identifiers (Conceptual View of the Service)

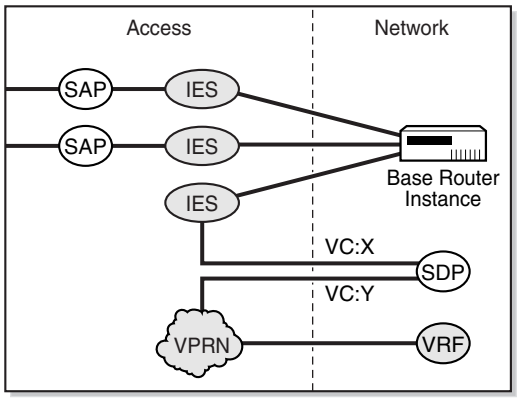
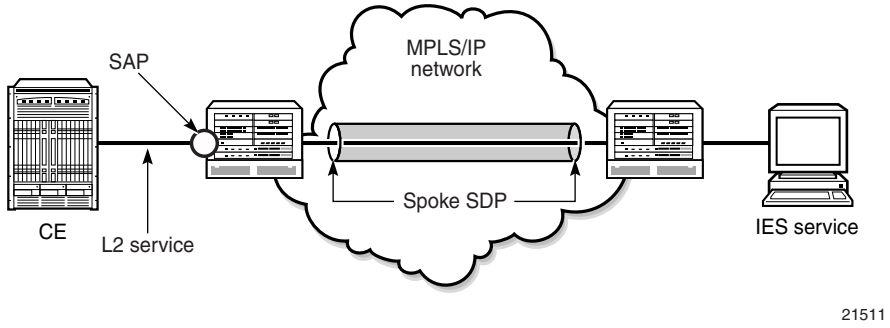


Figure 91 shows a spoke SDP terminating directly into an IES. In this case, a spoke SDP could be tied to an Epipe or a hierarchical VPLS service. There is no configuration required on the PE connected to the CE.

Figure 91: IES Spoke SDP Termination



Ethernet spoke SDP termination for IES is supported over the following network uplinks:

- Ethernet network ports (null or dot1q encapsulation)
- PPP/MLPPP network ports. For information on PPP/MLPPP ports, refer to the 7705 SAR OS Interface Configuration Guide, “Access, Network, and Hybrid Ports”
- GPON module ports and DSL module ports when the module is installed in the 7705 SAR-M (variants with module slots)
- POS ports

Spoke SDP termination for IES supports the following:

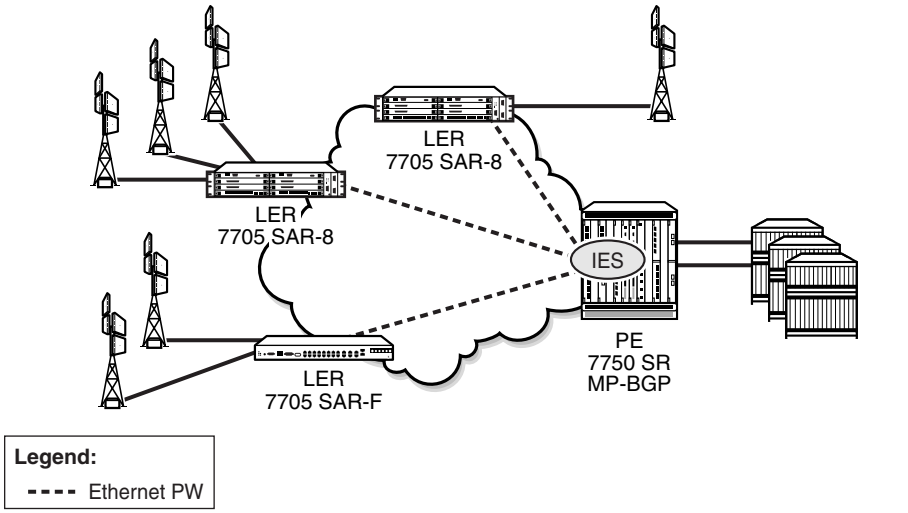
- Ethernet PW to VRF
- interface shutdown based on PW standby signaling
- spoke SDP ingress IP filtering with filter logging
- label withdrawal for spoke SDPs terminated on IES
- statistics collection
- VCCV ping (type 2)

A spoke SDP on an IES interface can be connected to the following entities:

- Epipe spoke SDP
- Epipe spoke SDP redundancy with standby-signal-master enabled
- IES interface
- VPRN interface
- VPLS spoke SDP
- VPLS spoke SDP redundancy with suppress-standby-signaling disabled

Figure 92 shows an example of backhauling from a given site that uses PW and IES on the 7705 SAR. An individual PW is configured on a per-CE device or a per-service basis. For routing services, this PW can be terminated to an IES at the 7750 SR end. This scenario offers per-service OAM and redundancy capabilities. Because there is no local communication on the remote 7705 SAR, traffic between any two devices connected to the 7705 SAR must traverse through the 7750 SR at the MTSO/CO.

Figure 92: Pseudowire-Based Backhaul (Spoke SDP Termination at 7750 SR)



Configuring IES with CLI

This section provides the information required to configure IP routing services; that is, direct forwarding of IP traffic between CE devices, and to configure IES for in-band management of the 7705 SAR over ATM links.

Topics in this section include:

- Common Configuration Tasks
- Configuring IES Components
- Service Management Tasks

Common Configuration Tasks

The following list provides a brief overview of the tasks that must be performed to configure IES.

- Associate the IES service with a customer ID.
- Create an IP interface on the 7705 SAR.
- Specify the IP address of the interface.
- Define interface parameters.
- Define SAP parameters.
- For IES spoke SDP applications only — define spoke SDP parameters.
- For IES applications only — configure VRRP (optional).
- For IES management service only—manually configure the remote address of the far-end router to which the 5620 SAM network manager is connected (far-end router must be enabled for IES service).*
- For IES management service only—create a static route to the remote router and to the 5620 SAM.*
- Enable the service.



Note: *Remote address and static route configuration is beyond the scope of this document. For information, refer to the 7705 SAR OS Router Configuration Guide.

Configuring IES Components

This section provides configuration examples for components of the IES service. Each component includes some or all of the following: introductory information, CLI syntax, a specific CLI example, and a sample CLI display output.

Topics in this section include:

- Creating an IES Service
- Configuring Interface Parameters
- Configuring IES SAP Parameters
- Configuring IES Spoke SDP Parameters
- Configuring VRRP
- Configuring a Security Zone within IES

Creating an IES Service

Use the following CLI syntax to create an IES service.

CLI Syntax:

```
config>service# ies service-id [customer customer-id]
  [create] [vpn vpn-id]
      description description-string
      interface ip-int-name [create]
      no shutdown
```

Example:

```
A:ALU-41>config>service# ies 5 customer 1 create
A:ALU-41>config>service>ies# description "IES for
 in-band management"
A:ALU-41>config>service>ies# interface "ATMoIP
 Management" create
A:ALU-41>config>service>ies# no shutdown
A:ALU-41>config>service>ies#
```

The following example displays the IES service creation output.

```
A:ALU-41>config>service# info
------------------------------------
...
        ies 5 customer 1 create
            description "IES for in-band management"
            interface "ATMoIP Management"
            no shutdown
        exit
...
```

Configuring Interface Parameters

Configure interface parameters for:

- IES Management Service
- IES Service
- IES IPv6 Service

IES Management Service

Use the following CLI syntax to configure interface parameters for the IES management service.

CLI Syntax:

```
config>service# ies service-id [customer customer-id]
  [create] [vpn vpn-id]
      interface ip-int-name
           address if-ip-address
           bfd transmit-interval [receive receive-
              interval] [multiplier multiplier] [type np]
           description description-string
           ip-mtu octets
           no shutdown
```

Example:

```
A:ALU-41>config>service# ies 5
A:ALU-41>config>service>ies# interface "ATMoIP
 Management"
A:ALU-41>config>service>ies>if# address 3.3.3.3/24
A:ALU-41>config>service>ies>if# ip-mtu 1524
A:ALU-41>config>service>ies>if# no shutdown
A:ALU-41>config>service>ies>if#
```

The following example displays the IES interface creation output for the IES management service.

```
A:ALU-41>config>service>ies>if# info detail
----------------------------------------
...
            no description
            address 3.3.3.3/24
            ip-mtu 1524
            no bfd
            exit
            no shutdown
...
----------------------------------
```

IES Service

Use the following CLI syntax to configure interface parameters for the IES service.

CLI Syntax:

```
config>service# ies service-id [customer customer-id]
  [create] [vpn vpn-id]
      interface ip-int-name
           address if-ip-address
           allow-directed-broadcasts
           arp-timeout
           bfd transmit-interval [receive receive-
             interval] [multiplier multiplier] [type np]
           description description-string
           dhcp
                description description-string
                option
                     action {replace | drop | keep}
                     circuit-id [ascii-tuple | ifindex |
                          sap-id | vlan-ascii-tuple]
                     remote-id [mac | string string]
                     vendor-specific option
                          client-mac-address
                          sap-id
                          service-id
                          string text
                          system-id
                     server server1 [server2...(up to 8
                          max)]
                no shutdown
                trusted
           icmp
                mask-reply
                ttl-expired [number seconds]
                unreachables
           ip-mtu octets
           ipcp
                dns ip-address [secondary ip-address]
                dns secondary ip-address
                peer-ip-address ip-address
           local proxy-arp
           mac ieee-address
           proxy-arp-policy policy-name [policy-
             name...(up to 5 max)]
           remote-proxy-arp
           secondary {ip-address/mask | ip-address
             netmask} [broadcast all-ones | host-ones]
             [igp-inhibit]
           no shutdown
           static-arp ip-address ieee-mac-address
```

```
                        static-arp ieee-mac-address unnumbered
                        unnumbered {ip-int-name | ip-address}
                    no shutdown
```

Example:

```
A:ALU-41>config>service# ies 4
A:ALU-41>config>service>ies$ interface "to Internet"
A:ALU-41>config>service>ies>if$ address 3.2.3.3/24
A:ALU-41>config>service>ies>if$ dhcp option
A:ALU-41>config>service>ies>if>dhcp>option$ circuit-id
  ifindex
A:ALU-41>config>service>ies>if>dhcp>option$ exit
A:ALU-41>config>service>ies>if$ ip-mtu 1524
```

The following example displays the IES interface creation output for the IES service.

```
A:ALU-41>config>service>ies>if# info detail
----------------------------------------
...
             no description
             address 3.2.3.3/24 broadcast host-ones
             no mac
             arp-timeout 14400
             no allow-directed-broadcasts
             icmp
                mask-reply
                unreachables 100 10
                ttl-expired 100 10
             exit
             dhcp
                shutdown
                no description
                option
                    action keep
                    circuit-id ifindex
                    no remote-id
                    no vendor-specific-option
                exit
                no server
                no trusted
             exit
             ip-mtu 1524
             no bfd
             ipcp
                no peer-ip-address
                no dns
             exit
             proxy-arp policy "proxyARPpolicy"
             local proxy-arp
             remote proxy-arp
             no shutdown...
-------------------------------------
```

IES IPv6 Service

Use the following CLI syntax to configure interface parameters for the IES IPv6 service.

CLI Syntax:

```
config>service# ies service-id [customer customer-id]
  [create] [vpn vpn-id]
      interface ip-int-name
           ipv6
                address ipv6-address/prefix-length
                   [eui-64]
                dhcp6-relay
                     description description-string
                     option
                          interface-id
                          interface-id ascii-tuple
                          interface-id ifindex
                          interface-id sap-id
                          interface-id string
                          remote-id
                     server ipv6-address
                          [ipv6-address...(up to 8 max)]
                     no shutdown
                icmp6
                     packet-too-big [number seconds]
                     param-problem [number seconds]
                     time-exceeded [number seconds]
                     unreachables [number seconds]
                neighbor ipv6-address mac-address
                reachable-time seconds
                stale-time seconds
```

Example:

```
A:ALU-41>config>service# ies 9
A:ALU-41>config>service>ies$ interface "ies_interface"
A:ALU-41>config>service>ies>if$ ipv6
A:ALU-41>config>service>ies>if>ipv6$ address
 1080:6809:8086:6502::/64
A:ALU-41>config>service>ies>if>ipv6$ dhcp6-relay
A:ALU-41>config>service>ies>if>ipv6>dhcp6-relay$ server
 2001:DB8::
A:ALU-41>config>service>ies>if>ipv6>dhcp6-relay$ option
A:ALU-41>config>service>ies>if>ipv6>dhcp6-relay>option$
 interface-id ascii-tuple
A:ALU-41>config>service>ies>if>ipv6>dhcp6-relay>option$
 exit
A:ALU-41>config>service>ies>if>ipv6$ icmp6
A:ALU-41>config>service>ies>if>ipv6>icmp6$ packet-too-
 big 80 10
A:ALU-41>config>service>ies>if>ipv6>icmp6$ exit
```

```
config>service>ies>>if>ipv6# neighbor 2001:DB8:CAFE::60
 00-50-56-A3-04-0C
config>service>ies>>if>ipv6>neighbor# exit
config>service>ies>>if>ipv6# reachable-time 30
config>service>ies>>if>ipv6# stale-time 14400
config>service>ies>>if>ipv6# exit
```

The following example displays the IES interface IPv6 output.

```
A:ALU-41>config>service>ies>if># info detail
----------------------------------------
...
            no description
            address 1080:6809:8086:6502::/64
            dhcp6-relay
                no description
                option
                    interface-id ascii-tuple
                    no remote-id
                server 2001:DB8::
            exit
            icmp6
                packet-too-big 80 10
                param-problem 100 10
                time-exceeded 100 10
                unreachables 100 10
                exit
            exit
            ,,.
          reachable-time 30
          stale-time 14400
       exit
...
```

Configuring IES SAP Parameters

Configure IES SAP parameters for:

- IES Management SAP
- IES Service SAP

IES Management SAP

Use the following CLI syntax to configure IES management SAP parameters.



Note: The encapsulation type is always aal5mux-ip.

CLI Syntax:

```
config>service# ies service-id [customer customer-id]
  [create] [vpn vpn-id]
     interface ip-int-name
          sap sap-id [create]
               atm
                    encapsulation encap-type
                    egress
                         traffic-desc traffic-desc-
                           profile-id
                    ingress
                         traffic-desc traffic-desc-
                           profile-id
                    oam
                         alarm-cells
               description description-string
               ingress
                    filter ip ip-filter-id
               no shutdown
```

Example:

```
A:ALU-41>config>service# ies 5
A:ALU-41>config>service>ies# interface "ATMoIP
 Management"
A:ALU-41>config>service>ies>if# sap 1/1/1.1:0/32 create
A:ALU-41>config>service>ies>if>sap# ingress
A:ALU-41>config>service>ies>if>sap>ingress# filter ip 3
A:ALU-41>config>service>ies>if>sap>ingress# exit
A:ALU-41>config>service>ies>if>sap# atm
A:ALU-41>config>service>ies>if>sap>atm# encapsulation
 aal5mux-ip
```

```
A:ALU-41>config>service>ies>if>sap>atm# egress
A:ALU-41>config>service>ies>if>sap>atm>egress# traffic-
 desc 3
A:ALU-41>config>service>ies>if>sap>atm>egress# exit
A:ALU-41>config>service>ies>if>sap>atm# ingress
A:ALU-41>config>service>ies>if>sap>atm>ingress#
 traffic-desc 2
A:ALU-41>config>service>ies>if>sap>atm>ingress# exit
A:ALU-41>config>service>ies>if>sap>atm# oam
A:ALU-41>config>service>ies>if>sap>atm>oam# alarm-cells
A:ALU-41>config>service>ies>if>sap>atm>oam# exit
A:ALU-41>config>service>ies>if>sap>atm# exit
A:ALU-41>config>service>ies>if>sap# exit
A:ALU-41>config>service>ies>if# exit
A:ALU-41>config>service>ies#
```

IES Service SAP

The following example displays the IES SAP creation output.

```
A:ALU-41>config>service>ies>if>sap# info detail
----------------------------------------
...
          no description
          ingress
              filter ip 3
          exit
          atm
              encapsulation aal5mux-ip
              ingress
                  traffic-desc 2
              exit
              egress
                 traffic-desc 3
              exit
              oam
                  alarm-cells
              exit
          exit
          no shutdown
------------------------------------
```

Use the following CLI syntax to configure SAP parameters for the IES service.

CLI Syntax: `config>service# ies` *`service-id`* `[customer` *`customer-id`*`]`
`[create] [vpn` *`vpn-id`*`]`
`interface` *`ip-int-name`*
`sap` *`sap-id`* `[create]`
`accounting policy` *`acct-policy-id`*

```
                    collect stats
                    description description-string
                    egress
                         filter ip ip-filter-id
                         filter ipv6 ipv6-filter-id
                         qos policy-id
                    ingress
                         filter ip ip-filter-id
                         filter ipv6 ipv6-filter-id
                         qos policy-id
                    no shutdown
```

Example:
```
A:ALU-41>config>service# ies 4
A:ALU-41>config>service>ies$ interface "to Internet"
A:ALU-41>config>service>ies>if$ sap 1/4/1 create
A:ALU-41>config>service>ies>if>sap$ egress
A:ALU-41>config>service>ies>if>sap>egress$ qos 3
A:ALU-41>config>service>ies>if>sap$ ingress
A:ALU-41>config>service>ies>if>sap>ingress$ filter ip 3
```

The following example displays the IES SAP creation output.

```
A:ALU-41>config>service>ies>if>sap# info detail
----------------------------------------
...
          no description
          egress
              qos 3
          ingress
              filter ip 3
          exit
          no shutdown
------------------------------------
```

Configuring IES Spoke SDP Parameters

Use the following CLI syntax to configure spoke SDP parameters for the IES service.

CLI Syntax:
```
config>service# ies service-id [customer customer-id]
  [create] [vpn vpn-id]
      interface ip-int-name
           spoke-sdp sdp-id:vc-id [create]
                egress
                     vc-label egress-vc-label
                ingress
                     filter ip ip-filter-id
                     vc-label ingress-vc-label
                [no] shutdown
```

Example:

```
A:ALU-41>config>service# ies 6
A:ALU-41>config>service>ies$ interface "ies6_interface"
A:ALU-41>config>service>ies>if$ spoke-sdp 5:6 create
A:ALU-41>config>service>ies>if>spoke-sdp$ ingress
A:ALU-41>config>service>ies>if>spoke-sdp>ingress$
 filter ip 56
A:ALU-41>config>service>ies>if>spoke-sdp>ingress$ vc-
 label 5566
```

The following example displays the IES spoke SDP creation output.

```
A:ALU-41>config>service>ies>if>spoke SDP# info detail
----------------------------------------
...
          no description
          egress
              no vc-label
          ingress
              filter ip 56
              vc-label 5566
          exit
          no shutdown
------------------------------------
```

Configuring VRRP

Configuring VRRP policies and instances on service interfaces is optional. The basic owner and non-owner VRRP configurations on an IES interface must specify the backup `ip-address` parameter.

VRRP helps eliminate the single point of failure in a routed environment by using virtual router IP address(es) shared between two or more routers connecting the common domain. VRRP provides dynamic failover of the forwarding responsibility to the backup router if the master becomes unavailable.

The VRRP implementation allows one master per IP subnet. All other VRRP instances in the same domain must be in backup mode.

For further information about VRRP CLI syntax and command descriptions, see the IES Command Reference.

The following displays an IES interface VRRP owner configuration:

```
config>service>ies> info
#------------------------------------------
interface "vrrpowner"
address 10.10.10.23/24
vrrp 1 owner
backup 10.10.10.23
```

```
authentication-key "testabc"
exit
exit
#-------------------------------------------
config>service>ies#
```

Configuring a Security Zone within IES

To configure NAT or firewall security, you must:

- configure a NAT or firewall security profile and policy in the `config>security` context
  - → in the `config>security>profile` context, specify the timeouts for the TCP/UDP/ICMP protocols and configure logging and application assurance parameters. This step is optional. If you do not configure the profile, a default profile is assigned.
  - → in the `config>security>policy` context, configure a security policy, specify the match criteria and the action to be applied to a packet if a match is found.
- configure a security zone and apply the policy ID to the zone, as shown in the CLI syntax below

CLI Syntax:

```
config>service
     ies service-id [customer customer-id] [create]
     abort
     begin
     commit
     zone zone-id [create]
          description description-string
          interface ip-int-name [create]
          name zone-name
          nat
               pool pool-id [create]
                    description description-string
                    direction {zone-outbound | zone-
                         inbound | both}
                    entry entry-id [create]
                         ip-address ip-address [to ip-
                           address] interface ip-int-
                           name
                         port port [to port] interface
                           ip-int-name
                    name pool-name
          policy policy-id | policy-name
          shutdown
```

The following example displays a NAT zone configuration output.

```
A:ALU-B>config>service>ies# info
----------------------------------------------
        configure
            service ies 10 create
                zone 1 create
                begin
                    name "IES zone"
                    description "uplink zone from private"
                    interface ies-100-10.30.10.1
                    exit
                    nat
                        pool 1 create
                            description "pool 1"
                            direction zone-inbound
                            exit
                            entry 1 create
                                ip-addr interface ies-100-10.30.10.1
                            exit
                        exit
                    exit
                    policy 1 nat pool 1
                    commit
                exit
                no-shutdown
----------------------------------------------
A:ALU-B>config>service>ies#
```

Service Management Tasks

This section discusses the following service management tasks:

- Modifying IES Service Parameters
- Disabling an IES Service
- Re-enabling an IES Service
- Deleting an IES Service

Modifying IES Service Parameters

Existing IES service parameters can be modified, added, removed, enabled, or disabled.

To display a list of customer IDs, use the `show>service>customer` command.

Enter the parameters (such as description, interface information, or SAP information), and then enter the new information.

The following is an example of changing the IP MTU size.

Example:
```
A:ALU-41>config>service# ies 5
A:ALU-41>config>service>ies# interface "testname"
A:ALU-41>config>service>ies>if# ip-mtu 1517
A:ALU-41>config>service>ies>if# exit
```

Disabling an IES Service

An IES service can be shut down without deleting the service parameters.

Use the `shutdown` command to shut down an IES service.

CLI Syntax:
```
config>service# ies service-id
     shutdown
```

Example:
```
A:ALU-41>config>service# ies 5
A:ALU-41>config>service>ies# shutdown
A:ALU-41>config>service>ies# exit
```

Re-enabling an IES Service

Use the `no shutdown` command to re-enable a previously disabled IES service.

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

Example:
```
A:ALU-41>config>service# ies 5
A:ALU-41>config>service>ies# no shutdown
A:ALU-41>config>service>ies# exit
```

Deleting an IES Service

An IES service cannot be deleted until SAPs, spoke SDPs, and interfaces are shut down and deleted and the service is shut down on the service level.

Use the following CLI syntax to delete an IES service:

CLI Syntax:
```
config>service#
     ies service-id
          interface ip-int-name
               sap sap-id
                    shutdown
                    exit
               no sap sap-id
               spoke-sdp sdp-id:vc-id
                    shutdown
                    exit
               no spoke-sdp sdp-id:vc-id
          interface ip-int-name
               shutdown
               exit
          no interface ip-int-name
          shutdown
          exit
     no ies service-id
```

IES Command Reference

Command Hierarchies

- Configuration Commands
  - → IES Management Configuration Commands
  - → IES Service Configuration Commands
  - → Routed VPLS Commands
  - → VRRP Commands
  - → IES Security Zone Configuration Commands
- Show Commands
- Clear Commands
- Debug Commands

Configuration Commands

IES Management Configuration Commands

```
config
    — service
        — ies service-id [customer customer-id] [create] [vpn vpn-id]
        — no ies service-id
            — description description-string
            — no description
            — interface ip-int-name [create]
            — no interface ip-int-name
                — address {ip-address/mask | ip-address netmask}
                — no address
                — bfd transmit-interval [receive receive-interval] [multiplier multiplier] [type
                    np]
                — no bfd
                — description description-string
                — no description
                — ip-mtu octets
                — no ip-mtu
                — sap sap-id [create]
                — no sap sap-id
                    — atm
                        — encapsulation atm-encap-type
                        — egress
                            — traffic-desc traffic-desc-profile-id
                            — no traffic-desc
                        — ingress
                            — traffic-desc traffic-desc-profile-id
                            — no traffic-desc
                        — oam
                            — [no] alarm-cells
                    — description description-string
                    — no description
                    — ingress
                        — filter ip ip-filter-id
                        — no filter ip
                        — no filter ip [ip ip-filter-id]
                    — [no] shutdown
                — [no] shutdown
            — service-name service-name
            — no service-name
            — [no] shutdown
```

IES Service Configuration Commands

```
config
    — service
        — ies service-id [customer customer-id] [create] [vpn vpn-id]
            — description description-string
            — no description
            — [no] interface ip-int-name [create]
                — address {ip-address/mask | ip-address netmask} [broadcast {all-ones |
                    host-ones}]
                — no address {ip-address/mask | ip-address netmask}
                — [no] allow-directed broadcasts
                — arp-retry-timer ms-timer
                — no arp-retry-timer
                — arp-timeout seconds
                — no arp-timeout
                — bfd transmit-interval [receive receive-interval] [multiplier multiplier] [type
                    np]
                — no bfd
                — description description-string
                — no description
                — 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 remote-id
                        — [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
                    — ttl-expired [number seconds]
                    — no ttl-expired
                    — unreachables [number seconds]
                    — no unreachables
                — ip-mtu octets
                — no ip-mtu
                — [no] ipcp
```

```
            — dns ip-address [secondary ip-address]
            — dns secondary ip-address
            — no dns [ip-address] [secondary ip-address]
            — peer-ip-address ip-address
            — no peer-ip-address
        — [no] ipv6
            — address ipv6-address/prefix-length [eui-64]
            — no address ipv6-address/prefix-length
            — [no] dhcp6-relay
                — description description-string
                — [no] description
                — [no] option
                    — interface-id
                    — interface-id ascii-tuple
                    — interface-id ifindex
                    — interface-id sap-id
                    — interface-id string
                    — no interface-id
                    — [no] remote-id
                — server ipv6-address [ipv6-address...(up to 8 max)]
                — no server ipv6-address [ipv6-address...(up to 8 max)]
                — [no] shutdown
            — icmp6
                — packet-too-big [number seconds]
                — no packet-too-big
                — param-problem [number seconds]
                — no param-problem
                — time-exceeded [number seconds]
                — no time-exceeded [number seconds]
                — unreachables [number seconds]
                — no unreachables
            — neighbor ipv6-address mac-address
            — no neighbor ipv6-address
            — reachable-time seconds
            — stale-time seconds
        — l4-load-balancing {includeL4 | excludeL4}
        — no l4-load-balancing
        — [no] local-dhcp-server local-server-name
        — [no] local-proxy-arp
        — [no] mac ieee-address
        — proxy-arp-policy policy-name [policy-name...(up to 5 max)]
        — no proxy-arp-policy
        — [no] remote-proxy-arp
        — [no] sap sap-id [create]
            — accounting-policy acct-policy-id
            — no accounting-policy [acct-policy-id]
            — [no] collect-stats
            — description description-string
            — no description
            — egress
                — agg-rate-limit agg-rate [cir cir-rate]
                — no agg-rate-limit
                — filter ip ip-filter-id
                — filter ipv6 ipv6-filter-id
```

```
                    — no filter [ip ip-filter-id | ipv6 ipv6-filter-id]
                    — [no] qinq-mark-top-only
                    — qos policy-id
                    — no qos
                    — scheduler-mode {4-priority | 16-priority}
                    — [no] shaper-group shaper-group-name [create]
                — ingress
                    — agg-rate-limit agg-rate [cir cir-rate]
                    — no agg-rate-limit
                    — filter ip ip-filter-id
                    — filter ipv6 ipv6-filter-id
                    — no filter [ip ip-filter-id | ipv6 ipv6-filter-id]
                    — match-qinq-dot1p {top | bottom}
                    — no match-qinq-dot1p
                    — qos policy-id
                    — no qos
                    — scheduler-mode {4-priority | 16-priority}
                    — [no] shaper-group shaper-group-name [create]
                — [no] shutdown
            — secondary {ip-address/mask | ip-address netmask} [broadcast all-ones | host-
              ones] [igp-inhibit]
            — no secondary {ip-address/mask | ip-address netmask}
            — [no] shutdown
            — spoke-sdp sdp-id:vc-id [create]
            — no spoke-sdp sdp-id:vc-id
                — egress
                    — vc-label egress-vc-label
                    — no [egress-vc-label]
                — ingress
                    — filter ip ip-filter-id
                    — no filter
                    — vc-label ingress-vc-label
                    — no vc-label [ingress-vc-label]
                — [no] shutdown
            — static-arp ip-address ieee-address
            — no static-arp ip-address [ieee-address]
            — static-arp ieee-address unnumbered
            — no static-arp [ieee-address] unnumbered
            — [no] teid-load-balancing
            — unnumbered {ip-int-name | ip-address}
            — no unnumbered
    — service-name service-name
    — no service-name
    — [no] shutdown
```

Routed VPLS Commands

```
config
    — service
        — ies service-id
            — interface ip-interface-name [create]
            — no interface ip-interface-name
                — vpls service-name
                — no vpls
                    — ingress
                        — v4-routed-override-filter ipv4-filter-id
                        — no v4-routed-override-filter
                        — v6-routed-override-filter ipv6-filter-id
                        — no v6-routed-override-filter
```

VRRP Commands

```
config
    — service
        — ies service-id [customer customer-id] [create] [vpn vpn-id]
            — [no] interface ip-int-name
                — [no] ipv6
                    — vrrp virtual-router-id [owner]
                    — no vrrp virtual-router-id
                        — [no] backup ip-address
                        — init-delay seconds
                        — no init-delay
                        — mac mac-address
                        — no mac
                        — [no] master-int-inherit
                        — message-interval {[seconds] [milliseconds milliseconds]}
                        — no message-interval
                        — [no] ping-reply
                        — policy vrrp-policy-id
                        — no policy
                        — [no] preempt
                        — priority base-priority
                        — no priority
                        — [no] shutdown
                        — [no] standby-forwarding
                        — [no] telnet-reply
                        — [no] traceroute-reply
                — vrrp virtual-router-id [owner]
                — no vrrp virtual-router-id
                    — authentication-key [authentication-key | hash-key] [hash | hash2]
                    — no authentication-key
                    — [no] backup ip-address
                    — [no] bfd-enable [base | service-id ] interface interface-name dst-ip
                          ip-address
                    — init-delay seconds
                    — no init-delay
```

```
— mac ieee-address
— no mac
— [no] master-int-inherit
— message-interval {[seconds] [milliseconds milliseconds]}
— no message-interval
— [no] ping-reply
— policy vrrp-policy-id
— no policy
— [no] preempt
— priority priority
— no priority
— [no] shutdown
— [no] ssh-reply
— [no] standby-forwarding
— [no] telnet-reply
— [no] traceroute-reply
```

IES Security Zone Configuration Commands

```
config
    — service
        — ies service-id [customer customer-id] [create]
        — no ies service-id
            — zone {zone-id | zone-name} [create]
            — no zone {zone-id | zone-name}
                — abort
                — begin
                — commit
                — description description-string
                — no description
                — inbound
                    — limit
                        — concurrent-sessions {tcp | udp | icmp | other} sessions
                        — no concurrent-sessions {tcp | udp | icmp | other}
                — [no] interface interface-name
                    — [no] shutdown
                — log {log-id | name}
                — no log
                — name zone-name
                — no name
                — nat
                    — pool pool-id [create]
                    — no pool pool-id
                        — description description-string
                        — no description
                        — direction {zone-outbound | zone-inbound | both}
                        — no direction
                        — entry entry-id [create]
                        — no entry entry-id
                            — ip-address ip-address [to ip-address] interface ip-int-
                              name
```

```
                                        — no ip-address
                                        — port port [to port]
                                        — no port
                                    — name pool-name
                                    — no name
                            — outbound
                                — limit
                                    — concurrent-sessions {tcp | udp | icmp | other} sessions
                                    — no concurrent-sessions {tcp | udp | icmp | other}
                            — policy {policy-id | policy-name}
                            — no policy
                            — [no] shutdown
```

Show Commands

```
show
    — service
        — customer [customer-id]
        — egress-label start-label [end-label]
        — id service-id
            — all
            — arp [ip-address] | [mac ieee-address] | [sap sap-id] | [interface ip-int-name]
            — base
            — dhcp
                — statistics [interface interface-name | ip-address]
                — summary [interface interface-name | saps]
            — interface [{[ip-address | ip-int-name] [interface-type] [detail] [family]} | summary]
            — sap [sap-id] [detail]
        — ingress-label start-label [end-label]
        — sap-using [sap sap-id]
        — sap-using interface [ip-address | ip-int-name]
        — sap-using description
        — sap-using [ingress | egress] atm-td-profile td-profile-id
        — sap-using [ingress | egress] scheduler-mode {4-priority | 16-priority}
        — sap-using [ingress] filter filter-id
        — sap-using [ingress | egress] qos-policy qos-policy-id
        — service-using [ies] [customer customer-id]
```

Clear Commands

```
clear
    — service
        — id service-id
            — dhcp
                — statistics [ip-int-name | ip-address]
            — dhcp6
                — statistics [ip-int-name | ip-address]
```

Debug Commands

```
debug
    — service
        — id service-id
```

Command Descriptions

- IES Management Configuration Commands
- IES Service Configuration Commands
- Show Commands
- Clear Commands
- Debug Commands

IES Management Configuration Commands

- Generic Commands
- IES Global Commands
- IES Management Interface Commands
- IES Management SAP Commands

Generic Commands

description

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

Context config>service>ies
config>service>ies>interface
config>service>ies>if>dhcp
config>service>ies>if>ipv6>dhcp6-relay
config>service>ies>if>sap
config>service>ies>zone
config>service>ies>zone>nat>pool

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

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

The **dhcp** and **dhcp6-relay** commands do not apply to IES when used for in-band management.

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

shutdown

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

Context config>service>ies
config>service>ies>interface
config>service>ies>if>dhcp
config>service>ies>if>ipv6>dhcp6-relay
config>service>ies>if>sap
config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many objects must be shut down before they may be deleted. Many entities must be explicitly enabled using the **no shutdown** command.

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

The **dhcp** and **dhcp6-relay** commands do not apply to IES when used for in-band management.

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

Special Cases

IES — the default administrative status of an IES service is down. While the service is down, its associated interface is operationally down.

For example, if:

1) An IES service is operational and its associated interface is shut down

2) The IES service is administratively shut down and brought back up

3) The interface that is shut down remains in the administrative shutdown state

A service is regarded as operational provided that one IP interface is operational.

IES IP Interfaces — when the IP interface is shut down, it enters the administratively and operationally down states. For a SAP bound to the IP interface, no packets are transmitted out of the SAP and all packets received on the SAP are dropped and the packet discard counter is incremented.

IES Global Commands

ies

Syntax **ies** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*]
no ies *service-id*

Context config>service

Description This command enables Internet Enhanced Service (IES). On the 7705 SAR, IES is used for direct IP connectivity between customer access points as well as in-band management of the 7705 SAR over ATM links.

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 shut down and deleted.

Parameters *service-id* — uniquely identifies a 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 or name used for every 7705 SAR on which this service is defined.

Values 1 to 2147483647 or *service-name*

customer-id — specifies the customer ID number to be associated with the service. This parameter is required on service creation and is optional for service editing or deleting.

Values 1 to 2147483647

vpn-id — specifies the VPN ID number, which allows you to identify virtual private networks (VPNs) by a VPN identification number. If this parameter is not specified, the VPN ID uses the service ID number. This parameter is not the same as the VRF ID used with VPRN services.

Values 1 to 2147483647

service-name

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

Context config>service>ies

Description This command configures a service name that can be used in other configuration commands and show commands that reference the service.

Parameters *service-name* — up to 64 characters

IES Management Interface Commands

interface

Syntax **interface** *ip-int-name* [**create**]
no interface *ip-int-name*

Context config>service>ies

Description This command creates a logical IP routing interface for an Internet Enhanced Service (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, creates and maintains IP routing interfaces within IES service IDs. The interface command can be executed in the context of an IES service ID. Two SAPs can be assigned to a single group interface.

Interface names are case-sensitive and must be unique within the group of IP interfaces defined for **config router interface** and **config service ies interface** (that is, the network core router instance). Interface names cannot 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.

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 the IP interface and all the associated configurations. The interface must be administratively shut down before issuing the no interface command. The IP interface must be shut down before the SAP on that interface can be removed.

Default no interface

Parameters *ip-int-name* — the name of the IP interface. Interface names must be unique within the group of IP interfaces defined for the network core router instance. 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 within double quotes.

Values 1 to 32 characters (must start with a letter)

If the *ip-int-name* already exists, the context is changed to maintain that IP interface. If the *ip-int-name* already exists as an IP interface defined within the `config router` command, an error will occur and the context will not be changed to that IP interface. If the *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*}
no address

Context config>service>ies>interface

Description This command assigns an IP address and IP subnet to an IES IP interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each IP interface. An IP address and a mask combine to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. The IP prefix cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7705 SAR.

The IP address for the interface can be entered in either CIDR (classless inter-domain routing) notation or traditional dotted-decimal notation. **Show** commands display CIDR notation and are stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

The **no** form of the command removes the IP address assignment from the IP interface. The **no** form of this command can only be performed when the IP interface is administratively shut down. Shutting down the IP interface brings the interface operationally down.

Default no address

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.

Values 1.0.0.0 to 223.255.255.255

/ — the forward slash is a parameter delimiter that 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* parameter. If a forward slash does not immediately follow the *ip-address*, a dotted-decimal mask must follow the prefix.

mask — 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* parameter. The *mask* parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address determines the host portion of the IP address.

Values 1 to 32 (mask length of 32 is reserved for system IP addresses)

netmask — the subnet mask in dotted-decimal notation

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

bfd

Syntax **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*] [**type np**]
no bfd

Context config>service>ies>interface
config>service>ies>if>ipv6

Description This command configures the time interval in which BFD control messages are transmitted and received on the interface. The *multiplier* parameter specifies the number of consecutive BFD messages that must be missed by the peer node before the BFD session closes and the upper layer protocols (OSPF, IS-IS, BGP, PIM) are notified of the fault.

Default no bfd

Parameters *transmit-interval* — the number of milliseconds between consecutive BFD sent messages

Values 10 to 100000

Default 100

receive-interval — the number of milliseconds between consecutive BFD received messages

Values 10 to 100000

Default 100

multiplier — the number of consecutive BFD messages that must be missed before the interface is brought down

Values 3 to 20

Default 3

type np — controls the value range of the *transmit-interval* and *receive-interval* parameters. If the **type np** option is not specified, the range of the *transmit-interval* and *receive-interval* parameter values is from 100 ms to 100000 ms. If the **type np** option is specified, the range of the *transmit-interval* and *receive-interval* parameter values is from 10 ms to 1000 ms, with the restriction that the maximum receiving detection time for the missing BFD packets must be less than or equal to 3000 ms. The maximum receiving detection time is the *receive-interval* parameter multiplied by the *multiplier* parameter.

Note: The BFD session must be disabled before the **type np** parameter can be changed.

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>ies>interface

Description This command configures the IP maximum transmit unit (packet size) for this interface.

The **no** form of the command returns the default value.

Parameters *octets* — the MTU for the interface

Values 512 to 2048

IES Management SAP Commands

sap

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

Context config>service>ies>interface

Description This command creates a SAP within an IES service. Each SAP must be unique.

All SAPs must be explicitly created with the **create** keyword. 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.

A SAP can only be associated with a single service. The SAP is owned by the service in which it was created. An IES SAP can only be defined on an ATM port or IMA group that has been configured as an access port in the **config>port** *port-id* context using the `mode access` command. Fractional TDM ports are always access ports. Refer to the 7705 SAR OS Interface Configuration Guide for information on access ports.

If a port is shut down, all SAPs on that port become operationally down. When a service is shut down, 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 sap

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

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

ingress

Syntax **ingress**

Context config>service>ies>if>sap

Description This command enables access to the context to associate ingress filter policies with the SAP.

If an ingress filter is not defined, no filtering is performed.

filter ip

Syntax **filter ip** *ip-filter-id*
no filter
no filter [**ip** *ip-filter-id*]

Context config>service>ies>if>sap>ingress

Description This command associates an IP filter policy with an ingress SAP. Filter policies control the forwarding and dropping of packets based on the IP match criteria. Only one filter ID can be specified.

The filter policy must already be defined before the filter command is executed. If the filter policy does not exist, the operation fails and an error message is returned. Filters applied to the ingress SAP apply to all IP packets on the SAP.

The **no** form of this command removes any configured filter ID association with the SAP.

Default no filter

Parameters *ip-filter-id* — specifies the IP filter policy. The filter ID or filter name must already exist within the created IP filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

Note: For information on configuring IP filter IDs, refer to the 7705 SAR OS Router Configuration Guide, “Filter Policies”.

atm

Syntax **atm**

Context config>service>ies>if>sap

Description This command enables access to the context to configure ATM-related attributes. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:

- configuring ATM port or ATM port-related functionality on T1/E1 ASAP Adapter cards on a 7705 SAR-8 or 7705 SAR-18, or on T1/E1 ports on a 7705 SAR-F, or 7705 SAR-M (variants with T1/E1 ports)
- configuring ATM-related configuration for ATM-based SAPs that exist on T1/E1 ASAP Adapter cards or T1/E1 ports on a 7705 SAR-8 or 7705 SAR-18, or on T1/E1 ports on a 7705 SAR-F, or 7705 SAR-M (variants with T1/E1 ports)

If ATM functionality is not supported for a given context, the command returns an error.

encapsulation

Syntax **encapsulation** *atm-encap-type*

Context config>service>ies>if>sap>atm

Description This command configures an ATM VC SAP for encapsulation in accordance with RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. This command is only supported in the IP over ATM management context.

The only supported encapsulation type is aal5mux-ip.

Ingress traffic that does not match the configured encapsulation is dropped.

Default aal5mux-ip

Parameters *atm-encap-type* — aal5mux-ip (routed IP encapsulation for a VC multiplexed circuit as defined in RFC 2684)

egress

Syntax **egress**

Context config>service>ies>if>sap>atm

Description This command provides access to the context to configure egress ATM traffic policies for the SAP.

ingress

Syntax **ingress**

Context config>service>ies>if>sap>atm

Description This command provides access to the context to configure ingress ATM traffic policies for the SAP.

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
no traffic-desc

Context config>service>ies>if>sap>atm>egress
config>service>ies>if>sap>atm>ingress

Description This command assigns an ATM traffic descriptor profile to an egress or ingress SAP.

When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction.

When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.

Note: Proper configuration of the traffic descriptor profiles is essential for proper operation of the IES SAP. If no profile is assigned, the default UBR service category is assumed. All IES 7705 SAR traffic is scheduled; no shaping is supported in this mode. To ensure that IP traffic transported over the IES SAP is prioritized fairly, ATM layer traffic descriptors should be assigned.

The **no** form of the command reverts the traffic descriptor to the default traffic descriptor profile.

Default The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created ATM VC SAPs.

Parameters *traffic-desc-profile-id* — specifies a defined traffic descriptor profile (for information on defining traffic descriptor profiles, see the 7705 SAR OS Quality of Service Guide)

Values 1 to 1000

oam

Syntax **oam**

Context config>service>ies>if>sap>atm

Description This command enables the context to configure OAM functionality for an IES SAP.

The T1/E1 ASAP Adapter cards support F4 and F5 end-to-end OAM functionality (AIS, RDI, Loopback).

alarm-cells

Syntax [**no**] **alarm-cells**

Context config>service>ies>if>sap>atm>oam

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC terminations to monitor and report the status of their connection by propagating fault information through the network and by driving the PVCC operational status.

Layer 2 OAM AIS/RDI cells that are received on the IES SAP will cause the IP interface to be disabled.

The **no** command disables alarm-cells functionality for the SAP. When alarm-cells functionality is disabled, OAM cells are not generated as result of the SAP going into the operationally down state.

Default enabled

IES Service Configuration Commands

- Generic Commands
- IES Global Commands
- IES Service Interface Commands
- IES Service IPv6 Commands
- IES Service VRRP Commands
- IES Service SAP Commands
- IES Service Spoke SDP Commands
- Routed VPLS Commands
- IES Service Security Zone Configuration Commands

IES Service Interface Commands

interface

Syntax [**no**] **interface** *ip-int-name* [**create**]

Context config>service>ies

Description This command creates a logical IP routing interface for Internet Enhanced Service (IES). When 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, creates and maintains 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 the default routing table. Two SAPs can be assigned to a single group interface.

Interface names are case-sensitive and must be unique within the group of IP interfaces defined for **config router interface** and **config service ies interface** (that is, the network core router instance). Interface names cannot be in the dotted-decimal format 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.

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 the IP interface and all the associated configurations. The interface must be administratively shut down before issuing the no interface command. The IP interface must be shut down before the SAP on that interface can be removed.

Default no interface

Parameters *ip-int-name* — the name of the IP interface. Interface names must be unique within the group of IP interfaces defined for the network core router instance. 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 within double quotes.

Values 1 to 32 characters (must start with a letter)

If the *ip-int-name* already exists, the context is changed to maintain that IP interface. If the *ip-int-name* already exists as an IP interface defined within the `config router` command, an error will occur and the context will not be changed to that IP interface. If the *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*} [**broadcast** {**all-ones** | **host-ones**}]
no address {*ip-address*/*mask* | *ip-address netmask*}

Context config>service>ies>interface

Description This command assigns an IP address, IP subnet, and broadcast address format to an IES IP interface.

An IP address must be assigned to each IES IP interface. An IP address and a mask combine to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. The IP prefix cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7705 SAR.

The IP address for the interface can be entered in either CIDR (classless inter-domain routing) notation or traditional dotted-decimal notation. Show commands display CIDR notation and are stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

The **no** form of the command removes the IP address assignment from the IP interface. The **no** form of this command can only be performed when the IP interface is administratively shut down. Shutting down the IP interface brings the interface operationally down.

Default no address

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.

Values 1.0.0.0 to 223.255.255.255

/ — the forward slash is a parameter delimiter that 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* parameter. If a forward slash does not immediately follow the *ip-address*, a dotted-decimal mask must follow the prefix.

mask — 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* parameter. The *mask* parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address determines the host portion of the IP address.

Values 1 to 32 (mask length of 32 is reserved for system IP addresses)

netmask — the subnet mask in dotted-decimal notation

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

broadcast — 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 indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to the default 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.

all-ones — specifies that the broadcast address used by the IP interface for this IP address is 255.255.255.255 (also known as the local broadcast)

host-ones — specifies that the broadcast address used by the IP interface for this IP address is the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the IP address and mask with all host bits set to 1. This IP address is the default broadcast address used by an IP interface.

allow-directed broadcasts

Syntax [**no**] **allow-directed broadcasts**

Context config>service>ies>interface

Description 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 of another IP interface. The **allow-directed-broadcasts** command on an IP interface enables or disables the transmission of packets destined for the subnet broadcast address of the egress IP interface.

When enabled, a frame destined for the local subnet on this IP interface is sent as a subnet broadcast out this interface.

Note: Allowing directed broadcasts is a well-known mechanism used for denial-of-service attacks.

By default, directed broadcasts are not allowed and are discarded at this egress IP interface.

The **no** form of the command disables directed broadcasts forwarding out of the IP interface.

Default no allow-directed broadcasts

arp-retry-timer

Syntax **arp-retry-timer** *ms-timer*
no arp-retry-timer

Context config>service>ies>interface

Description This command specifies the length of time, in 100s of milliseconds, that the system waits before reissuing a failed ARP request.

The **no** form of the command resets the interval to the default value.

Note: The ARP retry default value of 5000 ms is intended to protect CPU cycles on the 7705 SAR, especially when it has a large number of interfaces. Configuring the ARP retry timer to a value shorter than the default should be done only on mission-critical links, such as uplinks or aggregate spoke SDPs transporting mobile traffic; otherwise, the retry interval should be left at the default value.

Default 50 (in 100s of ms)

Parameters *ms-timer* — the time interval, in 100s of milliseconds, the system waits before retrying a failed ARP request

Values 1 to 300

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>ies>interface

Description This command configures the minimum interval, in seconds, that an ARP entry learned on the IP interface is 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 the **arp-timeout** value is set to 0 s, ARP aging is disabled.

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

Note: The 7705 SAR will attempt to refresh an ARP entry 30 s prior to its expiry. This refresh attempt occurs only if the ARP timeout is set to 45 s or more.

Default no arp-timeout

Parameters *seconds* — the minimum number of seconds a learned ARP entry is stored in the ARP table, expressed as a decimal integer. A value of 0 specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 to 65535

Default 14400 s (4 h)

bfd

Syntax **bfd** {*transmit-interval*} [**receive** *receive-interval*] [**multiplier** *multiplier*] [**type np**]
no bfd

Context config>service>ies>interface
config>service>ies>if>ipv6

Description This command configures the time interval in which BFD control messages are transmitted and received on the interface. The *multiplier* parameter specifies the number of consecutive BFD messages that must be missed by the peer node before the BFD session closes and the upper layer protocols (OSPF, IS-IS, BGP, PIM) are notified of the fault.

Default no bfd

Parameters *transmit-interval* — the number of milliseconds between consecutive BFD sent messages

Values 10 to 100000

Default 100

receive-interval — the number of milliseconds between consecutive BFD received messages

Values 10 to 100000

Default 100

multiplier — the number of consecutive BFD messages that must be missed before the interface is brought down

Values 3 to 20

Default 3

type np — controls the value range of the *transmit-interval* and *receive-interval* parameters. If the **type np** option is not specified, the range of the *transmit-interval* and *receive-interval* parameter values is from 100 ms to 100000 ms. If the **type np** option is specified, the range of the *transmit-interval* and *receive-interval* parameter values is from 10 ms to 1000 ms, with the restriction that the maximum receiving detection time for the missing BFD packets must be less than or equal to 3000 ms. The maximum receiving detection time is the *receive-interval* parameter multiplied by the *multiplier* parameter.



Note: The BFD session must be disabled before the **type np** parameter can be changed.

dhcp

Syntax **dhcp**

Context config>service>ies>interface

Description This command enables the context to configure DHCP parameters.

gi-address

Syntax **gi-address** *ip-address* [**src-ip-addr**]
no gi-address

Context config>service>ies>if>dhcp

Description This command configures the gateway interface address for the DHCP Relay Agent. By default, the GIADDR used in the relayed DHCP packet is the primary address of an interface. Specifying the GIADDR allows the user to choose a secondary address.

Default no gi-address

Parameters *ip-address* — the IP address of the gateway interface in dotted-decimal notation

Values a.b.c.d (host bits must be 0)

src-ip-addr — specifies that the GIADDR is to be used as the source IP address for DHCP relay packets

option

Syntax [**no**] **option**

Context config>service>ies>if>dhcp

Description This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 suboptions.

The **no** form of this command returns the system to the default.

Default no option

action

Syntax **action** {**replace** | **drop** | **keep**}
no action

Context config>service>ies>if>dhcp>option

Description This command configures the Relay Agent Information Option (Option 82) processing.

The **no** form of this command returns the system to the default value.

Default keep

Parameters **replace —** in the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (toward the user), the Option 82 field is stripped (in accordance with RFC 3046).

drop — the DHCP packet is dropped if an Option 82 field is present, and a counter is incremented

keep — the existing information is kept in the packet and the router does not add any additional information. In the downstream direction, the Option 82 field is not stripped and is forwarded toward the client.

The behavior is slightly different in the case of Vendor Specific Options (VSOs). When the **keep** parameter is specified, the router will insert its own VSO into the Option 82 field. This will only be done if the incoming message already has an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field. In this 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>ies>if>dhcp>option

Description This command sends either an ASCII tuple or the interface index (If Index) on the specified SAP ID in the **circuit-id** suboption of the DHCP packet.

If disabled, the **circuit-id** suboption of the DHCP packet is left empty.

The **no** form of the command returns the system to the default.

Default ascii-tuple

Parameters **ascii-tuple —** specifies that the ASCII-encoded concatenated tuple, which consists of the access node identifier, service ID, and interface name, separated by "/", will be used

ifindex — specifies that the interface index will be used. The If Index of a router interface can be displayed using the command **show>router>if>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 already included in **ascii-tuple**. The format is supported on dot1q and qinq ports only. Therefore, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

remote-id

Syntax **remote-id** [**mac** | **string** *string*]
no remote-id

Context config>service>ies>if>dhcp>option

Description This command sends the MAC address of the remote end (typically the DHCP client) in the **remote-id** suboption of the DHCP packet. This command identifies the host at the other end of the circuit.

If disabled, the **remote-id** suboption of the DHCP packet is left empty.

The **no** form of this command returns the system to the default.

Default remote-id

Parameters **mac —** specifies that the MAC address of the remote end is encoded in the suboption

string — the remote ID

vendor-specific option

Syntax [**no**] **vendor-specific-option**

Context config>service>ies>if>dhcp>option

Description This command configures the **vendor-specific** suboption of the DHCP relay packet.

client-mac-address

Syntax [**no**] **client-mac-address**

Context config>service>ies>if>dhcp>option>vendor-specific-option

Description This command enables the sending of the MAC address in the vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address.

sap-id

Syntax [**no**] **sap-id**

Context config>service>ies>if>dhcp>option>vendor-specific-option

Description This command enables the sending of the SAP ID in the vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID.

service-id

Syntax [**no**] **service-id**

Context config>service>ies>if>dhcp>option>vendor-specific-option

Description This command enables the sending of the service ID in the vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID.

string

Syntax **string** *text*
no string

Context config>service>ies>if>dhcp>option>vendor-specific-option

Description This command specifies the string in the vendor-specific suboption of the DHCP relay packet.

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

Default no string

Parameters *text* — any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, the entire string must be enclosed within double quotes.

system-id

Syntax [**no**] **system-id**

Context config>service>ies>if>dhcp>option>vendor-specific-option

Description This command specifies whether the system ID is encoded in the vendor-specific suboption of the DHCP relay packet.

server

Syntax **server** *server1* [*server2*...(up to 8 max)]
no server

Context config>service>ies>if>dhcp>option

Description This command specifies a list of servers where requests will be forwarded. The list of servers can be entered either as 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, 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* — the DHCP server IP address

trusted

Syntax [**no**] **trusted**

Context config>service>ies>if>dhcp>option

Description As specified in RFC 3046, *DHCP Relay Agent Information Option*, a DHCP request where the giaddr is 0.0.0.0 and that 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 router) will modify the request's giaddr to be equal to the ingress interface and forward the request.

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, and therefore there is no reason to enable the **trusted** option.

The **no** form of this command returns the system to the default.

Default not enabled

icmp

Syntax **icmp**

Context config>service>ies>interface

Description This command enables access to the context to configure Internet Control Message Protocol (ICMP) parameters on a network IP interface. ICMP is a message control and error reporting protocol that also provides information relevant to IP packet processing.

mask-reply

Syntax [**no**] **mask-reply**

Context config>service>ies>if>icmp

Description This command enables or disables responses to 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.

The **no** form of the command disables replies to ICMP mask requests on the router interface.

Default mask-reply

ttl-expired

Syntax **ttl-expired** [*number seconds*]
no ttl-expired

Context config>service>ies>if>icmp

Description This command configures the rate that ICMP Time To Live (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-s time interval.

The **no** form of the command disables the generation of TTL expired messages.

Default ttl-expired 100 10—maximum of 100 TTL expired message in 10 s

Parameters *number* — the maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. The *seconds* parameter must also be specified.

Values 10 to 100

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 to 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>ies>if>icmp

Description 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 unreachable messages on the router interface. The rate at which ICMP unreachable messages are issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of destination unreachable messages that can be issued on the interface for a specified time interval.

By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 100 per 10-s time interval.

The **no** form of the command disables the generation of ICMP destination unreachable messages on the router interface.

Default unreachables 100 10—maximum of 100 unreachable messages in 10 s

Parameters *number* — the maximum number of ICMP unreachable messages to send, expressed as a decimal integer. The *seconds* parameter must also be specified.

Values 10 to 100

seconds — the time frame, in seconds, used to limit the *number* of ICMP unreachable messages that can be issued, expressed as a decimal integer

Values 1 to 60

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>ies>interface

Description This command configures the IP maximum transmit unit (packet size) for this interface.

The **no** form of the command returns the default value.

Parameters *octets* — the MTU for the interface

Values 512 to 9732

ipcp

Syntax [**no**] **ipcp**

Context config>service>ies>interface

Description This command enables the context to configure IPCP. Within this context, IPCP extensions can be used to signal the remote IP address and DNS IP address to the PPP peer over the PPP/MLPPP interface. This command is only applicable if the associated SAP is a PPP/MLPPP interface.

dns

Syntax **dns** *ip-address* [**secondary** *ip-address*]
dns secondary *ip-address*
no dns [*ip-address*] [**secondary** *ip-address*]

Context config>service>ies>if>ipcp

Description This command defines the DNS address(es) to be assigned to the far end of the associated PPP/MLPPP link via IPCP extensions. This command is only applicable if the associated SAP or port is a PPP/MLPPP interface with an IPCP encapsulation.

The **no** form of the command deletes the specified primary DNS address, secondary DNS address, or both addresses from the IPCP extension **peer-ip-address** configuration.

Default no dns

Parameters *ip-address* — specifies a unicast IPv4 address for the primary DNS server to be signaled to the far end of the associated PPP/MLPPP link via IPCP extensions

secondary *ip-address* — specifies a unicast IPv4 address for the secondary DNS server to be signaled to the far end of the associated PPP/MLPPP link via IPCP extensions

peer-ip-address

Syntax **peer-ip-address** *ip-address*
no peer-ip-address

Context config>service>ies>if>ipcp

Description This command defines the remote IP address to be assigned to the far end of the associated PPP/MLPPP link via IPCP extensions. This command is only applicable if the associated SAP or port is a PPP/MLPPP interface with an IPCP encapsulation.

The **no** form of the command deletes the IPCP extension peer-ip-address configuration.

Default no peer-ip-address (0.0.0.0)

Parameters *ip-address* — a unicast IPv4 address to be signaled to the far end of the associated PPP/MLPPP link by IPCP extensions

l4-load-balancing

Syntax **l4-load-balancing** {**includeL4** | **excludeL4**}
no l4-load-balancing

Context config>service>ies>interface

Description This command configures Layer 4 load balancing at the interface level. Configuration must be done on the ingress network interface (that is, the interface on the node that the packet is received on). When enabled, Layer 4 source and destination port fields of incoming TCP/UDP packets are included in the hashing calculation to randomly determine the distribution of packets.

You can add additional fields to generate more randomness and more equal distribution of packets with the teid-load-balancing command.

The default configuration on the interface is to match the Layer 4 load balancing configuration in the **config>system** context. Using this command to modify Layer 4 load-balancing configuration on an interface overrides the system-wide load-balancing settings for that interface.

Parameters **includeL4 —** layer 4 source and destination port fields are included in the hashing calculation

excludeL4 — layer 4 source and destination port fields are not included in the hashing calculation

local-dhcp-server

Syntax [**no**] **local-dhcp-server** *local-server-name*

Context config>service>ies>interface

Description This command associates the interface with a local DHCP server configured on the system. A routed VPLS interface may not be associated with a local DHCP server.

The **no** form of the command removes the association of the interface with the local DHCP server.

Default n/a

Parameters *local-server-name* — the name of the local DHCP server

Values up to 32 alphanumeric characters

local-proxy-arp

Syntax [**no**] **local-proxy-arp**

Context config>service>ies>interface

Description This command enables local proxy ARP on the interface.

Local proxy ARP allows the 7705 SAR to respond to ARP requests received on an interface for an IP address that is part of a subnet assigned to the interface. The router responds to all requests for IP addresses within the subnet with its own MAC address and forwards all traffic between the hosts in the subnet.

Local proxy ARP is used on subnets where hosts are prevented from communicating directly.

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-address*

Context config>service>ies>interface

Description This command assigns a specific MAC address to an IES IP interface.

The **no** form of the command returns the MAC address to the default value.

Default the physical MAC address associated with the Ethernet interface on which the SAP is configured (default MAC address assigned to the interface by the system)

Parameters *ieee-address* — 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-policy

Syntax **proxy-arp-policy** *policy-name* [*policy-name*...(up to 5 max)]
no proxy-arp-policy

Context config>service>ies>interface

Description This command enables proxy ARP on the interface and specifies an existing policy statement that controls the flow of routing information by analyzing match and action criteria. The policy statement is configured in the **config>router>policy-options** context (refer to the 7705 SAR OS Router Configuration Guide, "Route Policy Command Reference, Route Policy Options"). When proxy ARP is enabled, the 7705 SAR responds to ARP requests on behalf of another device.

Default no proxy-arp-policy

Parameters *policy-name* — the route policy statement 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 policy statement must already be defined.

remote-proxy-arp

Syntax [**no**] **remote-proxy-arp**

Context config>service>ies>interface

Description This command enables remote proxy ARP on the interface, allowing a router on one network to respond to ARP requests intended for another node that is physically located on another network. The router effectively pretends to be the destination node by sending an ARP response to the originating node that associates the router's MAC address with the destination node's IP address (acts as a proxy for the destination node). The router then takes responsibility for routing traffic to the real destination.

Default no remote-proxy-arp

secondary

Syntax **secondary** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**igp-inhibit**]
no secondary {*ip-address/mask* | *ip-address netmask*}

Context config>service>ies>interface

Description This command assigns an secondary IP address, IP subnet, and broadcast address format to the interface.

Default no secondary

Parameters *ip-address* — the IP address of the IP interface. The *ip-address* portion of the **secondary** 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 a.b.c.d

(1.0.0.0 to 223.255.255.255 (with support of /31 subnets)

The "/" (forward slash) is a parameter delimiter that separates the *ip-address* portion of the IP address from the *mask*, which defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "/", and the *mask*. If a forward slash does not immediately follow the *ip-address*, a dotted-decimal *mask* must follow the prefix.

mask — 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*. The *mask* 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.

Values 0 to 32

netmask — the subnet mask, in dotted-decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted-decimal mask. The *netmask* parameter indicates the complete mask that will be used in a logical "AND" function to derive the local subnet of the IP address.

Values 128.0.0.0 to 255.255.255.252

(network bits all 1 and host bits all 0)

(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 indicates 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 — specifies that 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 — 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*, or the *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 **secondary** command does not have a negation 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 **secondary** command must be executed with the **broadcast** parameter defined.

igp-inhibit — specifies that this secondary IP interface should not be recognized as a local interface by the running IGP. For OSPF and IS-IS, this means that the secondary IP interface will not be injected and used as a passive interface and will not be advertised as an internal IP interface into the IGP link state database. For RIP, this means that the secondary IP interface will not source RIP updates.

static-arp

Syntax **static-arp** *ip-address ieee-address*
no static-arp *ip-address* [*ieee-address*]
static-arp *ieee-address* **unnumbered**
no static-arp [*ieee-address*] **unnumbered**

Context config>service>ies>interface

Description This command configures a static ARP entry associating an 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 is replaced by the new MAC address.

A router interface can only have one static ARP entry configured for it.

Static ARP is used when a 7705 SAR needs to know about a device on an interface that cannot or does not respond to ARP requests. Therefore, the 7705 SAR OS configuration can specify to send a packet with a particular IP address to the corresponding ARP address.

The **no** form of the command removes a static ARP entry.

Default no static-arp

Parameters *ip-address* — the IP address for the static ARP in dotted-decimal notation

ieee-mac-address — the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff*, where *aa*, *bb*, *cc*, *dd*, *ee,* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC, and non-IEEE reserved MAC addresses.

unnumbered — specifies the static ARP MAC addresses for an unnumbered interface. Unnumbered interfaces also support dynamic ARP. If this parameter is configured, it overrides any dynamic ARP.

teid-load-balancing

Syntax **[no] teid-load-balancing**

Context config>service>ies>interface

Description This command configures TEID load balancing at the interface level. Configuration must be done on the ingress network interface (that is, the interface on the node that the packet is received on). The TEID attribute is included in the header of GTP (general packet radio system tunneling protocol) packets. When TEID load balancing is enabled, the TEID field of incoming TCP/UDP packets is included in the hashing calculation to randomly determine the distribution of packets.

You can add additional fields to generate more randomness and more equal distribution of packets with the l4-load-balancing command.

Default no teid-load-balancing

unnumbered

Syntax **unnumbered** {*ip-int-name* | *ip-address*}
no unnumbered

Context config>service>ies>interface

Description This command configures an IP interface as an unnumbered interface and specifies an IP address or interface name to be used for the interface. Unnumbered interfaces are point-to-point interfaces that are not explicitly configured with a dedicated IP address and subnet; instead, they borrow (or link to) an IP address from another interface on the system (the system IP address, another loopback interface, or any other numbered interface) and use it as the source IP address for packets originating from the interface.

By default, no IP address exists on an IP interface until it is explicitly created.

The **no** form of the command removes the IP address assignment from the IP interface.

Default no unnumbered

Parameters *ip-int-name* | *ip-address* — the IP interface name or address to associate with the unnumbered IP interface

Values *ip-int-name*: 1 to 32 characters (must start with a letter)
ip-address: a.b.c.d

IES Service IPv6 Commands

ipv6

Syntax [**no**] **ipv6**

Context config>service>ies>interface

Description 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 This command assigns an IPv6 address to the IES interface.

Default n/a

Parameters *ipv6-address/prefix-length* — the IPv6 address on the interface

Values

ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
prefix-length:	1 to 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 the MAC address on Ethernet interfaces.

dhcp6-relay

Syntax [**no**] **dhcp6-relay**

Context config>service>ies>if>ipv6

Description This command enables the context to configure DHCPv6 Relay parameters for the IES interface.

option

Syntax [**no**] **option**

Context config>service>ies>if>ipv6>dhcp6-relay

Description This command enables the context to configure DHCPv6 Relay information options.

interface-id

Syntax **interface-id**
interface-id ascii-tuple
interface-id ifindex
interface-id sap-id
interface-id string
no interface-id

Context config>service>ies>if>ipv6>dhcp6-relay>option

Description This command enables the sending of interface ID options in the DHCPv6 Relay packet.

Default ascii-tuple

Parameters **ascii-tuple —** specifies that the ASCII-encoded concatenated tuple, which consists of the access node identifier, service ID, and interface name, separated by "/", will be used

ifindex — specifies that the interface index will be used. The If Index of a router interface can be displayed using the command **show>router>if>detail**.

sap-id — specifies that the SAP ID will be used

string — specifies that a string of up to 32 printable, 7-bit ASCII characters, will be used. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

remote-id

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

Context config>service>ies>if>ipv6>dhcp6-relay>option

Description This command enables the sending of the remote ID option in the DHCPv6 Relay packet. The client DHCP Unique Identifier (DUID) is used as the remote ID.

server

Syntax **server** *ipv6-address* [*ipv6-address*...(up to 8 max)]
no server *ipv6-address* [*ipv6-address*...(up to 8 max)]

Context config>service>ies>if>ipv6>dhcp6-relay

Description This command specifies a list of servers where DHCPv6 requests will be forwarded. The list of servers can be entered either as IP addresses or fully qualified domain names. At least one server must be specified in order for DHCPv6 Relay to work. If there are multiple servers, the request is forwarded to all of them. A maximum of eight servers can be configured.

Default n/a

Parameters *ipv6-address* — the IPv6 addresses of the DHCP servers

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 to FFFF]H
d: [0 to 255]D

icmp6

Syntax **icmp6**

Context config>service>ies>if>ipv6

Description This command enables the context to configure ICMPv6 parameters on the IES interface.

packet-too-big

Syntax **packet-too-big** [*number seconds*]
no packet-too-big

Context config>service>ies>if>ipv6>icmp6

Description This command configures the rate for ICMPv6 packet-too-big messages.

The **no** form of the command disables the sending of ICMPv6 packet-too-big messages.

Default 100 10

Parameters *number* — the maximum number of packet-too-big messages to send, expressed as a decimal integer, in the time frame specified by the *seconds* parameter

Values 10 to 1000

seconds — the time frame, in seconds, used to limit the number of packet-too-big messages that can be issued, expressed as a decimal integer

Values 1 to 60

param-problem

Syntax **param-problem** [*number seconds*]
no param-problem

Context config>service>ies>if>ipv6>icmp6

Description This command configures the rate for ICMPv6 param-problem messages.

The **no** form of the command disables the sending of ICMPv6 param-problem messages.

Default 100 10

Parameters *number* — the maximum number of param-problem messages to send, expressed as a decimal integer, in the time frame specified by the *seconds* parameter

Values 10 to 1000

seconds — the time frame, in seconds, used to limit the number of param-problem messages that can be issued, expressed as a decimal integer

Values 1 to 60

time-exceeded

Syntax **time-exceeded** [*number seconds*]
no time-exceeded

Context config>service>ies>if>ipv6>icmp6

Description This command configures the rate for ICMPv6 time-exceeded messages.

The **no** form of the command disables the sending of ICMPv6 time-exceeded messages.

Default 100 10

Parameters *number* — the maximum number of time-exceeded messages to send, expressed as a decimal integer, in the time frame specified by the *seconds* parameter

Values 10 to 1000

seconds — the time frame, in seconds, used to limit the number of time-exceeded messages that can be issued, expressed as a decimal integer

Values 1 to 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>ies>if>ipv6>icmp6

Description This command enables and configures the rate for ICMPv6 host and network destination unreachable messages issued on the router interface.

The **no** form of the command disables the generation of ICMPv6 destination unreachables on the router interface.

Default 100 10

Parameters *number* — the maximum number of destination unreachable messages to send, expressed as a decimal integer, in the time frame specified by the *seconds* parameter

Values 10 to 1000

seconds — the time frame, in seconds, used to limit the number of destination unreachable messages that can be issued, expressed as a decimal integer

Values 1 to 60

neighbor

Syntax **neighbor** *ipv6-address mac-address*
no neighbor *ipv6-address*

Context config>service>ies>if>ipv6

Description This command configures an IPv6-to-MAC address mapping on the IES interface. Use this command if a directly attached IPv6 node does not support ICMPv6 neighbor discovery or a static address must be used. This command can only be used on Ethernet interfaces. The *ipv6-address* must be on the subnet that was configured from the IPv6 address command or a link-local address.

Parameters *ipv6-address* — the IPv6 address on the interface

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 to FFFF]H
d: [0 to 255]D

mac-address — the MAC address for the neighbor in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

reachable-time

Syntax [**no**] **reachable-time** *seconds*

Context config>service>ies>if>ipv6

Description This command specifies the time an IPv6 neighbor remains in reachable state.

Default no reachable-time

Parameters *seconds* — specifies the number of seconds that an IPv6 neighbor remains in reachable state

Values 30 to 3600

Default 30

stale-time

Syntax [**no**] **stale-time** *seconds*

Context config>service>ies>if>ipv6

Description This command specifies the time that an IPv6 neighbor cache entry remains in stale state. When the specified time elapses, the system removes the neighbor cache entry.

Default no stale-time

Parameters *seconds* — specifies the number of seconds that an IPv6 neighbor remains in stale state

Values 60 to 65535

Default 14400

IES Service VRRP Commands

vrrp

Syntax **vrrp** *virtual-router-id* [**owner**]
no vrrp *virtual-router-id*

Context config>service>ies>interface
config>service>ies>if>ipv6>vrrp

Description This command creates or edits a virtual router ID on the service IP interface. A virtual router ID is internally represented in conjunction with the IP interface name. This allows the virtual router ID to be used on multiple IP interfaces while representing different virtual router instances.

Two VRIDs can be defined on an IP interface. One, both, or none may be defined as owner.

The **no** form of this command removes the specified virtual router ID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the virtual router ID. The virtual router ID does not need to be shut down in order to remove the virtual router instance.

Default n/a

Parameters *virtual-router-id* — specifies a new virtual router ID or one that can be modified on the IP interface

Values 1 to 255

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>service>ies>if>vrrp

Description This command assigns a simple text password authentication key to generate master VRRP advertisement messages and validate received VRRP advertisement messages.

If the command is re-executed with a different password key defined, the new key will be used immediately. If a **no authentication-key** command is executed, the password authentication key is restored to the default value. The **authentication-key** command may be executed at any time.

To change the current in-use password key on multiple virtual router instances:

- identify the current master
- shut down the virtual router instance on all backups
- execute the **authentication-key** command on the master to change the password key
- execute the **authentication-key** command and **no shutdown** command on each backup

The **no** form of this command restores the default value of the key.

Default The authentication data field contains the value 0 in all octets.

Parameters *authentication-key* — identifies the simple text password used when VRRP Authentication Type 1 is enabled on the virtual router instance. Type 1 uses a string 8 octets long that is inserted into all transmitted VRRP advertisement messages and compared against all received VRRP advertisement messages. The authentication data fields are used to transmit the key.

The *authentication-key* parameter is expressed as a string consisting up to eight alpha-numeric characters. Spaces must be contained in quotation marks (" "). The quotation marks are not considered part of the string.

The string is case-sensitive and is left-justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet containing the first character. The second field holds the fifth through eighth characters. Any unspecified portion of the authentication data field is padded with the value 0 in the corresponding octet.

Values any 7-bit printable ASCII character

exceptions:	double quote	ASCII 34
	carriage return	ASCII 13
	line feed	ASCII 10
	tab	ASCII 9
	backspace	ASCII 8

hash-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 option 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** keyword 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** keyword specified.

hash2 — specifies the key is entered in a more complex encrypted form. If the **hash2** keyword is not used, the less-encrypted hash form is assumed.

backup

Syntax [**no**] **backup** *ip-address*

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command configures virtual router IP addresses for backup.

Parameters *ip-address* — specifies the destination address for backup

bfd-enable

Syntax **bfd-enable** [**base** | *service-id*] **interface** *interface-name* **dst-ip** *ip-address*
no bfd-enable

Context config>service>ies>if>vrrp

Description This commands assigns a BFD session that provides a heart-beat mechanism for the given VRRP instance. Only one BFD session can be assigned to any given VRRP instance, but multiple VRRP sessions can use the same BFD session.

BFD controls the state of the associated 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 session are set with the **bfd-enable** command under the IP interface.

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

Default n/a

Parameters *service-id* — specifies the service ID of the interface running BFD

Values 1 to 2147483648 or *service-name*

interface-name — specifies the name of the interface running BFD

ip-address — specifies the destination address to be used for the BFD session

init-delay

Syntax **init-delay** *seconds*
no init-delay

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command configures a VRRP initialization delay timer.

Default no init-delay

Parameters *seconds* — specifies the number of seconds for the initialization delay timer for VRRP

Values 1 to 65535

mac

Syntax **mac** *ieee-address*
no mac

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command assigns a specific MAC address to an IES IP interface.

The **no** form of the 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 (the default MAC address assigned to the interface, assigned by the system)

Parameters *mac-address* — specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bbcc-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.

master-int-inherit

Syntax [**no**] **master-int-inherit**

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command allows the master instance to dictate the master down timer (non-owner context only).

The master down interval is the time that the master router can be down before backup takes over. The master down interval is used to specify the master down timer. If the master down timer expires, the backup virtual router enters the master state.

Default no master-int-inherit

message-interval

Syntax **message-interval** {[*seconds*] [**milliseconds** *milliseconds*]}
no message-interval

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command sets the advertisement timer and indirectly sets the master down timer on the virtual router instance. The message-interval setting must be the same for all virtual routers with the same VRID. Any VRRP advertisement message received with an advertisement interval field different from the virtual router instance configured message-interval value will be silently discarded.

The **message-interval** command is available for both non-owner and owner virtual routers. If the **message-interval** command is not executed, the default message interval is 1 s.

The **no** form of this command restores the default message-interval value of 1 s to the virtual router instance.

Parameters *seconds* — number of seconds that will transpire before the advertisement timer expires

Values 1 to 255

Default 1

milliseconds — the time interval, in milliseconds, between sending advertisement messages. This parameter is not supported on non-redundant chassis.

Values 100 to 900

ping-reply

Syntax [**no**] **ping-reply**

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command enables the non-owner master to reply to ICMP echo requests directed to the virtual router instance IP addresses. The ping request can be received on any routed interface.

Ping must not have been disabled at the management security level (either on the parental IP interface or based on the ping source host address). When ping reply is not enabled, ICMP echo requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP echo requests regardless of the setting of the ping reply configuration.

The **ping-reply** command is only available for non-owner virtual routers. If the **ping-reply** command is not executed, ICMP echo requests to the virtual router instance IP addresses are silently discarded.

The **no** form of this command restores the default operation of discarding all ICMP echo request messages destined for the non-owner virtual router instance IP addresses.

Default no ping-reply

policy

Syntax **policy** *vrrp-policy-id*
no policy

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command associates a VRRP priority control policy with the virtual router instance (non-owner context only).

Default n/a

Parameters *vrrp-policy-id* — specifies a VRRP priority control policy. The VRRP policy ID must already exist in the system for the **policy** command to be successful.

Values 1 to 9999

preempt

Syntax [**no**] **preempt**

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command provides the ability to override an existing non-owner master with a virtual router backup that has a higher priority. Enabling preempt mode enhances the operation of the base priority and VRRP policy ID definitions on the virtual router instance. If the virtual router cannot preempt an existing non-owner master, the affect of the dynamic changing of the in-use priority is greatly diminished.

The **preempt** command is only available for non-owner VRRP virtual routers. The owner cannot be preempted because the priority of non-owners can never be higher than the owner. The owner will always preempt all other virtual routers when it is available.

Non-owner backup virtual router instances will only preempt when preempt is set and the current master has an in-use message priority value less than the backup virtual router instance in-use priority.

A master non-owner virtual router will only allow itself to be preempted when the incoming VRRP advertisement message priority field value is one of the following:

- greater than its in-use priority value
- equal to the in-use priority value, and the source IP address (primary IP address) is greater than its primary IP address

The **no** form of this command prevents a non-owner virtual router instance from preempting another, less-desirable, virtual router.

Default preempt

priority

Syntax **priority** *priority*
no priority

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command configures a specific priority value for the virtual router instance. In conjunction with the optional **policy** command, the base priority derives the in-use priority of the virtual router instance.

The **priority** command is only available for non-owner VRRP virtual routers. The priority of owner virtual router instances is permanently set to 255 and cannot be changed. For non-owner virtual router instances, if the priority command is not executed, the base priority is set to 100.

The **no** form of this command restores the default value of 100.

Parameters *priority* — specifies the priority used by the virtual router instance. If a VRRP priority control policy is not defined, the base priority will be the in-use priority for the virtual router instance.

Values 1 to 254

ssh-reply

Syntax [**no**] **ssh-reply**

Context config>service>ies>if>vrrp

Description This command enables the non-owner master to reply to SSH requests directed at the IP addresses of the virtual router instances. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parental IP interface or based on the SSH source host address). Proper login and CLI command authentication are enforced.

When the **ssh-reply** command is not enabled, SSH packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to SSH requests regardless of the SSH reply configuration.

The **ssh-reply** command is only available for non-owner VRRP virtual routers.

The **no** form of this command restores the default operation of discarding all SSH packets destined to the non-owner virtual router instance IP addresses.

Default no ssh-reply

standby-forwarding

Syntax [**no**] **standby-forwarding**

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command allows the forwarding of packets by a standby router when sent to the virtual router MAC address.

The **no** form of the command specifies that a standby router should not forward traffic sent to the virtual router MAC address. The standby router should forward traffic sent to the real MAC address of the standby router.

Default no standby-forwarding

telnet-reply

Syntax [**no**] **telnet-reply**

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command enables the non-owner master to reply to TCP port 23 Telnet requests directed at the IP addresses of the virtual router instance. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Proper login and CLI command authentication are enforced.

If the **telnet-reply** command is not enabled, TCP port 23 Telnet packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet requests regardless of the Telnet reply configuration.

The **telnet-reply** command is only available for non-owner VRRP virtual routers.

The **no** form of this command restores the default operation of discarding all Telnet packets destined for the non-owner virtual router instance IP addresses.

Default no telnet-reply

traceroute-reply

Syntax [**no**] **traceroute-reply**

Context config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description This command enables a non-owner master to reply to traceroute requests directed to the virtual router instance IP addresses. The command is valid only if the VRRP virtual router instance associated with this entry is a non-owner. A non-owner backup virtual router never responds to traceroute requests regardless of the traceroute reply status.

Default no traceroute-reply

IES Service SAP Commands

sap

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

Context config>service>ies>interface

Description This command creates a SAP within an IES service. Each SAP must be unique.

All SAPs must be explicitly created with the create keyword. If no SAPs are created within a service or an IP interface, a SAP will not exist on that object.

To edit SAP parameters, enter an existing SAP without the **create** keyword.

A SAP can only be associated with a single service. The SAP is owned by the service in which it was created. A SAP can only be defined on a port that has been configured as an access port in the **config>port** *port-id* context using the **mode access** command. Refer to the 7705 SAR OS Interface Configuration Guide, "Access Ports".

If a port is shut down, all SAPs on that port become operationally down. When a service is shut down, 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 SAP types are supported:

- PPP IPCP encapsulation of an IPv4 packet for IES (RFC 1332)
- MLPPP bundle
- LAG
- Ethernet SAPs supporting null, dot1q, and qinq

To configure an IES interface SAP that is used for a public IPSec tunnel interface, see sap in Service Interface Tunnel Commands.

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 sap

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

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>ies>if>sap

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

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

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

Default no accounting-policy

Parameters *acct-policy-id* — the accounting *policy ID a*s configured in the **config>log>accounting-policy** context

Values 1 to 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>ies>if>sap

Description This command enables accounting and statistical data collection for the SAP**.** 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 CSM. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued, the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default collect-stats

egress

Syntax **egress**

Context config>service>ies>if>sap

Description This command enables the context to configure egress SAP QoS policies and IP 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 IP filter policy is defined, no filtering is performed.

ingress

Syntax **ingress**

Context config>service>ies>if>sap

Description This command enables the context to configure ingress SAP QoS policies and IP 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 IP filter policy is defined, no filtering is performed.

agg-rate-limit

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

Context config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress

Description This command sets the aggregate rate limits (PIR and CIR) for the SAP. The *agg-rate* sets the PIR value. The *cir-rate* sets the CIR value. When aggregate rate limits are configured on a second-generation (Gen-2) adapter card, the scheduler mode must be set to 16-priority. On the third-generation (Gen-3) Ethernet adapter card, only the PIR component (*agg-rate*) of the **agg-rate-limit** can be set and the scheduler mode is always 4-priority-hqos. For information on adapter card generations, refer to the "Evolution of Ethernet Adapter Cards, Modules, and Platforms" section in the 7705 SAR OS Interface Configuration Guide.

Configuring the *cir-rate* is optional. If a *cir-rate* is not entered, then the *cir-rate* is set to its default value (0 kb/s). If a *cir-rate* has been set and the *agg-rate* is changed without re-entering the *cir-rate*, the *cir-rate* automatically resets to 0 kb/s. For example, to change the *agg-rate* from 2000 to 1500 while maintaining a *cir-rate* of 500, use the command **agg-rate-limit** 1500 **cir** 500.

The *cir-rate* for this command is blocked on Gen-3 Ethernet adapter cards and Gen-3 platforms with Ethernet ports, such as the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X. Gen-3 Ethernet adapter cards and Gen-3 platforms with Ethernet ports support 4-priority-hqos scheduling mode. Gen-3 platforms with a TDM MDA support 4-priority scheduling mode.

If the specified SAP is a LAG SAP, then *agg-rate* can be configured regardless of the scheduler mode setting, and *cir-rate* can be configured but is not applied to Gen-3 Ethernet adapter cards or platforms. If the active port is on a Gen-3 card or platform, then only agg-rate is applicable. If the active port is on a Gen-2 card, then agg-rate and cir-rate apply when the scheduler mode is set to 16-priority. If the active port is on a Gen-1 card, then *agg-rate* and *cir-rate* are not applicable.

The **no** form of the command sets the *agg-rate* to the maximum and the *cir-rate* to 0 kb/s.

Default no agg-rate-limit

Parameters *agg-rate* — sets the PIR for the aggregate of all the queues on the SAP. The **max** keyword applies the maximum physical port rate possible.

Values 1 to 10000000 kb/s, or **max**

Default max

cir-rate — sets the CIR for the aggregate of all the queues on the SAP

Values 0 to 10000000 kb/s, or **max**

Default 0 kb/s

filter

Syntax **filter ip** *ip-filter-id*
filter ipv6 *ipv6-filter-id*
no filter [**ip** *ip-filter-id* | **ipv6** *ipv6-filter-id*]

Context config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress

Description This command associates an IPv4 or IPv6 filter policy with an egress or ingress IES SAP.

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 *ip-filter-id* or *ipv6-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 will be displayed.

The **no** form of the command removes any configured filter ID association with the SAP. The filter policy cannot be deleted until it is removed from all SAPs where it is applied.

Default no filter

Parameters *ip-filter-id* — specifies the IPv4 filter policy. The filter ID or filter name must already exist within the created IP filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

ipv6-filter-id — specifies the IPv6 filter policy. The filter ID or filter name must already exist within the created IP filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

Note: For information on configuring IP filter IDs, refer to the 7705 SAR OS Router Configuration Guide, “Filter Policies”.

match-qinq-dot1p

Syntax **match-qinq-dot1p** {**top** | **bottom**}
no match-qinq-dot1p

Context config>service>ies>if>sap>ingress

Description This command specifies which dot1q tag position (top or bottom) in a qinq-encapsulated packet should be used when QoS evaluates dot1p classification.

The **no** form of the command restores the default dot1p evaluation behavior for the SAP, which means that the inner (bottom) tag (second tag) dot1p bits are used for classification.

By default, the dot1p bits from the inner tag service-delineating dot1q tag are used.

Table 87 shows which set of dot1p bits are used for QoS purposes when **match-qinq-dot1p** is configured. To use the table, find the row that represents the settings for **Port/SAP Type** and **Match-QinQ-Dot1q Setting**. Use the **Existing Packet Tags** column to identify which dot1q tags are available in the packet. Then use the **P-bits Used for Match** column to identify which dot1q tag contains the dot1p bits that are used for QoS dot1p classification.

Default no match-qinq-dot1p

Parameters **top —** the **top** parameter and **bottom** parameter are mutually exclusive. When the **top** parameter is specified, the outer tag's dot1p bits (topmost P-bits) are used (if existing) to match any **dot1p** *dot1p-value* entries.

bottom — the **bottom** parameter and **top** parameter are mutually exclusive. When the **bottom** parameter is specified, the bottommost P-bits (second tag's P-bits) are used (if existing) to match any **dot1p** *dot1p-value* entries.

Table 87: Match-QinQ-Dot1p Matching Behavior

Port/ SAP Type	Match-QinQ-Dot1p Setting [1]	Existing Packet Tags	P-bits Used for Match
Null	n/a	None	None
Null	n/a	Dot1p (VLAN ID 0)	None [2]
Null	n/a	Dot1q	None [2]
Null	n/a	TopQ BottomQ	None [2]
Dot1Q	n/a	None	None
Dot1Q	n/a	Dot1p (default SAP VLAN ID 0)	Dot1p P-bits
Dot1Q	n/a	Dot1q	Dot1q P-bits
QinQ/ X.Y	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.Y	Default or Bottom	TopQ BottomQ	BottomQ P-bits

Table 87: Match-QinQ-Dot1p Matching Behavior (Continued)

Port/ SAP Type	Match-QinQ-Dot1p Setting [1]	Existing Packet Tags	P-bits Used for Match
QinQ/ X.0	Top	TopQ	TopQ P-bits
QinQ/ X.0	Default or Bottom	TopQ	TopQ P-bits
QinQ/ X.0	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.0	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ X.*	Top	TopQ	TopQ P-bits
QinQ/ X.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ X.*	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ 0.*	Top	None	None
QinQ/ 0.*	Default or Bottom	None	None
QinQ/ 0.*	Top	TopQ	TopQ P-bits
QinQ/ 0.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ 0.*	Top	TopQ BottomQ	TopQ P-bits
QinQ/ 0.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ *.*	Top	None	None
QinQ/ *.*	Default or Bottom	None	None
QinQ/ *.*	Top	TopQ	TopQ P-bits
QinQ/ *.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ *.*	Top	TopQ BottomQ	TopQ P-bits
QinQ/ *.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits

Notes:

1. "Default" in this column refers to the **no** form of **match-qinq-dot1p** command.
2. For null encapsulation, the 7705 SAR does not process dot1p bits.

qinq-mark-top-only

Syntax **[no] qinq-mark-top-only**

Context config>service>ies>if>sap>egress

Description When enabled, the **qinq-mark-top-only** command specifies which P-bits to mark during packet egress. When disabled, both sets of P-bits are marked. When enabled, only the P-bits in the top Q-tag are marked. The **no** form of the command is the default state (disabled).

Table 88 shows the dot1p re-marking behavior for different egress port type/SAP type combinations and **qinq-mark-top-only** state, where "False" represents the default (disabled) state.

If a new tag is pushed, the dot1p bits of the new tag will be zero (unless the new tag is re-marked by the egress policy. The dot1p bits are configured using the **dot1p** parameter under the **configure**>**qos** context.

Table 88: Dot1P Re-marking Behavior for the QinQ-mark-top-only Command

Egress Port Type/SAP Type	QinQ-mark-top-only State	Egress P-Bits Marked or Re-marked
Null [1]	n/a	None
Dot1q/ X [1]	n/a	Outer tag
Dot1q/ * [2]	n/a	None
QinQ/ X.Y [1]	False	Two outer tags [3]
	True	Outer tag [3]
QinQ/ X.* [4]	True or False	Outer tag
QinQ/ X.0 [1]	True or False	Outer tag
QinQ/ 0.* [1]	True or False	None
QinQ/ *.* [2]	True or False	None

Notes:

1. This port type/SAP type is supported by the following services: Epipe, Ipipe, VPLS, IES, and VPRN.
2. This port type/SAP type is supported by the following services: Epipe and VPLS.
3. Normally, when a new tag is pushed, the dot1p bits of the new tag will be zero, unless the P-bits are remarked by the egress policy. However, an exception to this occurs when the egress SAP type is X.Y and only one new outer tag must be pushed. In this case, the new outer tag will have its dot1p bits set to the inner tag's dot1p bits.
4. This port type/SAP type is supported by the following services: Epipe, Ipipe, and VPLS.

Default no qinq-mark-top-only

qos

Syntax **qos** *policy-id*
no qos

Context config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress

Description This command associates a QoS policy with an ingress or egress IES SAP.

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the *policy-id* does not exist, an error will be returned.

The **qos** command associates both ingress and egress QoS policies. The **qos** command allows only ingress policies to be associated on the SAP ingress and only egress policies to be associated on the SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with an IES SAP at one time. Attempts to associate a second QoS policy of a given type returns an error.

By default, no specific QoS policy is associated with the SAP for ingress or egress; therefore, 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.

Parameters *policy-id* — associates the ingress or egress policy ID with the SAP. The policy ID or name must already exist.

Values 1 to 65535, or *policy-name* (up to 64 characters)

scheduler-mode

Syntax **scheduler-mode** {**4-priority** | **16-priority**}

Context config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress

Description This command sets the scheduler mode for the SAP and is part of the hierarchical QoS (H-QoS) feature on the 7705 SAR.

If the mode is 4-priority, then the SAP is considered an unshaped 4-priority SAP and the agg-rate-limit cannot be changed from its default values.

If the mode is 16-priority and the **agg-rate limit** parameters are configured to be non-default values, then the SAP is considered a shaped SAP. If the **agg-rate limit** parameters are left in their default settings, the SAP is considered an unshaped, 16-priority SAP.

This command is blocked on third-generation (Gen-3) Ethernet adapter cards and platforms, such as the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X with Ethernet ports, which use only 4-priority-hqos scheduling mode. It is also blocked on the 7705 SAR-X with a TDM MDA, which uses only 4-priority scheduling mode.

If the specified SAP is a LAG SAP, **scheduler-mode** can be configured but is not applied to Gen-3 adapter cards and platforms. If one of the ports in the LAG is on a Gen-1 adapter card, then **scheduler-mode** cannot be configured.

Default 4-priority

Parameters **4-priority —** sets the scheduler mode for the SAP to be 4-priority mode

16-priority — sets the scheduler mode for the SAP to be 16-priority mode

shaper-group

Syntax [**no**] **shaper-group** *shaper-group-name* [**create**]

Context config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress

Description This command applies a shaper group to a SAP. The shaper group must already be created and must be within the shaper policy assigned to the Ethernet MDA (for ingress) or port (for egress). A shaper group is a dual-rate aggregate shaper used to shape aggregate access ingress or egress SAPs at a shaper group rate. Multiple aggregate shaper groups ensure fair sharing of available bandwidth among different aggregate shapers.

The default shaper group cannot be deleted.

The **no** form of this command removes the configured **shaper-group**.

Default shaper-group “default”

Parameters *shaper-group-name* — the name of the shaper group. To access the default shaper group, enter “default”.

create — keyword used to create a shaper group

IES Service Spoke SDP Commands

spoke-sdp

Syntax **spoke-sdp** *sdp-id*:*vc-id* [**create**]
no spoke-sdp *sdp-id*:*vc-id*

Context config>service>ies>interface

Description This command binds a service to an existing Service Distribution Point (SDP).

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

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 is down.

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

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

Class-based forwarding is not supported on a spoke SDP used for termination on an IES or VPRN service. All packets are forwarded over the default LSP.

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 the service. Once the binding is removed, no packets are forwarded to the far-end router. The spoke SDP must be shut down first before it can be deleted from the configuration.

Default no sdp-id is bound to a service

Special Cases **IES —** only one *sdp-id* can be bound to an IES

Parameters *sdp-id* — the SDP identifier

Values 1 to 17407

vc-id — the virtual circuit identifier

Values 1 to 4294967295

egress

Syntax **egress**

Context config>service>ies>if>spoke-sdp

Description This command enables the context to configure egress SDP parameters.

vc-label

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

Context config>service>ies>if>spoke-sdp>egress

Description This command configures the static MPLS VC label used by the 7705 SAR to send packets to the far-end device in this service via this SDP.

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

Values 16 to 1048575

ingress

Syntax **ingress**

Context config>service>ies>if>spoke-sdp

Description This command enables the context to configure ingress SDP parameters.

filter

Syntax **filter ip** *ip-filter-id*
no filter

Context config>service>ies>if>spoke-sdp>ingress

Description This command associates an IP filter policy with an ingress spoke SDP. Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria.

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 will be returned.

In general, filters applied to ingress spoke SDPs apply to all packets on the spoke SDP. One exception is that 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 spoke SDP.

Parameters *ip-filter-id* — specifies the IP filter policy. The filter ID or filter name must already exist within the created IP filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

vc-label

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

Context config>service>ies>if>spoke-sdp>ingress

Description This command configures the static MPLS VC label used by the far-end device to send packets to the 7705 SAR in this service via this SDP.

Parameters *ingress-vc-label* — a VC ingress value that indicates a specific connection

Values 2048 to 18431

Routed VPLS Commands

vpls

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

Context config>service>ies>if

Description This command within the IP interface context binds the IP interface to the specified VPLS service name.

The system does not attempt to resolve the service name 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 service name. If the IP interface is already in the administratively up state, the system immediately attempts to resolve the given service name.

Parameters *service-name* — specifies the service name that the system attempts to resolve to an allow-ip-int-binding enabled VPLS service associated with the service name. The specified service name is an ASCII string of up to 32 characters.

ingress

Syntax **ingress**

Context config>service>ies>if>vpls

Description This command within the VPLS binding context defines the routed IPv4 optional filter override.

v4-routed-override-filter

Syntax **v4-routed-override-filter** *ipv4-filter-id*
no v4-routed-override-filter

Context config>service>ies>if>vpls>ingress

Description This command specifies an IPv4 filter ID applied to all ingress packets entering the VPLS service. The filter overrides the existing ingress IPv4 filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional, and if not defined or removed, the IPv4 routed packets use the existing ingress IPv4 filter on the VPLS virtual ports.

The **no** form of the command removes the IPv4 routed override filter from the ingress IP interface.

Default n/a

Parameters *ipv4-filter-id* — specifies the IPv4 filter policy. The filter ID or filter name must already exist within the created IP filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

v6-routed-override-filter

Syntax **v6-routed-override-filter** *ipv6-filter-id*
no v6-routed-override-filter

Context config>service>ies>if>vpls>ingress

Description This command specifies an IPv6 filter ID applied to all ingress packets entering the VPLS service. The filter overrides the existing ingress IPv6 filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional, and if it is not defined or it is removed, the IPv6 routed packets use the existing ingress IPv6 filter on the VPLS virtual ports.

The **no** form of the command removes the IPv6 routed override filter from the ingress IP interface.

Default n/a

Parameters *ipv6-filter-id* — specifies the IPv6 filter policy. The filter ID or filter name must already exist within the created IPv6 filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

IES Service Security Zone Configuration Commands

zone

Syntax **zone** {*zone-id* | *zone-name*} [**create**]
no zone {*zone-id* | *zone-name*}

Context config>service>ies

Description This command creates or specifies a security zone within an IES context. Each zone must have a unique ID.

All zones must be explicitly created with the **create** keyword.

Enter an existing zone without the **create** keyword to edit zone parameters.

The **no** form of this command deletes the zone. When a zone is deleted, all configuration parameters for the zone are also deleted.

Parameters *zone-id* — the zone ID number. The zone ID must be unique within the system.

Values 1 to 65534

abort

Syntax **abort**

Context config>service>ies>zone

Description This command discards changes made to a security feature.

Default n/a

begin

Syntax **begin**

Context config>service>ies>zone

Description This command enters the mode to create or edit security features.

Default n/a

commit

Syntax **commit**

Context config>service>ies>zone

Description This command saves changes made to security features.

Default n/a

inbound

Syntax **inbound**

Context config>service>ies>zone

Description This command enables the context to configure limit parameters on inbound firewall sessions.

Default n/a

outbound

Syntax **outbound**

Context config>service>ies>zone

Description This command enables the context to configure limit parameters for outbound firewall sessions on the CSM.

Default n/a

limit

Syntax **limit**

Context config>service>ies>zone>inbound
config>service>ies>zone>outbound

Description This command enables the context to configure limits on concurrent sessions for inbound or outbound firewall sessions on the CSM.

Default n/a

concurrent-sessions

Syntax **concurrent-sessions** {**tcp** | **udp** | **icmp** | **other**} *sessions*
no concurrent-sessions {**tcp** | **udp** | **icmp** | **other**}

Context config>service>ies>zone>inbound>limit
config>service>ies>zone>outbound>limit

Description This command configures the maximum number of concurrent firewall sessions that can be established per zone, in either the inbound or outbound direction.

Default n/a

Parameters **tcp —** specifies that TCP connection traffic is to be firewalled

udp — specifies that UDP connection traffic is to be firewalled

icmp — specifies that ICMP connection traffic is to be firewalled

other — specifies that the traffic to be firewalled is other than TCP, UDP, or ICMP

sessions — the maximum number of concurrent firewall sessions that can be created in a zone for the configured direction and protocol

Values 1 to 16383

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>service>ies>zone

Description This command creates a logical IP routing interface for a zone. Once created, attributes such as an IP address can be associated with the IP interface. Multiple interfaces can be configured on a zone.

The **no** form of this command removes the IP interface and all the associated configurations.

Parameters *ip-int-name* — the name of the interface to be configured within the zone

Values 1 to 32 characters (must start with a letter)

log

Syntax **log** {*log-id* | *name*}
no log

Context config>service>ies>zone

Description This command configures a log identifier for the specified zone. A log identifier can be configured in the `config>router>zone` context and the `config>security>policy` context.

The **no** form of this command removes logging for the zone.

Parameters *log-id* — the identifier for the log

Values 1 to 32 characters

name

Syntax **name** *zone-name*
no name

Context config>service>ies>zone

Description This command configures a zone name. The zone name is unique within the system. It can be used to refer to the zone under configure, show, and clear commands.

Parameters *zone-name* — the name of the zone

Values 1 to 32 characters (must start with a letter). If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

nat

Syntax **nat**

Context config>service>ies>zone

Description This command enters the context to configure NAT parameters for a zone.

pool

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

Context config>service>ies>zone>nat

Description This command configures the NAT pool for the security zone within an IES service. Each pool must have a unique ID.

All pools must be explicitly created with the **create** keyword.

Enter an existing pool without the **create** keyword to edit pool parameters.

The **no** form of this command deletes the specified NAT pool. When a pool is deleted, all configuration parameters for the pool will also be deleted.

Parameters *pool-id* — the pool ID number

Values 1 to 100

direction

Syntax **direction** {**zone-outbound** | **zone-inbound** | **both**}
no direction

Context config>service>ies>zone>nat>pool

Description This command configures the NAT pool direction for the security zone. A specific NAT pool can be configured for different directions while using the same policy. For example, if the security policy entry direction is set to **both**, separate inbound and outbound pools can be created for that policy.

Parameters **zone-outbound —** configures a pool for the policy outbound traffic

zone-inbound — configures a pool for the policy inbound traffic

both — configures a pool for policy inbound and outbound traffic

entry

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

Context config>service>ies>zone>nat>pool

Description This command configures a NAT pool entry within an IES service.

The **no** form of this command deletes the entry with the specified ID. When an entry is deleted, all configuration parameters for the entry will also be deleted.

Parameters *entry-id* — the entry ID number

Values 1 to 65535

ip-address

Syntax **ip-address** *ip-address* [**to** *ip-address*] **interface** *ip-int-name*
no ip-address

Context config>service>ies>zone>nat>pool>entry

Description This command configures the source IP address or IP address range to which packets that match NAT policy are routed using NAT. An interface can also be configured, in which case all packets that match NAT policy are routed to the interface IP address. If the interface IP address is changed dynamically, NAT is updated accordingly. Only one IP address can be associated with an IP interface. Source IP addresses and interfaces cannot be used together in a single NAT pool.

The IP address for the interface must be entered in dotted-decimal notation. The maximum IP address range limit is 255.

The **no** form of the command removes the IP address assignment. The **no** form of this command can only be performed when the IP interface is administratively shut down. Shutting down the IP interface brings the interface operationally down.

Parameters *ip-address* — the source IP address to be used by NAT. The *ip-address* portion of the **ip-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 1.0.0.0 to 223.255.255.255

ip-int-name — the name of the interface to be used by NAT

port

Syntax **port** *port* [**to** *port*]
no port

Context config>service>ies>zone>nat>pool>entry

Description This command configures the UDP/TCP port or port range. Packets that match NAT policy undergo network port address translation (NPAT) and are routed to their source UDP/TCP port. Configuring a UDP/TCP port pool requires an IP-address pool because the 7705 SAR does not support port address translation (PAT) alone.

The **no** form of this command deletes the port or port range.

Parameters *port* — the UDP/TCP port or range of ports to which NPAT is applied

name

Syntax **name** *pool-name*
no name

Context config>service>ies>zone>nat>pool

Description This command configures a zone pool name. Pool names must be unique within the group of pools defined for a zone. It can be used to refer to the pool under configure, show, and clear commands.

Parameters *pool-name* — the name of the pool. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Values 1 to 32 characters (must start with a letter).

policy

Syntax **policy** {*policy-id* | *policy-name*}
no policy

Context config>service>ies>zone

Description This command sets the policy to be used by the security zone to build its matching criteria for incoming packets.

The **no** form of this command deletes the specified policy.

Parameters *policy-id* — the number of the referenced policy

Values 1 to 65535

policy-name — the name of the referenced policy

Show Commands

Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

customer

Syntax **customer** [*customer-id*]

Context show>service

Description This command displays service customer information.

Parameters *customer-id* — specifies the customer ID number to be displayed

Values 1 to 2147483647

Output The following output is an example of service customer information, and Table 89 describes the fields.

Sample Output

```
A:ALU-2# show service customer 1
===============================================================================
Customer  1
===============================================================================
Customer-ID        : 1
Contact            : Tech Support
Description        : Default customer
Phone              : (613) 555-1122
===============================================================================
```

Table 89: Show Service Customer Output Fields

Label	Description
Customer-ID	ID that uniquely identifies the customer
Contact	Name or title of the primary contact person
Description	Generic information about the customer
Phone	Phone number by which to reach the contact person

egress-label

Syntax **egress-label** *start-label* [*end-label*]

Context show>service

Description This command displays service information 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 labels in the specified range are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *end-label* — the ending egress label value for which to display services using the label range

Values 2049 to 131071

Default the *start-label* value

start-label — the starting egress label value for which to display services using the label range. If only *start-label* is specified, only services using *start-label* are displayed.

Values 0, or 2048 to 131071

Output The following output is an example of service egress label information, and Table 90 describes the fields.

Sample Output

In the sample below, services 3, 5 and 6 are IES, and services 5000 and 5001 are VPLS services.

```
*A:ALU-12>show>service# egress-label 0 131071
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Binding        Type  I.Lbl                E.Lbl
-------------------------------------------------------------------------------
3          15:15              Spok  0                    0
5          5:5                Spok  0                    0
6          5:6                Spok  0                    0
5000       15:5000            Mesh  0                    0
5000       15:5001            Spok  0                    0
5001       5001:100           Spok  0                    0
-------------------------------------------------------------------------------
Number of Bindings Found : 6
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-12#
```

Table 90: Show Service Egress Output Fields

Label	Description
Svc Id	The ID that identifies a service
Sdp Binding	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 7705 SAR in this service by the SDP
E. Lbl	The VC label used by 7705 SAR 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 label range

id

Syntax **id** *service-id*

Context show>service

Description This command displays information for a particular service ID

Parameters *service-id* — identifies the service in the domain by service number or name

all

Syntax **all**

Context show>service>id

Description This command displays detailed information for all aspects of the service.

Output The following output is an example of service ID all information, and Table 91 describes the fields.

Sample Output (IES Management Service)

```
A:ALU-2# show service id 751 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 751
Service Type      : IES
Name              : IES751
Description       : ATM_Backhaul_SAM_Mgmt
Customer Id       : 10
```

```
Last Status Change: 09/09/2008 16:26:25
Last Mgmt Change  : 09/09/2008 16:25:04
Admin State       : Up                  Oper State        : Up
SAP Count         : 2
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP bundle-ima-1/3.1:0/75
-------------------------------------------------------------------------------
Service Id         : 751
SAP                : bundle-ima-1/3.1:0/75    Encap             : atm
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 09/09/2008 16:26:25
Last Mgmt Change   : 09/09/2008 16:25:04
Sub Type           : regular

Admin MTU          : 1572                     Oper MTU          : 1572
Ingr IP Fltr-Id    : 1                        Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
tod-suite          : None                     qinq-pbit-marking : both
Egr Agg Rate Limit : max

Acct. Pol          : None                     Collect Stats     : Disabled

Anti Spoofing      : None                     Nbr Static Hosts  : 0
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                        Egress qos-policy : 1
Shared Q plcy      : n/a                      Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        n/a
Off. HiPrio           : 802789                   n/a
Off. LowPrio          : n/a                      n/a

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        n/a
Dro. LowPrio          : n/a                      n/a
For. InProf           : 802789                   69039854
For. OutProf          : 0                        0

Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        n/a
Dro. OutProf          : n/a                      n/a
For. InProf           : 802829                   41753273
For. OutProf          : n/a                      n/a
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
```

```
                        Packets                 Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 802789                  n/a
Off. LoPrio           : n/a                     n/a
Dro. HiPrio           : 0                       n/a
Dro. LoPrio           : n/a                     n/a
For. InProf           : 802789                  69039854
For. OutProf          : 0                       0

Egress Queue 1
For. InProf           : 802829                  41753273
For. OutProf          : n/a                     n/a
Dro. InProf           : 0                       n/a
Dro. OutProf          : n/a                     n/a
-------------------------------------------------------------------------------
ATM SAP Configuration Information
-------------------------------------------------------------------------------
Ingress TD Profile : 32                       Egress TD Profile : 32
Alarm Cell Handling: Enabled                  AAL-5 Encap       : mux-ip
OAM Termination    : Enabled                  Periodic Loopback : Disabled

-------------------------------------------------------------------------------
Service Interfaces
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Interface
-------------------------------------------------------------------------------
If Name           : IP_10.75.11.0/24
Admin State       : Up                  Oper State        : Up
Protocols         : None
IP Addr/mask      : 10.75.11.2/24       Address Type      : Primary
IGP Inhibit       : Disabled            Broadcast Address : Host-ones
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
If Index          : 3                   Virt. If Index    : 3
Last Oper Chg     : 09/09/2008 16:26:25 Global If Index   : 32
SAP Id            : bundle-ima-1/3.1:0/75
TOS Marking       : Untrusted           If Type           : IES
SNTP B.Cast       : False               IES ID            : 751
MAC Address       : 00:00:00:00:00:10   Arp Timeout       : 14400
IP MTU            : 1524                ICMP Mask Reply   : True
Arp Populate      : Disabled            Host Conn Verify  : Disabled
LdpSyncTimer      : None
Proxy ARP Details
Rem Proxy ARP     : Disabled            Local Proxy ARP   : Disabled
Policies          : none

ICMP Details
Unreachables : Number - 100                     Time (seconds)   - 10
TTL Expired  : Number - 100                     Time (seconds)   - 10

IPCP Address Extension Details
Peer IP Addr      : Not configured
Peer Pri DNS Addr : Not configured
Peer Sec DNS Addr : Not configured
===============================================================================
```

Table 91: Show Service ID All Output Fields

Label	Description
Service Detailed Information	
Service Id	Service ID number
Service Type	Type of service (IES)
Name	The service name
Description	Generic information about the service
Customer Id	Customer ID number
Last Status Change	Date and time of the most recent status change to this service
Last Mgmt Change	Date and time of the most recent management-initiated change to this service
Admin State	Desired state of the service
Oper State	Operating state of the service
MTU	Service MTU
SAP Count	Number of SAPs specified for this service
Service Access Points	
Service Id	Service Identifier
SAP	ID of the access port where this SAP is defined
Encap	Encapsulation type for this SAP on the access port
Admin State	Desired state of the SAP
Oper State	Operating state of the SAP
Flags	Conditions that affect the operating status of this SAP. Display output includes ServiceAdminDown, PortOperDown, and so on.
Last Status Change	Date and time of the most recent status change to this SAP
Last Mgmt Change	Date and time of the most recent management-initiated change to this SAP
Admin MTU	Desired largest service frame size (in octets) that can be transmitted through this SAP to the far-end router, without requiring the packet to be fragmented

Table 91: Show Service ID All Output Fields (Continued)

Label	Description
Oper MTU	Actual largest service frame size (in octets) that can be transmitted through this SAP to the far-end router, without requiring the packet to be fragmented
Ingr IP Fltr-Id	Ingress IP filter policy ID assigned to the SAP
Egr IP Fltr-Id	Egress IP filter policy ID assigned to the SAP
Ingr Mac Fltr-Id	Ingress MAC filter policy ID assigned to the SAP (not applicable)
Egr Mac Fltr-Id	Egress MAC filter policy ID assigned to the SAP (not applicable)
Acct. Pol	Accounting policy applied to the SAP
Collect Stats	Specifies whether accounting statistics are collected on the SAP
QOS	
Ingress qos-policy	SAP ingress QoS policy ID
Egress qos-policy	SAP egress QoS policy ID
Sap Statistics	
Last Cleared Time	Date and time that a clear command was issued on statistics
Forwarding Engine Stats	
Dropped	Number of packets or octets dropped by the forwarding engine
Off. HiPrio	Number of high-priority packets or octets offered to the forwarding engine
Off. LowPrio	Number of low-priority packets offered to the forwarding engine
Queueing Stats (Ingress QoS Policy)	
Dro. HiPrio	Number of high-priority packets or octets discarded, as determined by the SAP ingress QoS policy
Dro. LowPrio	Number of low-priority packets discarded, as determined by the SAP ingress QoS policy
For. InProf	Number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP ingress QoS policy
For. OutProf	Number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP ingress QoS policy

Table 91: Show Service ID All Output Fields (Continued)

Label	Description
Queueing Stats (Egress QoS Policy)	
Dro. InProf	Number of in-profile packets or octets discarded, as determined by the SAP egress QoS policy
Dro. OutProf	Number of out-of-profile packets or octets discarded, as determined by the SAP egress QoS policy
For. InProf	Number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP egress QoS policy
For. OutProf	Number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP egress QoS policy
Sap per Queue stats	
Ingress Queue *n*	Index of the ingress QoS queue of this SAP, where *n* is the index number
Off. HiPrio	Number of packets or octets of high-priority traffic for the SAP (offered)
Off. LoPrio	Number of packets or octets count of low-priority traffic for the SAP (offered)
Dro. HiPrio	Number of high-priority traffic packets or octets dropped
Dro. LoPrio	Number of low-priority traffic packets or octets dropped
For. InProf	Number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	Number of out-of-profile packets or octets (rate above CIR) forwarded
Egress Queue *n*	Index of the egress QoS queue of the SAP, where *n* is the index number
For. InProf	Number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	Number of out-of-profile packets or octets (rate above CIR) forwarded
Dro. InProf	Number of in-profile packets or octets dropped for the SAP
Dro. OutProf	Number of out-of-profile packets or octets discarded
ATM SAP Configuration Information	
Ingress TD Profile	Profile ID of the traffic descriptor applied to the ingress SAP

Table 91: Show Service ID All Output Fields (Continued)

Label	Description
Egress TD Profile	Profile ID of the traffic descriptor applied to the egress SAP
Alarm Cell Handling	Indicates that OAM cells are being processed
AAL-5 Encap	AAL-5 encapsulation type—this is always mux-ip
OAM Termination	Indicates whether this SAP is an OAM termination point
Services Interfaces	
If Name	Name used to refer to the IES interface
Admin State	Administrative state of the interface
Oper State	Operational state of the interface
IP Addr/mask	IP address and subnet mask length of the interface
Address Type	Specifies whether the IP address for the interface is the primary or secondary address on the interface (this is always primary)
Broadcast Address	Broadcast address of the interface
If Index	Interface index corresponding to the IES interface
Virt. If Index	Virtual interface index of the IES interface
Last Oper Chg	Date and time of the last operating state change on the interface
Global IF Index	Global interface index of the IES interface
SAP Id	SAP identifier
TOS Marking	Specifies whether the ToS marking state is trusted or untrusted for the IP interface
If Type	Type of interface: IES
IES ID	Service identifier
MAC Address	IEEE 802.3 MAC address
Arp Timeout	Timeout for an ARP entry learned on the interface
IP MTU	IP maximum transmit unit for the interface
ICMP Mask Reply	Specifies whether the IP interface replies to a received ICMP mask request
ARP Populate	Indicates if ARP is enabled or disabled
Proxy ARP Details	
Rem Proxy ARP	Indicates whether remote proxy ARP is enabled or disabled

Table 91: Show Service ID All Output Fields (Continued)

Label	Description
Local Proxy ARP	Indicates whether local proxy ARP is enabled or disabled
Policies	Specifies the policy statement(s) applied to proxy ARP
ICMP Details	
Unreachables	Maximum number of ICMP destination unreachable messages that the IP interface will issue in a given period of time, in seconds Disabled—indicates that the IP interface will not generate ICMP destination unreachable messages
TTL Expired	Maximum number of ICMP TTL expired messages that the IP interface will issue in a given period of time, in seconds Disabled—indicates that the IP interface will not generate ICMP TTL expired messages

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 IES instance.

Parameters *ip-address* — the IP address for which ARP entries will be displayed

Default all IP addresses

ieee-address — the 48-bit MAC address for which ARP entries will be displayed. 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-id — the SAP ID for which ARP entries will be displayed. See Table 39 for a full list of SAP IDs.

ip-int-name — the interface name for which ARP entries will be displayed

Output The following output is an example of service ID ARP information, and Table 92 describes the fields.

Sample Output

```
*A:ALU-2# show service id 4 arp
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Type    Expiry    Interface         SAP
```

```
-------------------------------------------------------------------------------
3.2.3.3                         Other   00h00m00s to Internet       n/a
===============================================================================
*A:ALU-2#
```

Table 92: Show Service ID ARP Output Fields

Label	Description
ARP Table	
IP Address	Specified IP address
MAC Address	Specified MAC address
Type	Static—FDB entries created by management
	Learned—dynamic entries created by the learning process
	OAM—entries created by the OAM process
	Other—local entries created for the IP interfaces
Expiry	Age of the ARP entry
Interface	Interface applied to the service
SAP	SAP ID

base

Syntax **base**

Context show>service>id

Description This command displays basic information about the service specified by the ID.

Output The following output is an example of service ID base information, and Table 93 describes the fields.

Sample Output

```
*A:ALU-2# show service id 4 base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 4
Service Type      : IES
Name              : IES4
Description       : Default IES description for service ID 4
Customer Id       : 1
Last Status Change: 01/07/2010 21:58:44
Last Mgmt Change  : 01/07/2010 22:14:40
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
===============================================================================
```

Table 93: Show Service ID Base Output Fields

Label	Description
Service Basic Information	
Service Id	Service ID number
Service Type	Type of service
Name	The service name
Description	Generic information about the service
Customer Id	Customer ID number
Last Status Change	Date and time of the most recent status change to this service
Last Mgmt Change	Date and time of the most recent management-initiated change to this service
Admin State	Desired state of the service
Oper State	Operating state of the service
SAP Count	Number of SAPs specified for this service
Service Access & Destination Points	
Identifier	SAP ID
Type	Signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received
AdmMTU	Desired largest service frame size (in octets) that can be transmitted to the far-end router without requiring the packet to be fragmented
OprMTU	Actual largest service frame size (in octets) that can be transmitted to the far-end router without requiring the packet to be fragmented
Adm	Administrative state of the SAP
Opr	Operating state of the SAP

dhcp

Syntax **dhcp**

Context show>service>id

Description This command enables the context to display DHCP information for the IES service.

statistics

Syntax **statistics** [**interface** {*interface-name* | *ip-address*}]

Context show>service>id>dhcp

Description This command displays DHCP statistics information.

Parameters *interface-name* — the interface name for which DHCP statistics will be displayed

ip-address — the IP address of the interface for which to display information

Values a.b.c.d (host bits must be 0)

Output The following output is an example of service ID DHCP statistics information, and Table 94 describes the fields.

Sample Output

```
*A:ALU-2# show service id 4 dhcp statistics
===============================================================
DHCP Global Statistics, service 4
===============================================================
Rx Packets                          : 0
Tx Packets                          : 0
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 0
Client Packets Discarded            : 0
Client Packets Relayed              : 0
Server Packets Discarded            : 0
Server Packets Relayed              : 0
===============================================================
```

Table 94: Show Service ID DHCP Statistics Output Fields

Label	Description
DHCP Global Statistics, service x	
Rx Packets	Number of packets received
Tx Packets	Number of packets transmitted
Rx Malformed Packets	Number of malformed packets received

Table 94: Show Service ID DHCP Statistics Output Fields (Continued)

Label	Description
Rx Untrusted Packets	Number of untrusted packets received
Client Packets Discarded	Number of packets from the DHCP client that were discarded
Client Packets Relayed	Number of packets from the DHCP client that were forwarded
Server Packets Discarded	Number of packets from the DHCP server that were discarded
Server Packets Relayed	Number of packets from the DHCP server that were forwarded

summary

Syntax **summary** [**interface** *interface-name* | **saps**]

Context show>service>id>dhcp

Description This command displays a summary of DHCP configuration.

Parameters *interface-name* — the interface name for which DHCP summary information will be displayed

saps — displays SAPs per interface

Output The following output is an example of service ID DHCP summary information, and Table 95 describes the fields.

Sample Output

```
*A:ALU-2 show service id 4 dhcp summary
===============================================================================
DHCP Summary, service 4
===============================================================================
Interface Name                   Arp      Used/                Info    Admin
  SapId/Sdp                      Populate Provided             Option  State
-------------------------------------------------------------------------------
to Internet                      No       0/0                  Keep    Down
-------------------------------------------------------------------------------
Interfaces: 1
===============================================================================
*A:ALU-2
```

Table 95: Show Service ID DHCP Summary Output Fields

Label	Description
DHCP Summary, service x	
Interface Name SapID/ Sdp	Name of the interface
Arp Populate	Specifies whether ARP populate is enabled
Used/Provided:	Used—number of lease-states that are currently in use on the specified interface; that is, the number of clients on the interface that got an IP address by DHCP. This number is always less than or equal to the "Provided" field.
	Provided—lease-populate value configured for the specified interface
Info Option	Specifies whether Option 82 processing is enabled on the interface
Admin State	Administrative state

interface

Syntax **interface** [{[*ip-address* | *ip-int-name*] [*interface-type*] [**detail**] [*family*]} | **summary**]

Context show>service>id

Description This command displays information for the IP interfaces associated with the IES service.

Parameters *ip-address* — only displays the interface information associated with the specified IP address

Values *ipv4-address*: a.b.c.d (host bits must be 0)
ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D

ip-int-name — the IP interface name for which to display information

interface-type — displays either group or subscriber interfaces

detail — displays detailed IP interface information

family — displays the specified router IP interface family

Values ipv4—displays only those peers that have the IPv4 family enabled
ipv6—displays the peers that are IPv6-capable

summary — displays summary IP interface information

Output The following output is an example of service ID interface information, and Table 96 describes the fields.

Sample Output

```
*A:ALU-2 show service id 4 interface
===============================================================================
Interface Table
===============================================================================
Interface-Name                   Adm         Opr(v4/v6)  Type     Port/SapId
   IP-Address                                                     PfxState
-------------------------------------------------------------------------------
to Internet                      Up          Down/Down   IES      n/a
   3.2.3.3/24                                                     n/a
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
*A:ALU-2
```

Table 96: Show Service ID Interface Output Fields

Label	Description
Interface Table	
Interface-Name	Name of the interface
IP-Address	IP address of the interface
Adm	Administrative state of the interface
Opr (v4/v6)	Operational state of the interface
Type	Service type
Port/SapId PfxState	Port or SAP associated with the interface

sap

Syntax **sap** [*sap-id*] [**detail**]

Context show>service>id

Description This command displays information for the SAP associated with the IES service.

Parameters *sap-id* — the SAP ID for which SAP information will be displayed. See Table 39 for a full list of SAP IDs.

detail — displays detailed SAP information

Output The following output is an example of IES service SAP information. See Table 50 in VLL Services Command Reference for field descriptions.

Sample Output

```
*A:7705custDoc:Sar18>show>service# id 6000 sap 1/12/6 detail
```

```
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 6000
SAP                : 1/12/6                    Encap             : null
Description        : (Not Specified)
Admin State        : Up                        Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Multi Svc Site     : None
Last Status Change : 10/01/2012 19:47:49
Last Mgmt Change   : 10/02/2012 17:21:04
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                    QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)
Admin MTU          : 1514                      Oper MTU          : 1514
Ingr IP Fltr-Id    : n/a                       Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                       Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                       Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                      qinq-pbit-marking : n/a
Ing Scheduler Mode : 16-priority               Egr Scheduler Mode: 16-priority
Ing Agg Rate Limit : 1000                      Egr Agg Rate Limit: 2000
Ing Agg cir        : 100                       Egr Agg cir       : 200
Q Frame-Based Acct : Disabled
Acct. Pol          : None                      Collect Stats     : Disabled
Anti Spoofing      : None                      Avl Static Hosts  : 0
                                               Tot Static Hosts  : 0
Calling-Station-Id : n/a
Application Profile: None
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                         Egress qos-policy : 1
Shared Q plcy      : n/a                       Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A
                        Packets                  Octets
Forwarding Engine Stats
Dropped               : 0                        0
Off. HiPrio           : 0                        0
Off. LowPrio          : 0                        0
Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                        0
Dro. LowPrio          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0
Queueing Stats(Egress QoS Policy 1)
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
For. InProf           : 0                        0
For. OutProf          : 0                        0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                  Octets
Ingress Queue 1 (Priority)
Off. HiPrio           : 0                        0
```

```
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
===============================================================================
*A:7705custDoc:Sar18>show>service#
```

ingress-label

Syntax **ingress-label** *start-label* [*end-label*]

Context show>service

Description This command displays service information 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 labels in the specified range are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *end-label* — the ending ingress label value for which to display services using the label range

Values 2049 to 131071

Default the *start-label* value

start-label — the starting ingress label value for which to display services using the label range. If only *start-label* is specified, only services using *start-label* are displayed.

Values 0, or 2048 to 131071

Output The following output is an example of service ingress label information, and Table 97 describes the fields.

Sample Output

In the sample below, services 3, 5 and 6 are IES, and services 5000 and 5001 are VPLS services.

```
*A:ALU-12>show>service# ingress-label 0 131071
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Binding        Type  I.Lbl                E.Lbl
-------------------------------------------------------------------------------
3          15:15              Spok  0                    0
5          5:5                Spok  0                    0
```

```
6           5:6                 Spok  0                     0
5000        15:5000             Mesh  0                     0
5000        15:5001             Spok  0                     0
5001        5001:100            Spok  0                     0
-------------------------------------------------------------------------------
Number of Bindings Found : 6
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-12#
```

Table 97: Show Service ingress Output Fields

Label	Description
Svc Id	The ID that identifies a service
Sdp Binding	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 the 7705 SAR in this service by the SDP
E. Lbl	The VC label used by the 7705 SAR 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 label range

sap-using

Syntax **sap-using** [**sap** *sap-id*]
sap-using interface [*ip-address* | *ip-int-name*]
sap-using description
sap-using [**ingress** | **egress**] **atm-td-profile** *td-profile-id*
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**ingress** | **egress**] **qos-policy** [*qos-policy-id* | *qos-policy-name*]
sap-using [**ingress** | **egress**] **scheduler-mode** {**4-priority** | **16-priority**}

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 atm-td-profile command applies only to HSDPA offload (that is, IES management service).

Parameters *sap-id* — the SAP ID for which SAP information will be displayed. See Table 39 for a full list of SAP IDs.

ip-address — only displays the interface information associated with the specified IP address

Values *ipv4-address:* a.b.c.d (host bits must be 0)
ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D

ip-int-name — the IP interface name for which to display information

description — displays a SAP summary table with description information

ingress — specifies matching an ingress policy

egress — specifies matching an egress policy

td-profile-id **—** displays SAPs using this traffic description

filter-id — specifies the ingress filter policy for which to display matching SAP specifies. The filter ID or filter name must already exist within the created IP filters.

Values 1 to 65535 or filter-name (up to 64 characters)

qos-policy-id **—** the ingress or egress QoS policy ID for which to display matching SAPs

Values 1 to 65535

qos-policy-name **—** the ingress or egress QoS policy name for which to display matching SAPs

Values up to 64 characters

scheduler-mode — specifies the scheduler mode for which to display the SAPs

Output The following output is an example of service SAP-using information, and Table 98 describes the fields.

Sample Output

```
*A:ALU-48# show service sap-using
===============================================================================
Service Access Points
===============================================================================
PortId            SvcId     Ing.  Ing.    Egr.  Egr.   Adm  Opr
                            QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/2/7:1           103       1     none    1     none   Up   Up
1/2/7:2           104       1     none    1     none   Up   Up
1/2/7:3           105       1     none    1     none   Up   Up
1/1/1.1           303       1     none    1     none   Up   Up
1/1/1.2           304       1     none    1     none   Up   Up
1/1/1.3           305       1     none    1     none   Up   Up
1/1/9.1:10/50     701       1     none    1     none   Up   Down
1/1/9.1:20        702       1     none    1     none   Up   Down
1/1/9.1:10/51     703       1     none    1     none   Up   Down
1/1/9.1:30        704       1     none    1     none   Up   Down
1/1/9.1:10/52     705       1     none    1     none   Up   Down
1/1/9.1:40        706       1     none    1     none   Up   Down
1/1/9.1:11/50     805       1     none    1     none   Up   Down
```

```
1/1/9.1:21        806       1     none    1     none   Up   Down
1/1/9.1:12/52     807       1     none    1     none   Up   Down
1/1/9.1:41        808       1     none    1     none   Up   Down
1/1/1.9           903       1     none    1     none   Up   Up
1/1/1.10          904       1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 18
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48#

*A:ALU-48# show service sap-using sap 1/1/21:0
===============================================================================
Service Access Points Using Port 1/1/21:0
===============================================================================
PortId                  SvcId      Ing.  Ing.    Egr.  Egr.    Adm  Opr
                                   QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/1/21:0                1          1     none    1     none    Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48#

*A:ALU-48# show service sap-using description
===============================================================================
Service Access Points
===============================================================================
PortId                              SvcId        Adm  Opr  Description
-------------------------------------------------------------------------------
1/1/2                               1            Down Down (Not Specified)
1/2/1.1                             4            Up   Down (Not Specified)
1/10/4                              5            Up   Down (Not Specified)
-------------------------------------------------------------------------------
Number of SAPs : 3
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-48#

*A:ALU-48# show service sap-using egress atm-td-profile 1
===============================================================================
Service Access Point Using ATM Traffic Profile 1
===============================================================================
PortId             SvcId     Ing.  Ing.    Egr.  Egr.    Adm  Opr
                             QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/1/9.1:10/50      701       1     none    1     none    Up   Down
1/1/9.1:20         702       1     none    1     none    Up   Down
1/1/9.1:10/51      703       1     none    1     none    Up   Down
1/1/9.1:30         704       1     none    1     none    Up   Down
1/1/9.1:10/52      705       1     none    1     none    Up   Down
1/1/9.1:40         706       1     none    1     none    Up   Down
1/1/9.1:11/50      805       1     none    1     none    Up   Down
1/1/9.1:21         806       1     none    1     none    Up   Down
1/1/9.1:12/52      807       1     none    1     none    Up   Down
1/1/9.1:41         808       1     none    1     none    Up   Down
-------------------------------------------------------------------------------
```

```
Saps : 10
===============================================================================
*A:ALU-12#

*A:7705custDoc:Sar18>show>service# sap-using ingress scheduler-mode 4-priority
===============================================================================
Service Access Points Using Ingress 4-priority Scheduler Mode
===============================================================================
PortId                          SvcId      Scheduler Mode    Adm  Opr
-------------------------------------------------------------------------------
1/12/6                          6000       4-priority        Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:7705custDoc:Sar18>show>service#
```

Table 98: Show Service SAP-Using Output Fields

Label	Description
Service Access Point Using...	
PortID	ID of the access port where the SAP is defined
SvcID	Service identifier
Ing.QoS	SAP ingress QoS policy number specified on the ingress SAP
Ing. Fltr	IP filter policy applied to the ingress SAP
Egr.QoS	SAP egress QoS policy number specified on the egress SAP
Egr. Fltr	IP filter policy applied to the egress SAP
Scheduler Mode	The scheduler mode of the SAP: 4-priority or 16-priority
Adm	Desired state of the SAP
Opr	Actual state of the SAP
Description	The description of the SAP
Number of SAPs/Saps	Number of SAPs using this service

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-id — displays only those services associated with the specified customer ID

Values 1 to 2147483647

Output The following output is an example of service-using information, and Table 99 describes the fields.

Sample Output

```
*A:ALU-2# show service service-using ies
===============================================================================
Services [ies]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
4            IES       Down   Down       1                 01/07/2010 22:14:40
23           IES       Down   Down       1                 01/07/2010 21:58:44
-------------------------------------------------------------------------------
Matching Services : 2
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-2#
```

Table 99: Show Service Service-Using Output Fields

Label	Description
ServiceID	ID that defines the service
Type	Service type configured for the service ID
Adm	Administrative state of the service
Opr	Operational state of the service
CustomerId	ID of the customer owning the service
Last Mgmt Change	Date and time of the most recent management-initiated change to this service
Matching Services	Number of services of the same type

Clear Commands

id

Syntax **id** *service-id*

Context clear>service

Description This command clears commands for a specific service.

Parameters *service-id* — uniquely identifies a service by service number or name

dhcp

Syntax **dhcp**

Context clear>service>id

Description This command enables the context to clear DHCP parameters.

dhcp6

Syntax **dhcp6**

Context clear>service>id

Description This command enables the context to clear DHCPv6 parameters.

statistics

Syntax **statistics** [*ip-int-name* | *ip-address*]

Context clear>service>id>dhcp
clear>service>id>dhcp6

Description This command clears statistics for DHCP and DHCPv6 Relay.

If no interface name or IP address is specified, statistics are cleared for all configured interfaces.

If an interface name or IP address is specified, statistics are cleared only for that interface.

Parameters *ip-int-name* — 32 characters maximum

ip-address — IPv4 or IPv6 address

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

Debug Commands

id

Syntax **id** *service-id*

Context debug>service

Description This command debugs commands for a specific service. The **no** form of the command disables debugging.

Parameters *service-id* — the ID that uniquely identifies an IES service by service number or name

VPRN Services

In This Chapter

This chapter provides information about the Virtual Private Routed Network (VPRN) service and implementation notes.

Topics in this chapter include:

- VPRN Service Overview
- VPRN Features
- Configuring a VPRN Service with CLI
- VPRN Services Command Reference

VPRN Service Overview

Topics in this section include:

- Routing Prerequisites
- BGP Support
- IPSec Support
- Security Zones and VPRN
- Route Distinguishers
- In-Band Management using a VPRN

RFC 2547bis, an extension of RFC 2547, 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 particular VPN among the PE routers that are attached to that VPN. This is done in a way that ensures that routes from different VPNs remain distinct and separate, even if two VPNs have an overlapping address space. Within a particular VPN, the PE routers distribute route information from and to the CE routers. Since the CE routers do not peer with each other, there is no overlay visible to the VPN routing algorithm.

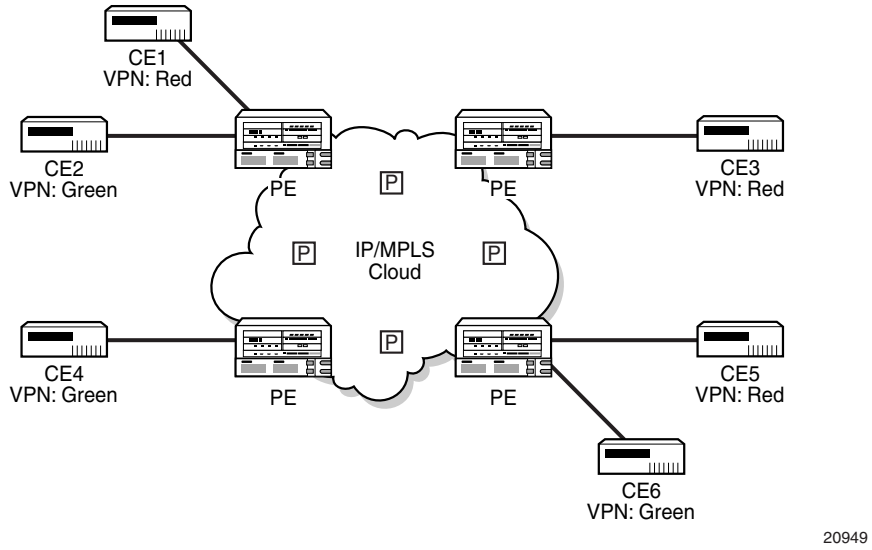
When BGP distributes a VPN route, it also distributes an MPLS label for that route. On an individual 7705 SAR, a single label is assigned to (advertised for) all routes in a VPN. A VRF lookup is used to determine the egress interface for a packet.

Before a customer data packet travels across the service provider's backbone network, it is encapsulated with the MPLS label that corresponds, in the customer's VPN, to the route that best matches the packet's destination address. That label (called the inner label) is the label that was advertised from the destination 7705 SAR, as described in the previous paragraph. The MPLS packet is further encapsulated with either another MPLS label or GRE tunnel 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 its VPRN association. Thus, the backbone core routers do not need to know the VPN routes.

Figure 93 shows an example of a VPRN network diagram, showing two VPNs (labeled "Red" and "Green") attached to PEs. The core routers are labeled "P".

Figure 93: Virtual Private Routed Network



Routing Prerequisites

RFC 2547bis requires the following features:

- multiprotocol extensions
- LDP support
- extended BGP community support
- BGP capability negotiation
- parameters defined in RFC 2918, *BGP Route Refresh*, and RFC 2796, *Route Reflector*
- a 4-byte autonomous system (AS) number

Tunneling protocol requirements are as follows:

- RFC 2547bis, *BGP/MPLS VPNs*, recommends implementing Label Distribution Protocol (LDP) to set up a full mesh of LSPs based on the IGP
- MPLS RSVP-TE tunnels can be used instead of LDP
- BGP route tunnels can be used as defined in RFC 3107
- alternatively, Generic Routing Encapsulation (GRE) tunnels can also be used

BGP Support

BGP is used with BGP extensions, as mentioned in Routing Prerequisites, to distribute VPRN routing information across the service provider's network.

BGP was initially designed to distribute IPv4 routing information. Therefore, multiprotocol extensions and the use of a VPN-IPv4 address were created to extend the ability of BGP 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. In addition, 128-bit VPN-IPv6 addresses extend the capability to distribute VPRN routing information.

BGP route tunnels can be used to distribute label mapping information for a particular route, as defined in RFC 3107. For more information on BGP route tunnels, refer to the 7705 SAR OS Routing Protocols Guide, "BGP Route Tunnel".



Note: The 7705 SAR supports 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. This allows up to 4 294 967 295 unique AS numbers.

BGP is configured through the `config>service>vprn>bgp` context.

BGP Fast Reroute with Prefix-Independent Convergence in a VPRN

BGP Fast Reroute (FRR) creates an alternate path to support fast rerouting of BGP traffic around failed or unreachable next hops. When BGP FRR is enabled, the system switches to a precalculated alternate path as soon as a failure is detected.

BGP Prefix-Independent Convergence (PIC) is supported on the 7705 SAR and is automatically enabled when a BGP backup path is enabled. With BGP FRR and PIC, alternate paths are precalculated and the FIB is updated with all alternate next hops. When a prefix has a backup path, and its primary paths fail, the affected traffic is rapidly diverted to the backup path without waiting for control plane reconvergence to occur. When many prefixes share the same primary paths, and in some cases also share the same backup path, the time to switch traffic to the backup path can be very fast and is independent of the number of prefixes.

In the VPRN context, BGP FRR is supported using unlabeled IPv4 and VPN-IPv4 routes. The supported VPRN scenarios are outlined in Table 100.

Table 100: BGP FRR Scenarios

Ingress Packet	Primary Route	Backup Route	PIC
IPv4 (ingress PE)	IPv4 route with next hop A resolved by an IPv4 route	IPv4 route with next hop B resolved by an IPv4 route	Yes
IPv4 (ingress PE)	VPN-IPv4 route with next hop A resolved by a GRE, LDP, RSVP or BGP tunnel	VPN-IPv4 route with next hop B resolved by a GRE, LDP, RSVP or BGP tunnel	Yes
MPLS (egress PE)	IPv4 route with next hop A resolved by an IPv4 route	IPv4 route with next hop B resolved by an IPv4 route	Yes
MPLS (egress PE)	IPv4 route with next hop A resolved by an IPv4 route	VPN-IPv4 route with next hop B resolved by a GRE, LDP, RSVP or BGP tunnel	Yes

If all IP prefixes require backup path protection, use a combination of the BGP context `backup-path` command and the VPRN context `enable-bgp-vpn-backup` command. If only specific IP prefixes require backup path protection, use route policies to apply the install backup path action to the best paths of the IP prefixes requiring protection.

For information about BGP FRR specific to the BGP context, refer to the 7705 SAR OS Routing Protocols Guide, “BGP FRR with Prefix-Independent Convergence”.

IPSec Support

The 7705 SAR supports IPSec and IPSec tunnels, where IES or VPRN is used as a public (untrusted) network-facing service and VPRN is used as a private (trusted) network-facing service. VPRN interfaces support provisioning of tunnel SAPs as part of IPSec provisioning. The *sap-id* for a VPRN tunnel SAP is `tunnel-1.private:`*tag*.

For more information, see the IPSec chapter in this guide.

Security Zones and VPRN

The 7705 SAR supports a number of mechanisms for node security, including Access Control Lists (ACLs), Network Address Translation (NAT), and stateful, zone-based firewalls. For information about ACLs, NAT, and firewalls, refer to the 7705 SAR OS Router Configuration Guide, “Configuring Security Parameters”.

NAT and firewall security configurations are both based on zones. Zones segment a network, making it easier to control and organize traffic. A zone consists of a group of Layer 3 interfaces with common criteria, bundled together. Security policies, which define a set of rules that determine how NAT or firewall should direct traffic, can be applied to the entire zone or to multiple zones. To enable NAT or firewall functionality, security policy, application group, host group, and profile parameters must be configured under the `config>security` context in the CLI, and a security zone must be configured under one or more of the following contexts:

- `config>router>zone`
- `config>service>vprn>zone`
- `config>service>ies>zone`

A zone is created by adding at least one Layer 3 interface to the zone configuration. Multiple zones can be created within each service context or within the router context. Layer 3 interfaces from different services cannot be grouped into a single common zone. Table 101 lists the supported interfaces that can be added to zones in each CLI context for NAT or firewall.

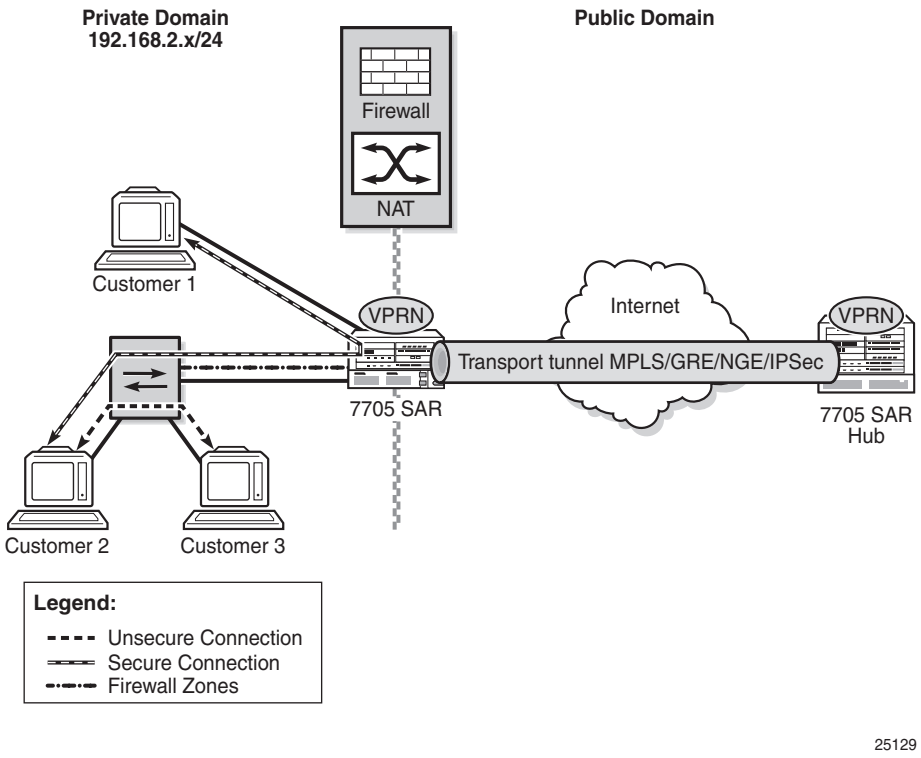
Table 101: Security Zone Interfaces per Context

CLI Context	Interface Type	NAT	Firewall
Router	Layer 3	✓	✓
VPRN	SAP	✓	✓
	Spoke-SDP termination	✓	✓
	IPSec private	✓	✓
	Routed VPLS	✓	✓
	MP-BGP	✓	✓
IES	SAP	✓	✓
	Spoke-SDP termination	✓	✓
	IPSec public	✓	
	Routed VPLS	✓	✓

A zone configured in the VPRN context could be used to create border security for Layer 3 service traffic traversing from the secure edge VPRN into the network core.

Figure 94 shows a firewall on a VPRN, where the core of the network is protected from access devices and from any traffic entering the VPRN through the transport tunnel. A VPRN configured with access security policies can also protect access networks or LANs from each other.

Figure 94: Firewall Protection for the Network Core

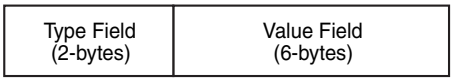


A security zone can also be created with spoke SDPs or auto-bind tunnels that have been configured for a VPRN service using MP-BGP. The auto-bind tunnels can be MPLS, LDP, or GRE. When a zone contains spoke SDPs or auto-bind MPLS tunnels, it cannot contain any other type of interface. A zone configured with spoke SDPs or auto-bind MPLS tunnels will firewall traffic arriving from the access side to this zone for any MP-BGP transport tunnel residing in that Layer 3 service. Once a security zone is created for MP-BGP transport tunnels, all MP-BGP transport tunnels going to far-end peers are part of this zone. Traffic entering or exiting the zone is firewalled; however, traffic traveling from one auto-bind tunnel to another within the same zone is not firewalled.

Route Distinguishers

The route distinguisher (RD) is an 8-byte value consisting of two major fields: the Type field and Value field. The Type field determines how the value field should be interpreted. The 7705 SAR OS implementation supports the three (3) Type-Value combinations, as defined in RFC 2547bis. Figure 95 illustrates the RD structure.

Figure 95: Route Distinguisher Structure



20950

The three Type-Value combinations supported are described in Table 102.

Table 102: Route Distinguisher Type-Value Fields

Type Field	Value Field	Notes
Type 0	Administrator subfield (2 bytes)	The Administrator field must contain an AS number (using private AS numbers is discouraged)
	Assigned number subfield (4 bytes)	The Assigned field contains a number assigned by the service provider
Type 1	Administrator subfield (4 bytes)	The Administrator field must contain an IP address (using private IP address space is discouraged)
	Assigned number subfield (2 bytes)	The Assigned field contains a number assigned by the service provider
Type 2	Administrator subfield (4 bytes)	The Administrator field must contain a 4-byte AS number (using private AS numbers is discouraged)
	Assigned number subfield (2 bytes)	The Assigned field contains a number assigned by the service provider

PE-to-CE Route Exchange

Routing information between the Provider Edge (PE) and Customer Edge (CE) can be exchanged by the following methods:

- EBGP
- OSPF
- RIP
- static routes

Each protocol provides controls to limit the number of routes learned from each CE router.

Route Redistribution

Routing information learned from the PE-to-CE routing protocols and configured static routes is injected into the associated local virtual routing and forwarding table (VRF). In the case of the 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 the PE-to-CE routing protocols is 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 MP-BGP instance is specified per VRF and is controlled by VRF route target associations or by VRF route policies.

A route belonging to a VPRN must use the protocol owner, VPN-IPv4, to denote that it is a 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 are removed from the service provider's routing tables dynamically, and wasted bandwidth is minimized.

Figure 96 and Figure 97 illustrate the use of CPE connectivity check in directly connected and multiple-hop connected routes.

Figure 96: Directly Connected IP Target

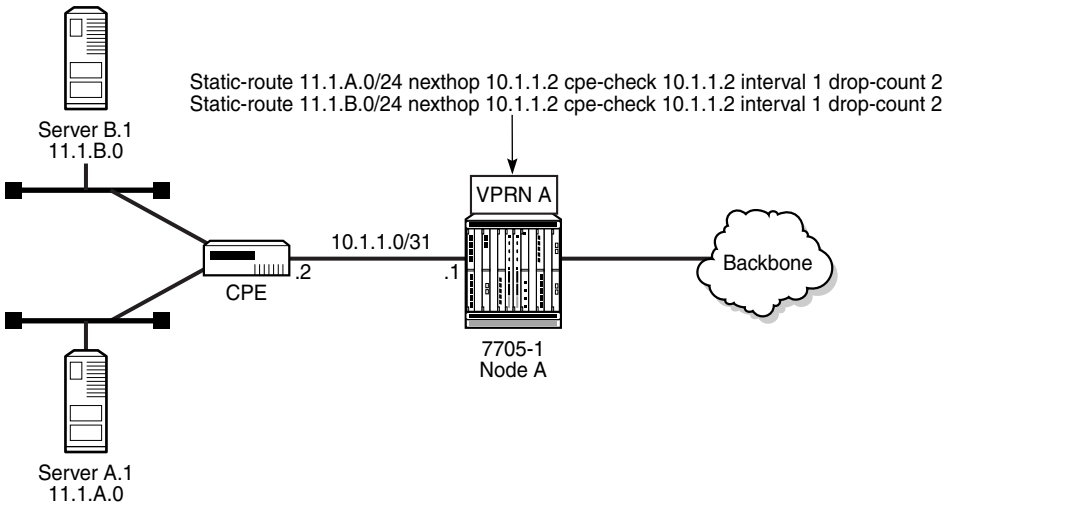
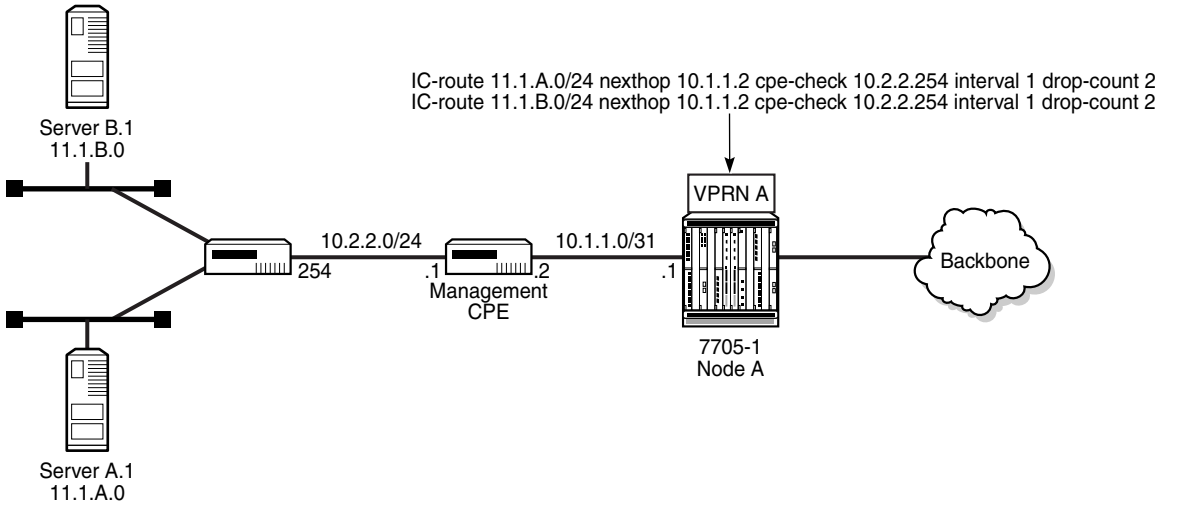


Figure 97: Multiple Hops to IP Target



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 7705 SAR router will continue to send polls and reactivate any routes that are restored.

In-Band Management using a VPRN

VPRN in-band management is supported on the 7705 SAR. In-band management of the 7705 SAR is performed using the global routing table (GRT) to perform a lookup on the system IP address of the 7705 SAR.

On network ingress, when a packet arrives from the transport tunnel to the VPRN, a lookup is performed within the VPRN on the inner customer packet IP header. If the destination IP address in the packet header matches that of the local 7705 SAR system IP address, and the command `enable-grt-local-management-only` is enabled under the VPRN instance, then the packet is extracted to the CSM for processing. If the `enable-grt-local-management-only` command is not enabled, the packet is routed using the corresponding 7705 SAR VRF FIB.

If the 7705 SAR system IP address is the same as any local IP address within the VPRN and the arriving packet destination IP matches this address, the packet is extracted to the CSM for processing only if the `enable-grt-local-management-only` command is enabled. Any ICMP packet destined for local interfaces will be processed by the system IP. If the local interface is operationally down, the system IP will still reply to ICMP packets successfully. Having a single IP address shared by the system IP and VPRN local interface is not recommended because some GRT-supported management protocols, such as Telnet and SSH, will not function with this configuration.

For MP-BGP VPRNs, the system IP address can be advertised to the far-end node using a static route whose next hop is configured as black-hole. If the command `enable-grt-local-management-only` is enabled under the VPRN instance, a packet arriving on a transport tunnel will be extracted to the CSM before hitting the black-hole route. In this case, the only effect of the black-hole route will be to advertise the system IP address to the far-end peer. If the command `enable-grt-local-management-only` is not enabled, packet forwarding will be the default forwarding mode; that is, all packets destined for the system IP address will be black-holed because of the static route configuration.

For MP-BGP VPRNs, when the command `enable-grt-local-management-only` is enabled, at least one interface (such as a loopback interface) must be configured on the VPRN and have an operational status of Up.

On network egress, the packets generated from the CSM with a source IP address that matches the 7705 SAR system IP address and destination IP address of either the far-end 5620 SAM or other management entity must perform a GRT route lookup in order to be resolved. A route policy can be configured with an IP address prefix of the far-end management entity and with the action to accept. This policy is configured for the GRT under the `config>router>policy-options` context and is installed in the GRT RIB (not the FIB) using the `export-grt-rib-only` command. The route installed in the GRT RIB will have a next hop of the corresponding VRF tunnel. This prevents any user data traffic in the GRT data path from leaking into the VPRN, and ensures that only the management traffic originating from the system IP address and the CSM gets transported through the VPRN. This forces the management packet to get routed by the corresponding VPRN transport tunnel, which means the VPRN route is leaked into the GRT so the GRT resolves the route using the corresponding VPRN.

Table 103 lists the management protocols supported by IPv4 and IPv6 GRT in the reverse and forward directions.

Table 103: IPv4 and IPv6 GRT-Supported Management Protocols

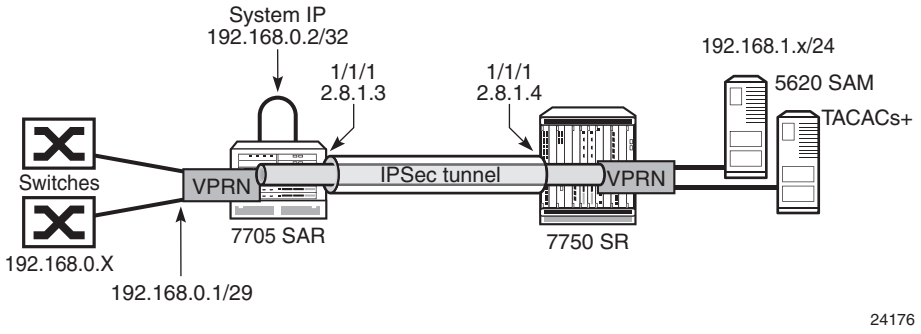
Protocol	Reverse Direction (Towards the 7705 SAR)	Forward Direction (From the 7705 SAR)
FTP Passive and Active	Yes [1]	Yes [2]
SFTP	Yes [1]	—
NTP	—	Yes [3]
RADIUS	Yes	—
SCP	Yes [1]	Yes [3]
SNMP	Yes	—
SNMP Trap	—	Yes
SSH	Yes [1]	Yes [3]
TACACS+	Yes	—
Telnet	Yes [1]	Yes [3]
TWAMP [4]	Yes [5]	Yes [6]
TWAMP Light	Yes [5]	Yes [6]

Notes:

1. Supported, if the 7705 SAR is acting as a server
2. Supported, if the 7705 SAR is acting as an active client
3. Supported, if the 7705 SAR is acting as a client
4. Supported on IPv4 only
5. Supported, if the 7705 SAR is acting as a server (control packets)
6. Supported, if the 7705 SAR is acting as a session reflector (test packets)

Figure 98 shows an example of IPv4 in-band management of the 7705 SAR and the switches behind it by the 5620 SAM and a TACACS+ server. In the example, an IPSec tunnel is being used as the VPRN in order to transport the management traffic via a secure and encrypted medium over the public internet.

Figure 98: IPv4 In-Band Management Using a VPRN Configured with GRT Lookup



In this example, the 7705 SAR system IP address is in the same subnet as the local interface; that is, subnet 192.168.0.x.

On network ingress in the above example, when `enable-grt-local-management-only` is configured for the VPRN, packets arriving on the 192.168.0 subnet are treated as follows.

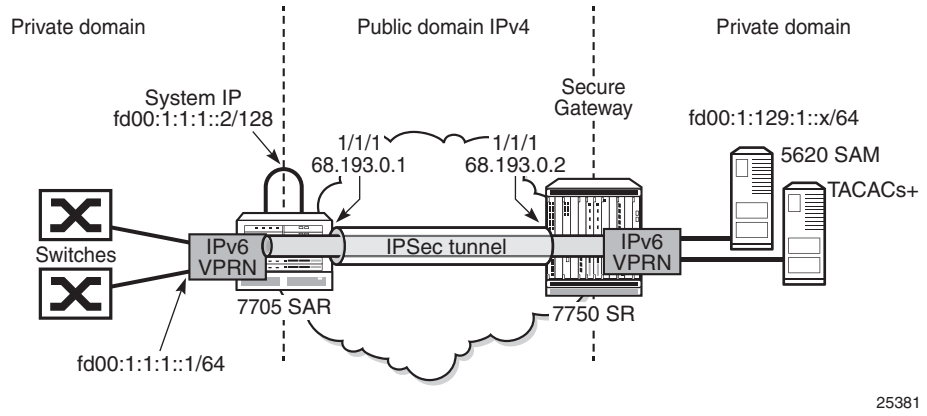
- If the packet destination IP address is 192.168.0.2, the packet is extracted to the CSM to be processed as management traffic.
- If the packet is destined for subnet 192.168.0.x/29, it is forwarded out of interface 192.168.0.1.

On network egress in the above example, routing is as follows.

- A static route can be installed in the 7705 SAR VPRN for subnet 192.168.1.0/24 with a next hop to IPSecTunnel_1.
- A route policy can be created with an IP address prefix of 192.168.1.0/24 and an action to accept. This route policy is configured under the `config>router>policy-options` context, and can be exported to the GRT RIB using the `export-grt-rib-only` command.
- The above configuration will add route 192.168.1.0/24 to the GRT RIB with the next hop being the corresponding VPRN IPSec tunnel. This entry will force the CSM-generated packets destined for the 192.168.1.x subnet to resolved by the VPRN IPSec tunnel.

Figure 99 shows an example of IPv6 in-band management of the 7705 SAR and the switches behind it by the 5620 SAM and a TACACS+ server.

Figure 99: IPv6 In-Band Management Using a VPRN Configured with GRT Lookup



On network ingress in the above example, the 7705 SAR system IP address is in the same subnet as the local interface IPv6 VPRN address. When `enable-grt-local-management-only` is configured for the IPv6 VPRN, packets arriving on the fd00:1:1:1 subnet are treated as follows.

- If the packet destination IP address is fd00:1:1:1::2, then the packet is extracted to the CSM to be processed as management traffic.
- If the packet is destined for subnet fd00:1:1:1::/64, it is forwarded out of interface fd00:1:1:1::1.

On network egress in the above example, routing is as follows.

- A static route can be installed at the 7705 SAR VPRN for subnet fd00:1:129:1::/64 with a next hop to IPSecTunnel_1.
- A route policy can be created with an IP address prefix of fd00:1:129:1::/64 and an action to accept. This route policy is configured under the `config>router>policy-options` context, and can be exported to the GRT RIB using the `export-grt-rib-only` command.

  The above configuration adds route fd00:1:129:1::/64 to the GRT RIB with the next hop being the corresponding VPRN IPSec tunnel. This entry forces the CSM-generated packets destined for the fd00:1:129:1::x subnet to be resolved by the VPRN IPSec tunnel.

VPRN Features

This section describes the 7705 SAR service features and any special capabilities or considerations as they relate to VPRN services:

- IP Interfaces
- SAPs
- PE-to-CE Routing Protocols
- PE-to-PE Tunneling Mechanisms
- Per-VRF Route Limiting
- RIP Metric Propagation in VPRNs
- Spoke SDPs
- Spoke SDP Termination to VPRN
- IPv6 on Virtual Private Edge Router
- IPv6 over IPv4 LAN-to-LAN IPSec tunnels

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:

- DHCP options (see DHCP and DHCPv6)
- Local DHCP server options (see Local DHCP and DHCPv6 Server)
- IPSec tunnel interfaces (see IPSec Support)
- IPCP options (see IPCP)
- ICMP options (see Troubleshooting and Fault Detection Services)

Configuration options found on core IP interfaces not supported on VPRN IP interfaces are:

- NTP broadcast receipt

Unnumbered Interfaces

Unnumbered interfaces are supported on VPRN and IES services for IPv4. Unnumbered interfaces are point-to-point interfaces that are not explicitly configured with a dedicated IP address and subnet; instead, they borrow (or link to) an IP address from another interface on the system (the system IP address, another loopback interface, or any other numbered interface) and use it as the source IP address for packets originating from the interface.

This feature is supported via both dynamic and static ARP for unnumbered interfaces to allow interworking with unnumbered interfaces that may not support dynamic ARP.

The use of unnumbered interfaces has no effect on IPv6 routes; however, the `unnumbered` command must only be used in cases where IPv4 is active (IPv4 only and mixed IPv4/IPv6 environments). When using an unnumbered interface for IPv4, the loopback address used for the unnumbered interface must have an IPv4 address. The interface type for the unnumbered interface is automatically point-to-point.

DHCP and DHCPv6

DHCP is a configuration protocol used to communicate network information and configuration parameters from a DHCP server to a DHCP-aware client. DHCP is based on the BOOTP protocol, with additional configuration options and the added capability of allocating dynamic network addresses. DHCP-capable devices are also capable of handling BOOTP messages.

A DHCP client is an IP-capable device (typically a computer or base station) that uses DHCP to obtain configuration parameters such as a network address. A DHCP server is an Internet host or router that returns configuration parameters to DHCP clients. A DHCP/BOOTP Relay agent is a host or router that passes DHCP messages between clients and servers.

DHCPv6 is not based on, and does not use, the BOOTP protocol.

The 7705 SAR can act as a DHCP client, a DHCP or DHCPv6 Relay agent, or a local DHCP or DHCPv6 server.

Home computers in a residential high-speed Internet application typically use the DHCP protocol to have their IP address assigned by their Internet service provider.

Since IP routers do not forward broadcast or multicast packets, this would suggest that the DHCP client and server must reside on the same network segment. However, for various reasons, it is sometimes impractical to have the server and client reside in the same IP network.

When the 7705 SAR is acting as a DHCP Relay agent, it processes these DHCP broadcast or multicast packets and relays them to a preconfigured DHCP server. Therefore, DHCP clients and servers do not need to reside on the same network segment.

When the 7705 SAR is acting as a local DHCP server, it processes these DHCP broadcast or multicast packets and allocates IP addresses for the DHCP client as needed.

The 7705 SAR supports a maximum of 16 servers per node on the 7705 SAR-8 with a CSMv1, and on the 7705 SAR-A, 7705 SAR-F, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-M, 7705 SAR-W, 7705 SAR-Wx, and 7705 SAR-X. The 7705 SAR supports a maximum of 62 servers per node on the 7705 SAR-8 with a CSMv2 and on the 7705 SAR-18. Any Layer 3 interface configured using the global routing table or Layer 3 services supports up to 8 servers.

DHCP Relay and DHCPv6 Relay

The 7705 SAR provides DHCP/BOOTP Relay agent services and DHCPv6 Relay agent services for DHCP clients. DHCP is used for IPv4 network addresses and DHCPv6 is used for IPv6 network addresses. Both DHCP and DHCPv6 are known as stateful protocols because they use dedicated servers to maintain parameter information.

Unless stated otherwise, DHCP is equivalent to “DHCP for IPv4”, or DHCPv4.

In the stateful autoconfiguration model, hosts obtain interface addresses and/or configuration information and parameters from a server. The server maintains a database that keeps track of which addresses have been assigned to which hosts.

The 7705 SAR supports DHCP Relay on access IP interfaces associated with IES and VPRN and on network interfaces. Each DHCP instance supports up to eight DHCP servers.

The 7705 SAR supports DHCPv6 Relay on access IP interfaces associated with IES and VPRN. Each DHCPv6 instance supports up to eight DHCPv6 servers.

Note:

- The 7705 SAR acts as a Relay agent for DHCP and DHCPv6 requests and responses, and can also be configured to function as a DHCP or DHCPv6 server. DHCPv6 functionality is only supported on network interfaces and on access IP interfaces associated with VPRN.
- When used as a CPE, the 7705 SAR can act as a DHCP client to learn the IP address of the network interface. Dynamic IP address allocation is supported on both network and system interfaces.
- For more information on DHCP and DHCPv6, refer to the 7705 SAR OS Router Configuration Guide, “DHCP and DHCPv6”.

DHCP Relay

The 7705 SAR provides DHCP/BOOTP Relay agent services for DHCP clients. DHCP is a configuration protocol used to communicate network information and configuration parameters from a DHCP server to a DHCP-aware client. DHCP is based on the BOOTP protocol, with additional configuration options and the added capability of allocating dynamic network addresses. DHCP-capable devices are also capable of handling BOOTP messages.

A DHCP client is an IP-capable device (typically a computer or base station) that uses DHCP to obtain configuration parameters such as a network address. A DHCP server is an Internet host or router that returns configuration parameters to DHCP clients. A DHCP/BOOTP Relay agent is a host or router that passes DHCP messages between clients and servers.

Home computers in a residential high-speed Internet application typically use the DHCP protocol to have their IP address assigned by their Internet service provider.

The DHCP protocol requires the client to transmit a request packet with a destination broadcast address of 255.255.255.255 that is processed by the DHCP server. Since IP routers do not forward broadcast packets, this would suggest that the DHCP client and server must reside on the same network segment. However, for various reasons, it is sometimes impractical to have the server and client reside in the same IP network. When the 7705 SAR is acting as a DHCP Relay agent, it processes these DHCP broadcast packets and relays them to a preconfigured DHCP server. Therefore, DHCP clients and servers do not need to reside on the same network segment.

DHCP OFFER messages are not dropped if they contain a yiaddr that does not match the local configured subnets on the DHCP relay interface. This applies only to regular IES and VPRN interfaces with `no lease-populate` configured on the DHCP relay interface.

DHCP Options

DHCP options are codes that the 7705 SAR inserts in packets being forwarded from a DHCP client to a DHCP server. Some options have additional information stored in suboptions.

The 7705 SAR supports the Relay Agent Information Option 82 as specified in RFC 3046. The following suboptions are supported:

- circuit ID
- remote ID
- vendor-specific options

DHCPv6 Relay

DHCPv6 Relay operation is similar to DHCP in that servers send configuration parameters such as IPv6 network addresses to IPv6 nodes, but DHCPv6 Relay is not based on the DHCP or BOOTP protocol. DHCPv6 can be used instead of stateless autoconfiguration (refer to the 7705 SAR OS Router Configuration Guide, "Neighbor Discovery") or in conjunction with it.

DHCPv6 is also oriented around IPv6 methods of addressing, especially the use of reserved, link-local scoped multicast addresses. DHCPv6 clients transmit messages to these reserved addresses, allowing messages to be sent without the client knowing the address of any DHCP server. This transmission allows efficient communication even before a client has been assigned an IP address. When a client has an address and knows the identity of a server, it can communicate with the server directly using unicast addressing.

The DHCPv6 protocol requires the client to transmit a request packet with a destination multicast address of ff02::1:2 (all DHCP servers and relay agents on the local network segment) that is processed by the DHCP server.

Similar to DHCP address allocation, if a client needs to obtain an IPv6 address and other configuration parameters, it sends a Solicit message to locate a DHCPv6 server, then requests an address assignment and other configuration information from the server. Any server that can meet the client's requirements responds with an Advertise message. The client chooses one of the servers and sends a Request message, and the server sends back a Reply message with the confirmed IPv6 address and configuration information.

If the client already has an IPv6 address, either assigned manually or obtained in some other way, it only needs to obtain configuration information. In this case, exchanges are done using a two-message process. The client sends an Information Request message, requesting only configuration information. A DHCPv6 server that has configuration information for the client sends back a Reply message with the information.

The 7705 SAR supports the DHCPv6 Relay Agent option in the same way that it supports the DHCP Relay Agent option. This means that when the 7705 SAR is acting as a DHCPv6 Relay Agent, it relays messages between clients and servers that are not connected to the same link.

DHCPv6 Options

DHCPv6 options are codes that the 7705 SAR inserts in packets being forwarded from a DHCPv6 client to a DHCPv6 server. DHCPv6 supports interface ID and remote ID options as defined in RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPV6)* and RFC 4649, *DHCPv6 Relay Agent Remote-ID Option*.

Local DHCP and DHCPv6 Server

The 7705 SAR supports local DHCP server functionality on the base router and on access IP interfaces associated with VPRN, by dynamically assigning IPv4 or IPv6 addresses to access devices that request them. This standards-based, full DHCP server implementation allows a service provider the option to decentralize IP address management into the network. The 7705 SAR can support public and private addressing in the same router, including overlapped private addressing in the form of VPRNs in the same router.

The 7705 SAR can act as a DHCP server or a DHCPv6 server.

An administrator creates pools of addresses that are available for assigned hosts. Locally attached hosts can obtain an address directly from the server. Routed hosts receive addresses through a relay point in the customer's network.

When a DHCP server receives a DHCP message from a DHCP Relay agent, the server looks for a subnet to use for assigning an IP address. If configured with the `use-pool-from-client` command, the server searches Option 82 information for a pool name. If a pool name is found, an available address from any subnet of the pool is offered to the client. If configured with the `use-gi-address` command, the server uses the gateway IP address (GIADDR) supplied by the Relay agent to find a matching subnet. If a subnet is found, an address from the subnet is offered to the client. If no pool or subnet is found, then no IP address is offered to the client.

When a DHCPv6 server receives a DHCP message from a DHCPv6 Relay agent, the server looks for a subnet to use for assigning an IP address. If configured with the `use-pool-from-client` command, the server searches Option 17 information for a pool name. If a pool name is found, an available address from any subnet of the pool is offered to the client. If configured with the `use-link-address` command, the server uses the address supplied by the Relay agent to find a matching subnet prefix. If a prefix is found, an address from the subnet is offered to the client. If no pool or prefix is found, then no IP address is offered to the client.

IPv4 and IPv6 address assignments are temporary and expire when the configured lease time is up. The server can reassign addresses after the lease expires.

If both the `no use-pool-from-client` command and the `no use-gi-address` command or `no use-link-address` command are specified, the server does not act.

DHCP and DHCPv6 Server Options

Options and identification strings can be configured on several levels.

DHCPv4 servers support the following options, as defined in RFC 2132:

- Option 1—Subnet Mask
- Option 3—Default Routers
- Option 6—DNS Name Servers
- Option 12—Host Name
- Option 15—Domain Name
- Option 44—Netbios Name Server
- Option 46—Netbios Node Type Option
- Option 50—IP Address
- Option 51—IP Address Lease Time
- Option 53—DHCP Message Type
- Option 54—DHCP Server IP Address
- Option 55—Parameter Request List
- Option 58—Renew (T1) Timer
- Option 59—Renew (T2) Timer

DHCPv4 servers also support Suboption 13 Relay Agent Information Option 82 as specified in RFC 3046, to enable the use of a pool indicated by the DHCP client.

DHCPv6 servers support the following options, as defined in RFC 3315:

- Option 1—OPTION_CLIENTID
- Option 2—OPTION_SERVERID
- Option 3—OPTION_IA_NA
- Option 4—OPTION_IA_TA
- Option 5—OPTION_IAADDR
- Option 6—OPTION_ORO
- Option 7—OPTION_PREFERENCE
- Option 8—OPTION_ELAPSED_TIME
- Option 9—OPTION_RELAY_MSG
- Option 11—OPTION_AUTH
- Option 12—OPTION_UNICAST
- Option 13—OPTION_STATUS_CODE
- Option 14—OPTION_RAPID_COMMIT

- Option 15—OPTION_USER_CLASS
- Option 16—OPTION_VENDOR_CLASS
- Option 17—OPTION_VENDOR_OPTS
- Option 18—OPTION_INTERFACE_ID
- Option 19—OPTION_RECONF_MSG
- Option 20—OPTION_RECONF_ACCEPT

These options are copied into the DHCP reply message, but if the same option is defined several times, the following order of priority is used:

1. subnet options
2. pool options
3. options from the DHCP client request

A local DHCP server must be bound to a specified interface by referencing the server from that interface. The DHCP server will then be addressable by the IP address of that interface. A normal interface or a loopback interface can be used.

A DHCP client is defined by the MAC address and the circuit identifier. This implies that for a certain combination of MAC and circuit identifier, only one IP address can be returned; if more than one request is made, the same address will be returned.

IPCP

Similar to DHCP over Ethernet interfaces, Internet Protocol Control Protocol (IPCP) extensions to push IP information over PPP/MLPPP VPRN (and IES) SAPs are supported. Within this protocol, extensions can be configured to define the remote IP address and DNS IP address to be signaled via IPCP on the associated PPP interface. The IPCP-based IP and DNS assignment process is similar to DHCP behavior; IPCP-based IP/DNS assignment is a natural use of PPP/MLPPP IP layer protocol handshake procedures. PPP/MLPPP connected devices hooked up to VPRN (and IES) can benefit from this feature for the assignment of IP and DNS to the associated interface.

Troubleshooting and Fault Detection Services

Bidirectional forwarding detection (BFD) can be configured on the VPRN interface. BFD is a simple protocol for detecting failures in a network. BFD uses a “hello” mechanism that sends control messages periodically to the far end and expects to receive periodic control messages from the far end. On the 7705 SAR, BFD is implemented for static routes in asynchronous mode only, meaning that neither end responds to control messages; rather, the messages are sent periodically from each end.

To support redundancy with fast switchover, BFD must be enabled to trigger the handoff to the other route in case of failure.

Due to the lightweight nature of BFD, it can detect failures faster than other detection protocols, making it ideal for use in applications such as mobile transport.

If BFD packets are not received in the configured amount of time, the associated static route is declared “not active”, causing a reroute to an alternative path, if any.



Note: Link failures detected by BFD will disable the IP interface.

The 7705 SAR also supports Internet Control Message Protocol (ICMP). ICMP is a message control and error reporting protocol that also provides information relevant to IP packet processing.

IP ECMP Load Balancing

IP ECMP allows the configuration of load balancing across all IP interfaces at the system level or interface level on the network side. You can include Layer 4 port attributes and/or the TEID attribute in the hashing algorithm with the `l4-load-balancing` and `teid-load-balancing` commands in the `config>service>vprn>interface` context. Configuration at the interface level overrides the system-level settings for the specific interface.

For more information on IP ECMP, refer to the 7705 SAR OS Router Configuration Guide, “Static Routes, Dynamic Routes, and ECMP”.

Proxy ARP

Proxy ARP is supported on VPRN interfaces.

Proxy ARP is a technique by which a router on one network responds to ARP requests intended for another node that is physically located on another network. The router effectively pretends to be the destination node by sending an ARP response to the originating node that associates the router's MAC address with the destination node's IP address (acts as a proxy for the destination node). The router then takes responsibility for routing traffic to the real destination.

For more information on proxy ARP, refer to the 7705 SAR OS Router Configuration Guide, "Proxy ARP".

Configurable ARP Retry Timer

A timer is available to configure a shorter retry interval when an ARP request fails. An ARP request may fail for a number of reasons, such as network connectivity issues. By default, the 7705 SAR waits 5000 ms before retrying an ARP request. The configurable retry timer makes it possible to shorten the retry interval to between 100 and 30 000 ms.

Note: The ARP retry default value of 5000 ms is intended to protect CPU cycles on the 7705 SAR, especially when it has a large number of interfaces. Configuring the ARP retry timer to a value shorter than the default should be done only on mission-critical links, such as uplinks or aggregate spoke SDPs transporting mobile traffic; otherwise, the retry interval should be left at the default value.

The configurable ARP retry timer is supported on VPRN and IES service interfaces, as well on the router interface.

SAPs

Topics in this section include:

- Encapsulations
- QoS Policies
- CoS Marking for Self-generated Traffic
- QinQ (VPRN)
- Filter Policies on a VPRN SAP

VPRN service also supports SAPs for IPSec tunnels (see IPSec Support).

Encapsulations

The following SAP encapsulations are supported on the 7705 SAR VPRN service:

- Ethernet null
- Ethernet dot1q
- Ethernet qinq
- PPP
- MLPPP
- MC-MLPPP
- LAG

QoS Policies

For each instance of VPRN service, QoS policies can be applied to the ingress and egress VPRN interface SAPs.

VPRN service ingress QoS policies only create the unicast queues defined in the policy. VPRN service egress QoS policies function in the same way as they do for other services, where the class-based queues are created as defined in the policy. Both the Layer 2 and Layer 3 criteria can be used in the QoS policies for traffic classification in a VPRN.

For VPRN services, the fabric mode needs to be set to aggregate mode as opposed to per-destination mode. VPRN services are only supported with aggregate-mode fabric profiles. When the fabric mode is set to per-destination mode, creation of VPRN service is blocked through the CLI. The user must change the fabric mode to aggregate mode before being able to configure VPRN services. As well, when a VPRN service is configured, changing from aggregate mode is blocked. The fabric mode is configured under the `config>qos>fabric-profile` context. For more information, refer to the 7705 SAR OS Quality of Service Guide.

CoS Marking for Self-generated Traffic

For each instance of VPRN service, DSCP marking and dot1p marking for self-generated traffic QoS can be configured for the applications supported by the 7705 SAR.

For VPRN service, DSCP marking is configured in the `vprn>sgt-qos>application` context. For more information about DSCP marking and self-generated QoS traffic, see "CoS Marking for Self-generated Traffic" in the 7705 SAR OS Quality of Service Guide.

QinQ (VPRN)

VPRN supports QinQ functionality. For details, see QinQ Support.

Filter Policies on a VPRN SAP

IPv4 and IPv6 filter policies can be applied to ingress and egress VPRN SAPs.

Refer to the 7705 SAR OS Router Configuration Guide, "Filter Policies", for information on configuring IP filters.

PE-to-CE Routing Protocols

The 7705 SAR supports the following PE-to-CE routing protocols for VPRN service:

- EBGP
- OSPF
- RIP
- static routes

EBGP is supported within both the router context and VPRN service context. OSPF is supported within the router context as well as within the VPRN service context, with some minor differences in the command set depending on the context.

Using OSPF in IP VPNs

Using OSPF as a PE-to-CE 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 PE-CE, or relying on static routes for the distribution of routes into the service provider's IP-VPN.

The following features are supported:

- transportation of OSPF learned routes as OSPF externals

  This feature uses OSPF as the protocol between the PE and CE routers; however, instead of transporting the OSPF LSA information across the IP-VPN, the OSPF routes are "imported" into MP-BGP as AS externals. As a result, other OSPF-attached VPRN sites on remote PEs receive these via type 5 LSAs.
- 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.

TTL Security

TTL security provides protection for EBGP peering sessions against CPU utilization-based attacks such as denial of service (DoS) attacks. This feature is supported for directly connected peering sessions and for multihop EBGP peering sessions. The BGP session can be over spoke-SDP terminated VPRN interfaces, SAP interfaces, and loopback interfaces, as well as over router interfaces and IPSec interface tunnels.

TTL security is most important for EBGP PE-CE sessions because CE devices can be multiple hops away, which adds a higher level of risk. TTL security provides a mechanism to better ensure the validity of BGP sessions from the CE device.

For more information on TTL security, refer to the 7705 SAR OS Routing Protocols Guide, "TTL Security".

PE-to-PE Tunneling Mechanisms

The 7705 SAR supports multiple mechanisms to provide transport tunnels for the forwarding of traffic between PE routers within the RFC 2547bis network.

The 7705 SAR VPRN implementation supports the use of:

- RSVP-TE protocol to create tunnel LSPs between PE routers
- LDP protocol to create tunnel LSPs between PE routers
- GRE tunnels 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 there is the ability to provide certain VPN services with their own transport tunnels by explicitly binding SDPs, if desired. When the `auto-bind` command is used, all services traverse the same LSPs and do not allow alternate tunneling mechanisms (like GRE) or the ability to craft sets of LSPs with bandwidth reservations for specific customers, as is available with explicit SDPs for the service.

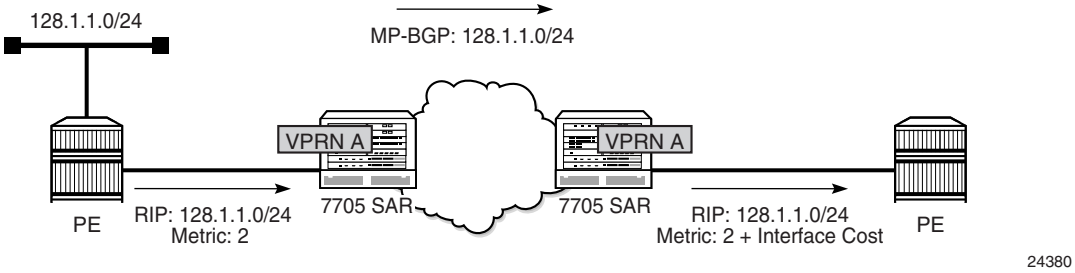
Per-VRF Route Limiting

The 7705 SAR 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 is nearly full and an option to disable additional route learning when the VRF is full or only generate an event.

RIP Metric Propagation in VPRNs

When RIP is used as the PE-CE protocol for VPRNs (IP-VPNs), the RIP metric is used only by the local node running RIP with the Customer Equipment (CE). The metric is not used with the MP-BGP path attributes that are exchanged between PE routers. Figure 100 shows an example of RIP metric propagation in a VPRN across two autonomous systems.

Figure 100: RIP Metric Propagation in VPRNs



The RIP metric can also be used to exchange routing information between PE routers if a customer network is dual homed to separate PEs. The RIP metric learned from the CE router can be used to choose the best route to the destination subnet. The RIP metric sets the BGP MED attribute, which allows remote PEs to choose the lowest MED and the PE with the lowest advertised RIP metric as the preferred egress point for the VPRN.

Spoke SDPs

For VPRN service, spoke SDPs can be used only for providing network connectivity between the PE routers.

Spoke SDP Termination to VPRN

This feature enables a customer to exchange traffic between a VLL or VPLS (Layer 2) service and an IES or VPRN (Layer 3) service. Customer premises traffic coming in from a VLL or VPLS service (SAP to spoke SDP) is forwarded over the IP/MPLS network to the IES or VPRN service, and vice versa. Network QoS policies can be applied to the spoke SDP to control traffic forwarding to the Layer 3 service.

In a Layer 3 spoke SDP termination to an IES or VPRN service, where the destination IP address resides within the IES or VPRN network, CE device-generated ARP frames must be processed by the Layer 3 interface. When an ARP frame is received over the spoke SDP at the Layer 3 interface endpoint, the 7705 SAR responds to the ARP frame with its own MAC address. When an ARP request is received from the routed network and the ARP entry for the CE device that is connected to the spoke SDP is not known, the 7705 SAR initiates an ARP frame to resolve the MAC address of the next hop or CE device.

Figure 101 shows traffic terminating on a specific IES or VPRN service that is identified by the SDP ID and VC label present in the service packet.

Figure 101: SDP ID and VC Label Service Identifiers (Conceptual View of the Service)

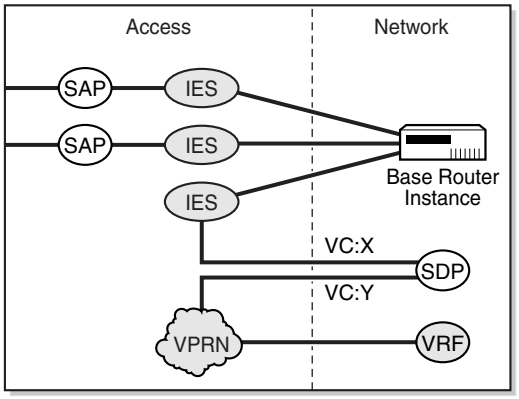
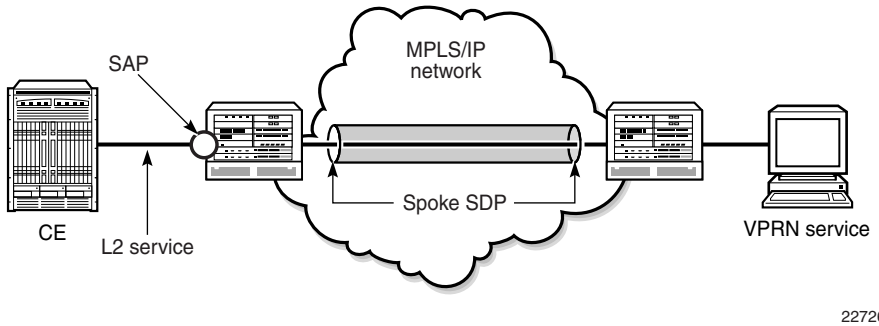


Figure 102 shows a spoke SDP terminating directly into a VPRN. In this case, a spoke SDP could be tied to an Epipe or a hierarchical VPLS service. There is no configuration required on the PE connected to the CE.

Figure 102: VPRN Spoke SDP Termination



Ethernet spoke SDP termination for VPRN service is supported over the following network uplinks:

- Ethernet network ports (null or dot1q encapsulation)
- PPP/MLPPP network ports. For information on PPP/MLPPP ports, refer to the 7705 SAR OS Interface Configuration Guide, “Access, Network, and Hybrid Ports”
- GPON module ports and DSL module ports when the module is installed in the 7705 SAR-M (variants with module slots)
- POS ports

Spoke SDP termination for VPRN supports the following:

- Ethernet PW to VRF
- interface shutdown based on PW standby signaling
- spoke SDP ingress IP filtering with filter logging
- label withdrawal for spoke SDPs terminated on VPRN
- statistics collection
- VCCV ping (type 2)

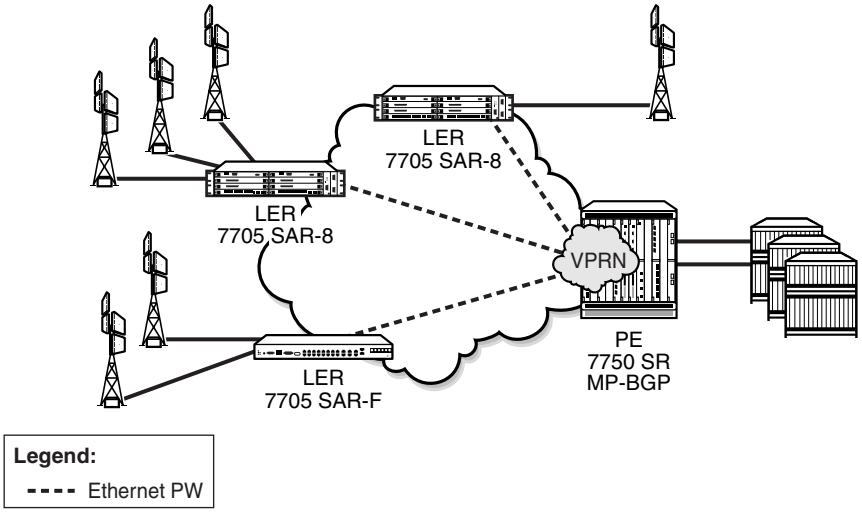
A spoke SDP on a VPRN interface service can be connected to the following entities:

- Epipe spoke SDP
- Epipe spoke SDP redundancy with standby-signal-master enabled
- IES interface
- VPRN interface
- VPLS spoke SDP
- VPLS spoke SDP redundancy with suppress-standby-signaling disabled

There are three scenarios to backhaul traffic from a given site that uses PWs and VPRN on a 7705 SAR.

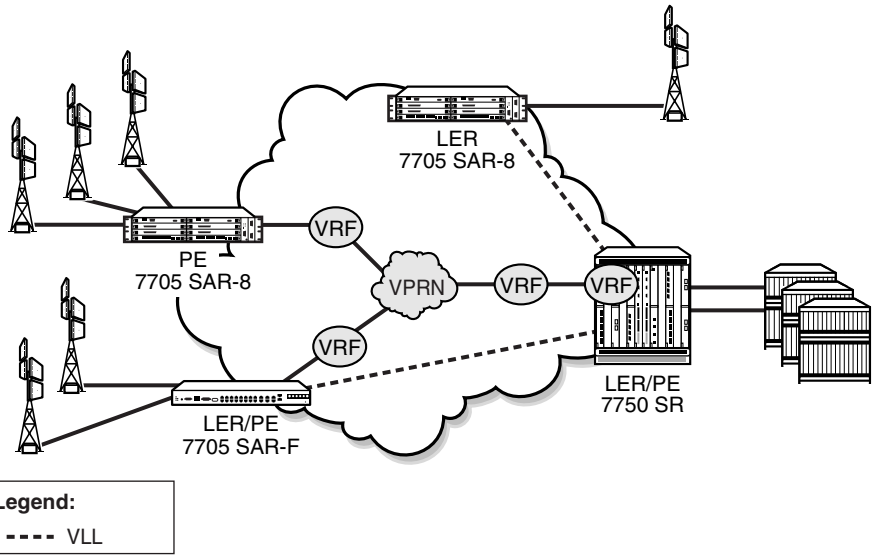
- Scenario 1 (Figure 103): An individual PW is configured on a per-CE device or a per-service basis. For routing services, this PW can be terminated to a VPRN at the 7750 SR end. This scenario offers per-service OAM and redundancy capabilities. Also, because there is no local communication on the remote 7705 SAR, traffic between any two devices connected to the 7705 SAR must traverse through the 7750 SR at the MTSO/CO.

Figure 103: Pseudowire-Based Backhaul (Spoke SDP Termination at 7750 SR)



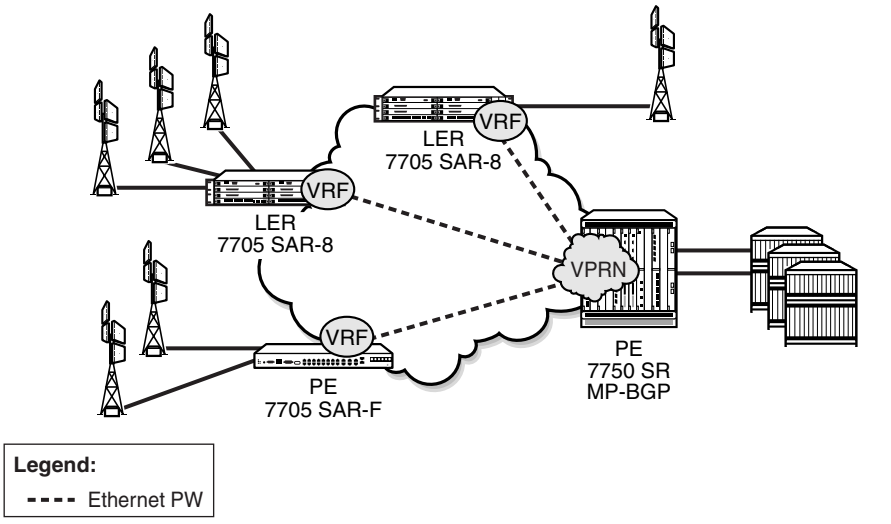
- Scenario 2 (Figure 104): An MP-BGP-based solution can provide a fully routed scenario.

Figure 104: VPRN in Mobile Backhaul Application



- Scenario 3 (Figure 105): In the hybrid scenario, IP forwarding among locally connected devices is handled by the 7750 SR directly, but instead of using MP-BGP to backhaul traffic, a PW is used to backhaul traffic to the MTSO/CO 7750 SR or possibly to a 7705 SAR node.

Figure 105: Spoke-SDP Termination to VPRN



IPv6 on Virtual Private Edge Router

The IPv6 on Virtual Private Edge Router (6VPE) feature allows customers that are migrating from an IPv4 to an IPv6 environment to use their existing IPv4 core infrastructure for transporting IPv6 traffic. Customers can migrate their access network to IPv6, including the eNodeBs, and keep the IPv4 core. The IPv4 core can be used for transporting eNodeB IPv6 traffic over MPLS or GRE tunnels. See IPv6 over IPv4 LAN-to-LAN IPSec tunnels for a description of how the 6VPE functionality is achieved.

Note: The 6VPE feature is not supported on the 16-port T1/E1 ASAP Adapter card or the 32-port T1/E1 ASAP Adapter card. This applies to both the access side (VPRN interfaces) and network side (MPLS/GRE tunnels).

On the network side, 6VPE is not supported on DS3/OC3 network interfaces, but is supported on SAR-A, SAR-M, SAR-H, and SAR-X T1/E1 ASAP network interfaces.

On the access side, 6VPE (VPRN SAP interfaces) is not supported on any T1/E1 ASAP adapter cards/blocks. VPRN spoke-SDP interfaces (spoke-SDP termination) are supported on SAR-A, SAR-M, SAR-H, and SAR-X T1/E1 ASAP blocks but not on T1/E1 adapter cards.

The classification of packets on a 6VPE access network is based on a customer packet Transaction Code (TC) field. The TC field is one byte long, but only the first six bits are used for the classification process. The use of six bits offers 64 different classes. The marking of the network outer tunnel DSCP/EXP bits is based on this access classification.

The supported protocols for 6VPE are listed in Table 103. The Access Control Lists for 6VPE are provided by Table 104, Table 105, and Table 106.

Table 104: 6VPE Access Control List, SAP

Service	IngV4	IngV6	IngMac	EgrV4	EgrV6	EgrMac
Network	Yes	Yes	No	Yes	Yes	No
Epipe	Yes	No	No	No	No	No
IES	Yes	Yes	No	Yes	Yes	No
Ipipe	Yes	No	No	No	No	No
VPLS	Yes	Yes	Yes	Yes	Yes	No
VPRN	Yes	Yes	No	Yes	Yes	No

Table 105: 6VPE Access Control List, SDP

Service	IngV4	IngV6	IngMac	EgrV4	EgrV6	EgrMac
Epipe	No	No	No	No	No	No
IES	Yes	No	No	No	No	No
Ipipe	No	No	No	No	No	No
VPLS	Yes	Yes	Yes	No	No	No
VPRN	Yes	Yes	No	No	No	No

Table 106: 6VPE Access Control List, r-VPLS Override

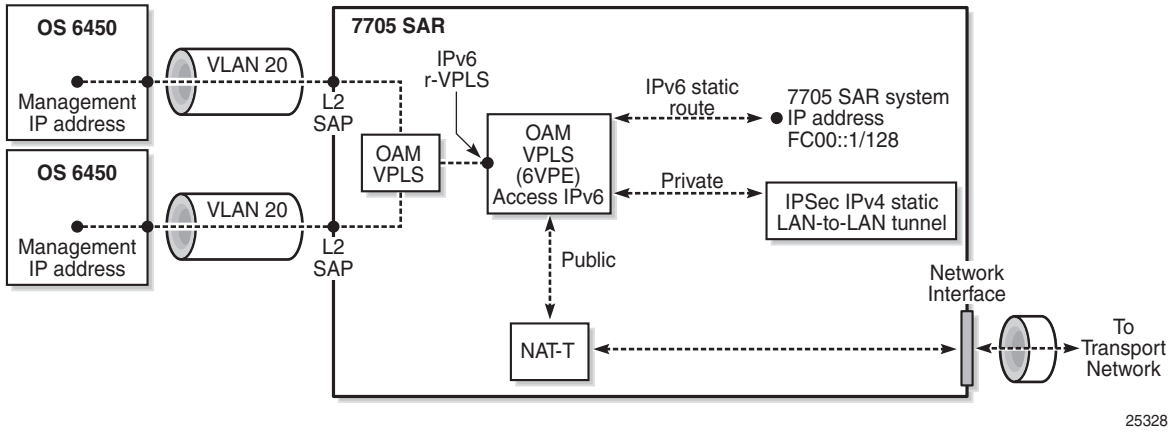
Service	Ingress Override-v4	IngV6
IES	Yes	Yes
VPRN	Yes	Yes

IPv6 over IPv4 LAN-to-LAN IPSec tunnels

In order to support the 6VPE functionality described in IPv6 on Virtual Private Edge Router, access (customer) IPv6 traffic is aggregated using a service VPRN and encrypted via an IPSec IPv4 static LAN-to-LAN tunnel, as shown in Figure 106. BGPv4 or BGPv6 can be configured over the IPSec IPv4 static LAN-to-LAN tunnel with an IPv6 address family to advertise the IPv6 VPRN routes to the peer VPRN.

The management IP addresses of all customer switches are migrated to IPv6. The system IP address of the 7705 SAR is configured as an IPv6 address and also migrated to IPv6. The OAM customer traffic is aggregated via an r-VPLS (IPv6 r-VPLS) into a 6VPE OAM VPRN. An IPv6 static route carries the IPv6 OAM traffic to an IPv4 static LAN-to-LAN tunnel, where the traffic is encrypted and encapsulated in an IPSec IPv4 transport tunnel. The IPSec IPv4 transport tunnel uses a NAT-to-single public IP address method, that is, NAT-T.

Figure 106: Access IPv6 Traffic Aggregation and Encryption



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
- Common Configuration Tasks
- Configuring VPRN Components
- Service Management Tasks

Basic Configuration

The following fields require specific input (there are no defaults) to configure a basic VPRN service:

- customer ID (refer to Configuring Customer Accounts)
- interface parameters
- spoke SDP parameters (at VPRN service level)

The following example displays a sample configuration of a VPRN service.

```
*A:ALU-1>config>service>vprn# info
----------------------------------------------
vrf-import "vrfImpPolCust1"
vrf-export "vrfExpPolCust1"
autonomous-system 10000
route-distinguisher 10001:1
auto-bind ldp
vrf-target target:10001:1
interface "to-ce1" create
    address 11.1.0.1/24
    exit
    sap 1/1/10:1 create
        ingress
            qos 100
            filter ip 10
        exit
        egress
            qos 1010
        exit
    exit
    dhcp
        description "DHCP test"
    exit
exit
exit
static-route 6.5.0.0/24 next-hop 10.1.1.2
no shutdown
----------------------------------------------
*A:ALU-1>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. Define BGP parameters and enable BGP in the `config>service>vprn>bgp` context. To configure BGP for the router, refer to the 7705 SAR OS Routing Protocols Guide, "BGP".
2. Associate a VPRN service with a customer ID.
3. Define a route distinguisher (RD) (mandatory).
4. Set the auto-bind parameter. Alternatively, define a VPRN spoke SDP. When using the spoke SDP command, you must enter the command for each peer PE router.
5. Define VRF route-target associations or VRF import/export policies.
6. Define OSPF parameters (optional).
7. Create a VPRN interface.
8. Define SAP parameters on the VPRN interface:
   - → select nodes and ports
   - → optional—select QoS policies other than the default (configured in `config>qos` context)
   - → optional—select filter policies (configured in `config>filter` context)
   - → optional—select accounting policy (configured in `config>log` context)
   - → optional—configure DHCP features
9. Define spoke SDP parameters on the VPRN interface.
10. Configure VRRP (optional).
11. Enable the service.

Configuring VPRN Components

Topics in this section include:

- Creating a VPRN Service
- Configuring Global VPRN Parameters
- Configuring Router Interfaces
- Configuring BGP for VPRN
- Configuring IPv6 Parameters for VPRN BGP
- Configuring VPRN IPv6 In-Band Management Parameters
- Configuring VPRN IPv6 Neighbor Discovery Parameters
- Configuring OSPF for VPRN
- Configuring RIP for VPRN
- Configuring a VPRN Interface
- Configuring a VPRN IPv6 Interface
- Configuring VPRN Interface Routed VPLS IPv6 Parameters
- Configuring VPRN Interface SAP Parameters
- Configuring VPRN Interface SAP IPv6 Parameters
- Configuring VPRN Interface Spoke SDP Parameters
- Configuring VPRN Interface Spoke SDP IPv6 Parameters
- Configuring VRRP
- Configuring a Security Zone within a VPRN
- Configuring VPRN Router Advertisement

Creating a VPRN Service

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

CLI Syntax:
```
config>service# vprn service-id [customer customer-id]
     route-distinguisher rd
     description description-string
     no shutdown
```

where *rd* is:

ip-addr:*comm-val* | *2byte-asnumber*:*ext-comm-val* | *4-byte-asnumber*:*comm-val*

ip-addr: a.b.c.d

comm-val: 0 to 65535

2-byte-asnumber: 1 to 65535

ext-comm-val: 0 to 4294967295

4-byte-asnumber: 1 to 4294967295

The following example displays a VPRN service configuration.

```
*A:ALU-1>config>service# info
----------------------------------------------
...
    vprn 1 customer 1 create
        route-distinguisher 100:76525
            no shutdown
        exit
...
----------------------------------------------
*A:ALU-1>config>service>vprn#
```

Configuring Global VPRN Parameters

See VPRN Services Command Reference for the CLI syntax to configure VPRN parameters.

The autonomous system (AS) number is configured in the `config>router>bgp` context.

A spoke SDP can be bound to the VPRN service using the `auto-bind` command or the `spoke-sdp sdp-id` command. However, when using the `spoke-sdp` command, you must create a spoke SDP for each peer PE router.

A VPRN spoke SDP can be any of the supported SDPs, except the IP SDP.

The following example displays a VPRN service with configured parameters.

```
*A:ALU-1>config>service# info
----------------------------------------------
...
    vprn 1 customer 1 create
        vrf-import "vrfImpPolCust1"
        vrf-export "vrfExpPolCust1"
        autonomous-system 10000
        router-id 2000
        route-distinguisher 10001:1
        spoke-sdp 2 create
        exit
        no shutdown
    exit
...
----------------------------------------------
*A:ALU-1>config>service#
```

Configuring Router Interfaces

Refer to the 7705 SAR OS Router Configuration Guide for command descriptions and syntax information to configure router interfaces.

The following example displays a router interface configuration:

```
ALU48>config>router# info
#------------------------------------------
echo "IP Configuration"
#------------------------------------------
...
    interface "if1"
        address 2.2.2.1/24
        port 1/1/33
    exit
    interface "if2"
        address 10.49.1.46/24
        port 1/1/34
    exit
    interface "if3"
        address 11.11.11.1/24
        port 1/1/35
    exit
...
#------------------------------------------
ALU48>config>router#
```

Configuring BGP for VPRN

Configuring BGP between the PE routers allows the PE routers to exchange information about routes originating and terminating in the VPRN. The PE routers use the information to determine which labels are used for traffic intended for remote sites.

The minimal parameters that should be configured for a VPRN BGP instance are:

- an autonomous system number for the router

  For an example of a VPRN service with a configured autonomous system number, refer to Configuring Global VPRN Parameters.
- a router ID

  For an example of a VPRN service with a configured router ID, refer to Configuring Global VPRN Parameters.
- a VPRN BGP peer group
- a VPRN BGP neighbor with which to peer
- a VPRN BGP peer-AS that is associated with the above peer

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. The VPRN BGP command hierarchy consists of three levels:

- global level
- group level
- neighbor level

Use the following CLI syntax to configure these three levels:

CLI Syntax:
```
config>service>vprn>bgp#
     group
          neighbor
```



Note: The `local-address` command must be explicitly configured if two systems have multiple BGP peer sessions between them.

BGP for MP-BGP purposes is configured under the `config>router>bgp` context. For more information about the BGP protocol, refer to the 7705 SAR OS Routing Protocols Guide, “BGP”.

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.

After a group name is created and options are configured, neighbors can be added in different autonomous systems, creating EBGP peers. All parameters configured for the peer group are inherited by each peer (neighbor), but a group parameter can be overridden on a specific neighbor-level basis.

VPRN BGP CLI Syntax

See VPRN Services Command Reference for the CLI syntax to configure VPRN BGP parameters.

The following example displays a VPRN BGP configuration:

```
*A:ALU-1>config>service# info
----------------------------------------------
...
    vprn 1 customer 1 create
        vrf-import "vrfImpPolCust1"
        vrf-export "vrfExpPolCust1"
        autonomous-system 10000
        route-distinguisher 10001:1
        auto-bind ldp
        vrf-target target:10001:1
        interface "to-ce1" create
            address 11.1.0.1/24
            sap 1/1/10:1 create
                ingress
                    qos 100
                    filter ip 6
                exit
                egress
                    qos 1010
                exit
            exit
        exit
        static-route 6.5.0.0/24 next-hop 10.1.1.2
        bgp
            router-id 10.0.0.1
            group "to-cel"
                export "vprnBgpExpPolCust1"
                multihop 3
                peer-as 65101
                ttl-security 10
                neighbor 10.1.1.2
                exit
            exit
```

```
        exit
        spoke-sdp 2 create
        exit
        no shutdown
    exit
...
-------------------------------------------
*A:ALU-1>config>service#
```

Configuring IPv6 Parameters for VPRN BGP

Use the following CLI syntax to configure IPv6 parameters for VPRN BGP:

CLI Syntax:

```
config>service# vprn service-id [customer customer-id]
     bgp
          family ipv6
               group name
                    family ipv6
                    neighbor ipv6-address
                         family ipv6
```

Example:

```
A:ALU>config>service# vprn 20
A:ALU>config>service>vprn$ bgp
A:ALU>config>service>vprn>bgp$ family ipv6
A:ALU>config>service>vprn>bgp>family$ group BGP1
A:ALU>config>service>vprn>bgp>family>group$ family ipv6
A:ALU>config>service>vprn>bgp>family>group>family$
 neighbor 2001:db8:a0b:12f0::1
A:ALU>config>service>vprn>bgp>family>group>family>
 neighbor$ family ipv6
A:ALU>config>service>vprn>bgp>family>group>family>
 neighbor$ exit
A:ALU>config>service>vprn>bgp>family>group>family$ exit
A:ALU>config>service>vprn>bgp>family>group$ exit
A:ALU>config>service>vprn>bgp>family$ exit
A:ALU>config>service>vprn>bgp$ exit
```

Configuring VPRN IPv6 In-Band Management Parameters

Use the following CLI syntax to configure IPv6 in-band management parameters for a VPRN service:

CLI Syntax:
```
config>service# vprn service-id [customer customer-id]
     grt-lookup
          enable-grt-local-management-only
          export-grt-rib-only policy-name [policy-name
             ...(up to 5 max)]
          export-v6-limit num-routes
```

Example:
```
config# service vprn 20
config>service>vprn# grt-lookup
config>service>vprn>grt-lookup# enable-grt-local-
 management-only
config>service>vprn>grt-lookup# export-grt-rib-only
 "policy-name 1"
config>service>vprn>grt-lookup# export-v6-limit 4
config>service>vprn>grt-lookup# exit
config>service>vprn# exit
```

The following example displays IPv6 in-band management parameters output.

```
A:ALU-A>config>service>vprn 20# info
#------------------------------------------
.....
            grt-lookup
                enable-grt-local-management-only
                export-grt-rib-only "policy-name 1"
                export-v6-limit 4
            exit
.....
```

Configuring VPRN IPv6 Neighbor Discovery Parameters

Use the following CLI syntax to configure IPv6 neighbor discovery parameters for a VPRN service:

CLI Syntax:
```
config# config>service# vprn service-id [customer
 customer-id]
     ipv6
          reachable-time seconds
          stale-time seconds
```

Example:
```
config# service vprn 20
config>service>vprn# ipv6
config>service>vprn>ipv6# reachable-time 30
config>service>vprn>ipv6# stale-time 14400
config>service>vprn>ipv6# exit
config>service>vprn# exit
```

The following example displays IPv6 neighbor discovery parameters output.

```
A:ALU-A>config>service>vprn 20# info
#------------------------------------------
        ...
            reachable-time 30
            stale-time 14400
        exit
        ...
```

Configuring OSPF for VPRN

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. For more information on the OSPF protocol and for the commands used to run OSPF over the backbone (router context), refer to the 7705 SAR OS Routing Protocols Guide.

Use the following CLI syntax to configure OSPF in the VPRN context:

CLI Syntax: `config>service>vprn>ospf#`

Refer to VPRN Services Command Reference for the CLI syntax to configure VPRN OSPF parameters.

The following example displays a VPRN OSPF configuration:

```
*A:ALU-1>config>service# info
----------------------------------------------
     vprn 2 customer 1 create
         interface "ospf_interface" create
         exit
         ospf
             area 0.0.0.0
                 interface "ospf_interface"
                     no shutdown
                 exit
             exit
         exit
----------------------------------------------
*A:ALU-1>config>service#
```

Configuring RIP for VPRN

PE routers need to advertise reachability information for each CE that is attached to a VRF. RIP can be used to exchange reachability information between PE and CE routers by establishing adjacency with a CPE router that supports RIP. Via this adjacency, RIP learns the subnet or subnets for the customer site and will advertise any routes learned from other CEs. The routing table is updated to reflect the new information.

RIP can be used to distribute routes between PE and CE routers. When PE and CE routers are RIP peers, the CE router can use RIP to transmit to the PE router the set of address prefixes that are reachable via the CE router. When RIP is configured on the CE, care must be taken to ensure that address prefixes from other sites, that is, address prefixes learned by the CE router from the PE router, are never advertised to the PE. Specifically, if a PE router receives a VPN-IPv4 route and distributes it to a CE, that route must never be distributed from the CE site to, either the originating PE router, or any other PE router.

The parameters configured at the VPRN RIP global level are inherited by the group and neighbor levels. Parameters can be modified and overridden on a level-specific basis. The VPRN RIP command hierarchy consists of three levels:

- global
- group
- neighbor

Hierarchical VPRN RIP commands can be modified on different levels. The most specific value is used. A group-specific command takes precedence over a global command. A neighbor-specific command takes precedence over a global or group-specific command.

Note: Careful planning is essential to implement commands that can affect the behavior of global, group, and neighbor levels. Because the RIP commands are hierarchical, analyze the values that can disable features on a particular level.

To enable a VPRN RIP instance, the RIP protocol must be enabled in the `config>service>vprn>rip` context of the VPRN. VPRN RIP is administratively enabled upon creation. Configuring other RIP commands and parameters is optional.

The minimum RIP configuration for a VPRN instance must define:

- one VPRN RIP peer group
- one VPRN RIP neighbor peer
- one VPRN RIP peer-AS associated with the neighbor peer

See VPRN Services Command Reference for the CLI syntax to configure VPRN RIP parameters.

The following example displays a VPRN RIP configuration:

```
*A:ALU-1>config>service# info
----------------------------------------------
...
        vprn 1 customer 1 create
            vrf-import "vrfImpPolCust1"
            vrf-export "vrfExpPolCust1"
            ecmp 8
            autonomous-system 10000
            route-distinguisher 10001:1
            auto-bind ldp
            vrf-target target:10001:1
            interface "to-ce1" create
                address 11.1.0.1/24
                sap 1/1/10:1 create
                    ingress
                        qos 100
                    exit
                    egress
                        qos 1010
                        filter ip 6
                    exit
                exit
            exit
            rip
                export "vprnRipExpPolCust1"
                group "ce1"
                    neighbor "to-ce1"
                    exit
                exit
            exit
            spoke-sdp 2 create
            exit
            no shutdown
        exit
...
----------------------------------------------
```

For more information about the RIP protocol, refer to the 7705 SAR OS Routing Protocols Guide.

Configuring a VPRN Interface

Interface names associate an IP address with the interface, and then associate the IP interface with a physical port. The logical interface can associate attributes such as an IP address, port, or Link Aggregation Group (LAG). There are no default interfaces.

Note:

- The VPRN interface can be configured 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.
- See Configuring a VPRN IPv6 Interface for the CLI required to configure VPRN IPv6 interface parameters.

See VPRN Services Command Reference for CLI commands and syntax.

The following example displays a VPRN interface configuration:

```
*A:ALU-1>config>service>vprn# info
----------------------------------------------
...
    vprn 1 customer 1 create
        vrf-import "vrfImpPolCust1"
        vrf-export "vrfExpPolCust1"
        autonomous-system 10000
        route-distinguisher 10001:1
        auto-bind ldp
        vrf-target target:10001:1
        interface "to-ce1" create
            address 11.1.0.1/24
              proxy-arp policy "proxyARPpolicy"
              local proxy-arp
              remote proxy-arp
            exit
        exit
        static-route 6.5.0.0/24 next-hop 10.1.1.2
        spoke-sdp 2 create
        exit
    no shutdown
    exit
...
----------------------------------------------
*A:ALU-1>config>service#
```

Use the following CLI syntax to configure interface parameters for the VPRN service.

CLI Syntax:

```
config>service# vprn service-id [customer customer-id]
  [create]
     interface ip-int-name
          address if-ip-address
          allow-directed-broadcasts
          arp-timeout
          bfd transmit-interval [receive receive-
            interval] [multiplier multiplier] [type np]
          description description-string
          dhcp
               description description-string
               option
                    action {replace | drop | keep}
                    circuit-id [ascii-tuple | ifindex |
                         sap-id | vlan-ascii-tuple]
                    remote-id [mac | string string]
                    vendor-specific-option
                         client-mac-address
                         sap-id
                         service-id
                         string text
                         system-id
                    server server1 [server2...(up to 8
                         max)]
               no shutdown
               trusted
          icmp
               mask-reply
               ttl-expired [number seconds]
               unreachables
          ip-mtu octets
          ipcp
               dns ip-address [secondary ip-address]
               dns secondary ip-address
               peer-ip-address ip-address
          local proxy-arp
          mac ieee-address
          proxy-arp-policy policy-name [policy-
            name...(up to 5 max)]
          remote-proxy-arp
          secondary {ip-address/mask | ip-address
            netmask} [broadcast all-ones | host-ones]
            [igp-inhibit]
          no shutdown
          static-arp ip-address ieee-mac-address
          static-arp ieee-mac-address unnumbered
          unnumbered {ip-int-name | ip-address}
     no shutdown
```

Example:
```
A:ALU-41>config>service# vprn 4
A:ALU-41>config>service>vprn$ interface "vprn_interface"
A:ALU-41>config>service>vprn>if$ address 10.10.3.3/24
A:ALU-41>config>service>vprn>if$ dhcp option
A:ALU-41>config>service>vprn>if>dhcp>option$ circuit-id
  ifindex
A:ALU-41>config>service>vprn>if>dhcp>option$ exit
A:ALU-41>config>service>vprn>if$ ip-mtu 1524
```

The following example displays the VPRN interface creation output.

```
A:ALU-41>config>service>vprn>if# info detail
----------------------------------------
...
              no description
              address 10.10.3.3/24 broadcast host-ones
              no mac
              arp-timeout 14400
              no allow-directed-broadcasts
              icmp
                 mask-reply
                 unreachables 100 10
                 ttl-expired 100 10
              exit
              dhcp
                 shutdown
                 no description
                 option
                     action keep
                     circuit-id ifindex
                     no remote-id
                     no vendor-specific-option
                 exit
                 no server
                 no trusted
              exit
              ip-mtu 1524
              no bfd
              ipcp
                 no peer-ip-address
                 no dns
              exit
              proxy-arp policy "proxyARPpolicy"
              local proxy-arp
              remote proxy-arp
              no shutdown...
```

Configuring a VPRN IPv6 Interface

Use the following CLI syntax to create a VPRN IPv6 interface and to configure optional VPRN IPv6 interface parameters:

CLI Syntax:

```
config>service
     vprn service-id
          interface ip-int-name
               ipv6
                    address ipv6-address/prefix-length
                         [eui-64] [preferred]
                    bfd transmit-interval [receive
                         receive-interval] [multiplier
                         multiplier]
                    dhcp6-relay
                         description description-string
                         option
                              interface-id
                                   interface-id ascii-
                                     tuple
                                   interface-id ifindex
                                   interface-id sap-id
                                   interface-id string
                              remote-id
                         server ipv6z-address...(upto 8
                           max)
                         shutdown
                         source-address ipv6-address
                    dhcp6-server
                         max-nbr-of-leases max-nbr-of-
                           leases
                         prefix-delegation
                              prefix ipv6-address/
                                prefix-length
                                   duid duid [iaid iaid]
                                   preferred-lifetime
                                     seconds
                                   preferred-lifetime
                                     infinite
                                   valid-lifetime seconds
                                   valid-lifetime
                                     infinite
                              no shutdown
                    icmp6
                         packet-too-big number seconds
                         param-problem number seconds
                         time-exceeded number seconds
                         unreachables number seconds
```

```
link-local-address ipv6-address
     [preferred]
local-dhcp-server server-name
     [create]
neighbor ipv6-address mac-address
reachable-time seconds
stale-time seconds
```

(The example below shows only some of the CLI VPRN IPv6 interface commands).

Example:

```
config>service# vprn 20
config>service>vprn# interface "int1"
config>service>vprn>if>ipv6#
config>service>vprn>if>ipv6>address# fd00:1:1:1::/64
config>service>vprn>if>ipv6>icmp6# packet-too-big 100 10
config>service>vprn>if>ipv6>icmp6# param-problem 100 10
config>service>vprn>if>ipv6>icmp6# time-exceeded 100 10
config>service>vprn>if>ipv6>icmp6# unreachables 100 10
config>service>vprn>if>ipv6>icmp6# exit
config>service>vprn>if>ipv6>neighbor#
 FE80::203:FAFF:FE78:5 C88 00:16:4d:50:17:a3
config>service>vprn>if>ipv6>reachable-time# 30
config>service>vprn>if>ipv6>stale-time# 14400
config>service>vprn>if>ipv6># exit
config>service>vprn>if># exit
config>service>vprn# exit
```

The following example displays a VPRN IPv6 interface configuration:

```
A:ALU-B>config>service>vprn 20# info detail
----------------------------------------------
.....
               ipv6
                   icmp6
                       packet-too-big 100 10
                       param-problem 100 10
                       time-exceeded 100 10
                       unreachables 100 10
                   exit
                   address fd00:1:1:1::/64
                   reachable-time 30
                   stale-time 14400
                   no dhcp6-relay
                   no local-dhcp-server
                   neighbor fe80::203:faff:fe78:5c88 00:16:4d:50:17:a3
                   no bfd
               exit
.....
```

Configuring VPRN Interface Routed VPLS IPv6 Parameters

Use the following CLI syntax to configure VPRN interface routed VPLS IPv6 parameters:

CLI Syntax:

```
config>service# vprn service-id [customer customer-id]
  [create]
     interface ip-int-name
          vpls service-name create
               ingress
                    v6-routed-override-filter ipv6-
                         filter-id
               [no] shutdown
```

Example:

```
A:ALU-41>config>service# vprn 20
A:ALU-41>config>service>vprn$ interface
 "vprn20_interface"
A:ALU-41>config>service>vprn>if$ vpls 2/2/2:1 create
A:ALU-41>config>service>vprn>if>vpls$ ingress
A:ALU-41>config>service>vprn>if>vpls>ingress$ v6-
 routed-override-filter 44
A:ALU-41>config>service>vprn>if>vpls>ingress$ exit
A:ALU-41>config>service>vprn>if>vpls$ exit
A:ALU-41>config>service>vprn>if$ exit
```

Configuring VPRN Interface SAP Parameters

A SAP is a combination of a port and encapsulation parameters that identify the service access point on the interface and within the 7705 SAR. Each SAP must be unique within a router. A SAP cannot be defined if the interface `loopback` command is enabled.

When configuring VPRN interface SAP parameters, a default QoS policy is applied to each ingress and egress SAP. 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.

A VPRN interface SAP is supported on the following ports and adapter cards:

- port in access mode with PPP or MLPPP encapsulation (by setting the port's channel-group encap-type to be ipcp):
  - → any T1/E1 ASAP port or bundle on a 16-port T1/E1 ASAP Adapter card or on a 32-port T1/E1 ASAP Adapter card on a 7705 SAR-8 or 7705 SAR-18:
    - fractional T1/E1
    - clear channel T1/E1
  - → any T1/E1 ASAP port or bundle on a 7705 SAR-F, 7705 SAR-X, or on a 7705 SAR-M or 7705 SAR-A (variants with T1/E1 ports):
    - fractional T1/E1
    - clear channel T1/E1
- port in access mode:
  - → any Ethernet port (null, dot1q, or qinq) on:
    - an 8-port Ethernet Adapter card, version 1 or version 2, or an 8-port Gigabit Ethernet Adapter card on a 7705 SAR-8
    - a 6-port Ethernet 10Gbps Adapter card on a 7705 SAR-8 Shelf V2 with CSMv2
    - a 6-port Ethernet 10Gbps Adapter card, 8-port Ethernet Adapter card, version 2, or an 8-port Gigabit Ethernet Adapter card, or a 10-port 1GigE/ 1-port 10GigE X-Adapter card on a 7705 SAR-18
  - → any Ethernet port (null, dot1q, or qinq) on a 7705 SAR-F, 7705 SAR-M (all variants), 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A (both variants), 7705 SAR-W, 7705 SAR-Wx (all variants), or 7705 SAR-X
  - → DSL module ports (null, dot1q, or qinq) on a 7705 SAR-Wx
  - → GPON module ports and DSL module ports (null or dot1q) when the module is installed in a 7705 SAR-M (variants with module slots)

Note:

- The 10-port 1GigE/1-port 10GigE X-Adapter card supports qinq only on version 2 when it is in 10-port 1GigE mode.
- The 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X support qinq only when the card is in access mode.

See VPRN Services Command Reference for CLI commands and syntax.

The following examples show the configuration of a VPRN interface SAP for:

- an access port on a 16-port T1/E1 ASAP Adapter card
- an MLPPP bundle on an access port on a 16-port T1/E1 ASAP Adapter card

```
*A:ALU-1>config>service# info
----------------------------------------------
...
    vprn 1 customer 1 create
        vrf-import "vrfImpPolCust1"
        vrf-export "vrfExpPolCust1"
        autonomous-system 10000
        route-distinguisher 10001:1
        auto-bind ldp
        vrf-target target:10001:1
        interface "to-ce1" create
            address 11.1.0.1/24
            sap 1/1/10:1 create
                ingress
                    qos 100
                    filter ip 6
                exit
                egress
                    qos 1010
                exit
            exit
        exit
        static-route 6.5.0.0/24 next-hop 10.1.1.2
        spoke-sdp 2 create
        exit
        no shutdown
    exit
...
----------------------------------------------
*A:ALU-1>config>service#
```

```
*A:ALU-1>config>service>vprn# info
----------------------------------------------
    description "Sasha's test VPRN for PPP SAPs"
    route-distinguisher 65101:25
    vrf-target target:65101:25
    interface "int_to_ixia" create
        address 210.177.253.1/24
        sap 1/2/4:25 create
        exit
    exit
    interface "to_ce2_ppp" create
        address 11.25.1.1/30
        bfd 100 receive 100 multiplier 3
        ipcp
            peer-ip-address 11.25.1.2
        exit
        sap 1/1/2.24 create
        exit
    exit
    interface "to_ce2_mlppp" create
        address 11.25.2.1/24
        bfd 100 receive 100 multiplier 3
        ipcp
            peer-ip-address 11.25.2.11
            dns 2.2.2.2 secondary 3.3.3.3
        exit
        sap bundle-ppp-1/1.1 create
        exit
    exit
    interface "to_ce2_eth" create
         address 11.25.3.1/24
         sap 1/2/1:25 create
         exit
    exit
    static-route 210.177.35.0/24 next-hop 210.177.253.100
    static-route 210.177.36.0/24 next-hop 210.177.253.100
    static-route 210.177.45.0/24 next-hop 11.25.1.2
    static-route 210.177.46.0/24 next-hop 11.25.2.2
    static-route 210.177.46.0/24 next-hop 11.25.3.2 disable
    no shutdown
----------------------------------------------
*A:ALU-1>config>service>vprn#
```

Configuring VPRN Interface SAP IPv6 Parameters

Use the following CLI syntax to configure VPRN interface SAP IPv6 parameters:

CLI Syntax:
```
config>service# vprn service-id [customer customer-id]
 [create]
    interface ip-int-name
         sap sap-id create
              ingress
                   filter ipv6 ipv6-filter-id
              [no] shutdown
```

Example:
```
A:ALU-41>config>service# vprn 20
A:ALU-41>config>service>vprn$ interface
 "vprn20_interface"
A:ALU-41>config>service>vprn>if$ sap 1/1/10:1 create
A:ALU-41>config>service>vprn>if>sap$ ingress
A:ALU-41>config>service>vprn>if>sap>ingress$ filter ipv6
 78
A:ALU-41>config>service>vprn>if>sap>ingress$ exit
A:ALU-41>config>service>vprn>if>sap$ exit
```

Configuring VPRN Interface Spoke SDP Parameters

Use the following CLI syntax to configure VPRN interface spoke SDP parameters:

CLI Syntax:
```
config>service# vprn service-id [customer customer-id]
 [create]
    interface ip-int-name
         spoke-sdp sdp-id:vc-id [create]
              egress
                   vc-label egress-vc-label
              ingress
                   filter ip ip-filter-id
                   vc-label ingress-vc-label
              [no] shutdown
```

Example:
```
A:ALU-41>config>service# vprn 6
A:ALU-41>config>service>vprn$ interface
 "vprn6_interface"
A:ALU-41>config>service>vprn>if$ spoke-sdp 7:8 create
A:ALU-41>config>service>vprn>if>spoke-sdp$ ingress
A:ALU-41>config>service>vprn>if>spoke-sdp>ingress$
 filter ip 78
```

```
A:ALU-41>config>service>vprn>if>spoke-sdp>ingress$ vc-
 label 7788
```

The following example displays the VPRN interface spoke SDP creation output.

```
A:ALU-41>config>service>vprn>if>spoke SDP# info detail
----------------------------------------
...
          no description
          egress
              no vc-label
          ingress
              filter ip 78
              vc-label 7788
          exit
          no shutdown
```

Configuring VPRN Interface Spoke SDP IPv6 Parameters

Use the following CLI syntax to configure VPRN interface spoke SDP IPv6 parameters:

CLI Syntax:

```
config>service# vprn service-id [customer customer-id]
 [create]
    interface ip-int-name
         spoke-sdp sdp-id:vc-id [create]
              egress
                   filter ipv6 ipv6-filter-id
              ingress
                   filter ipv6 ipv6-filter-id
              [no] shutdown
```

Example:

```
A:ALU-41>config>service# vprn 10
A:ALU-41>config>service>vprn$ interface
 "vprn10_interface"
A:ALU-41>config>service>vprn>if$ spoke-sdp 8:9 create
A:ALU-41>config>service>vprn>if>spoke-sdp$ egress
A:ALU-41>config>service>vprn>if>spoke-sdp>egress$
 filter ipv6 88
A:ALU-41>config>service>vprn>if>spoke-sdp>egress$ exit
A:ALU-41>config>service>vprn>if>spoke-sdp$ ingress
A:ALU-41>config>service>vprn>if>spoke-sdp>ingress$
 filter ipv6 89
A:ALU-41>config>service>vprn>if>spoke-sdp>ingress$ exit
A:ALU-41>config>service>vprn>if>spoke-sdp$ exit
```

Configuring VRRP

Configuring VRRP policies and instances on service interfaces is optional. The basic owner and non-owner VRRP configurations on a VPRN interface must specify the backup `ip-address` parameter.

VRRP helps eliminate the single point of failure in a routed environment by using virtual router IP address(es) shared between two or more routers connecting the common domain. VRRP provides dynamic failover of the forwarding responsibility to the backup router if the master becomes unavailable.

The VRRP implementation allows one master per IP subnet. All other VRRP instances in the same domain must be in backup mode.

For further information about VRRP CLI syntax and command descriptions, see the VPRN Services Command Reference.

The following displays a VPRN interface VRRP owner configuration:

```
config>service>vprn> info
#------------------------------------------
interface "vrrpowner"
address 10.10.10.23/24
vrrp 1 owner
backup 10.10.10.23
authentication-key "testabc"
exit
exit
#------------------------------------------
config>service>vprn#
```

Configuring a Security Zone within a VPRN

To configure NAT or firewall security functionality, you must:

- configure a NAT or firewall security profile and policy in the `config>security` context
  - → in the `config>security>profile` context, specify the timeouts for the TCP/UDP/ICMP protocols and configure logging and application assurance parameters. This step is optional. If you do not configure the profile, a default profile is assigned.
  - → in the `config>security>policy` context, configure a security policy, specify the match criteria and the action to be applied to a packet if a match is found.
- configure a security zone and apply the policy ID to the zone, as shown in the CLI syntax below

CLI Syntax:

```
config>service
     vprn service-id [customer customer-id] [create]
     abort
     begin
     commit
     zone zone-id [create]
          description description-string
          interface ip-int-name [create]
          name zone-name
          nat
               pool pool-id [create]
                    description description-string
                    direction {zone-outbound | zone-
                         inbound | both}
                    entry entry-id [create]
                         ip-address ip-address [to ip-
                           address] interface ip-int-
                           name
                         port port [to port] interface
                           ip-int-name
                    name pool-name
          policy policy-id | policy-name
          shutdown
```

The following example displays a NAT zone configuration output.

```
A:ALU-B>config>service>vprn# info
----------------------------------------------
        configure
            service vprn 1 create
                zone 1 create
                begin
                    name "VPRN zone"
                    description "uplink zone from private"
                    interface vprn-100-10.30.10.1
                    exit
                    nat
                        pool 1 create
                            description "pool 1"
                            direction zone-inbound
                            exit
                            entry 1 create
                                ip-addr interface vprn-100-10.30.10.1
                            exit
                        exit
                    exit
                    policy 1 nat pool 1
                    commit
                exit
                no-shutdown
----------------------------------------------
A:ALU-B>config>service>ies#
```

Configuring VPRN Router Advertisement

Use the following CLI syntax to enable VPRN router advertisement on all IPv6-enabled interfaces and to configure optional router advertisement parameters:

CLI Syntax:

```
config>service
     vprn service-id
          router-advertisement
               interface ip-int-name
                    current-hop-limit number
                    managed-configuration
                    max-advertisement-interval seconds
                    min-advertisement-interval seconds
                    mtu mtu-bytes
                    other-stateful-configuration
                    prefix ipv6-prefix/prefix-length
                         autonomous
                         on-link
                         preferred-lifetime {seconds |
                           infinite}
                         valid-lifetime {seconds |
                           infinite}
```

```
                                    reachable-time milli-seconds
                                    retransmit-time milli-seconds
                                    router-lifetime seconds
                                    no shutdown
```

Example:

```
config>service# vprn 1
config>service>vprn# router-advertisement
config>service>vprn>router-advert# interface "int1"
config>service>vprn>router-advert>if# prefix
 fd00:1:1:1::/64
config>router>router-advert>if>prefix# autonomous
config>router>router-advert>if>prefix# on-link
config>router>router-advert>if>prefix# preferred-
 lifetime 206800
config>router>router-advert>if>prefix# valid-lifetime
 1502000
config>router>router-advert>if>prefix# exit
config>router>router-advert>if# exit
config>router>router-advert# exit
```

The following example displays a VPRN router advertisement configuration:

```
A:ALU-A>config>service# info detail
----------------------------------------
            interface "n1"
                prefix 3::/64
                exit
                no shutdown
----------------------------------------
A:ALU-A>config>router>router-advert# interface n1
A:ALU-A>config>router>router-advert>if# prefix fd00:1:1:1::/64
A:ALU-A>config>router>router-advert>if>prefix# into detail
----------------------------------------

                    autonomous
                    on-link
                    preferred-lifetime 604800
                    valid-lifetime 2592000
----------------------------------------
A:ALU-A>config>router>router-advert>if>prefix#
```

Service Management Tasks

This section discusses the following service management tasks:

- Modifying VPRN Service Parameters
- Deleting a VPRN Service
- Disabling a VPRN Service
- Re-enabling a VPRN Service

Modifying VPRN Service Parameters

Use the CLI syntax to modify VPRN parameters (VPRN Services Command Reference).

The following example displays the VPRN service creation output.

```
*A:ALU-1>config>service# info
----------------------------------------------
...
    vprn 1 customer 1 create
        shutdown
        vrf-import "vrfImpPolCust1"
        vrf-export "vrfExpPolCust1"
        maximum-routes 2000
        autonomous-system 10000
        route-distinguisher 10001:1
        interface "to-ce1" create
            address 10.1.1.1/24
            sap 1/1/10:1 create
            exit
        exit
        static-route 6.5.0.0/24 next-hop 10.1.1.2
        spoke-sdp 2 create
        exit
    exit
...
----------------------------------------------
*A:ALU-1>config>service>vprn#
```

Deleting a VPRN Service

A VPRN service cannot be deleted until SAPs, interface spoke SDPs, and interfaces are shut down and deleted. If protocols and/or a service 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
          [no] sap sap-id
          [no] spoke-sdp sdp-id:vc-id
          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
```

```
*A:ALU-1>config>service# info
----------------------------------------------
...
    vprn 1 customer 1 create
        shutdown
        vrf-import "vrfImpPolCust1"
        vrf-export "vrfExpPolCust1"
        autonomous-system 10000
        route-distinguisher 10001:1
        auto-bind ldp
        vrf-target target:10001:1
        interface "to-ce1" create
            address 11.1.0.1/24
            sap 1/1/10:1 create
                ingress
                    qos 100
                    filter ip 6
                exit
```

```
                egress
                    qos 1010
                exit
            exit
        exit
        static-route 6.5.0.0/24 next-hop 10.1.1.2
        spoke-sdp 2 create
        exit
    exit
...
----------------------------------------------
*A:ALU-1>config>service#
```

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

Command Hierarchies

- Configuration Commands
  - → VPRN Service Configuration Commands
  - → BGP Commands
  - → OSPF Commands
  - → RIP Commands
  - → VPRN Security Zone Configuration Commands
  - → Router Advertisement Commands
  - → Local DHCP and DHCPv6 Server Commands
  - → Interface Commands
  - → IPv6 Interface Commands
  - → Interface DHCP Commands
  - → Interface ICMP Commands
  - → Interface SAP IPSec-Tunnel Commands
  - → Routed VPLS Commands
  - → Interface VRRP Commands
  - → TWAMP Light Commands
- Show Commands
- Clear Commands
- Debug Commands

Configuration Commands

VPRN Service Configuration Commands

```
config
    — service
        — vprn service-id [customer customer-id] [create]
        — no vprn service-id
            — aggregate ip-prefix/ip-prefix-length [summary-only]
            — no aggregate ip-prefix/ip-prefix-length
            — auto-bind {ldp | gre | rsvp-te | mpls}
            — no auto-bind
            — autonomous-system as-number
            — no autonomous-system
            — [no] bgp
            — description description-string
            — no description
            — dhcp      (refer to the 7705 SAR OS Router Configuration Guide, "Local DHCP
                Server Commands")
            — dhcp6      (see Local DHCP and DHCPv6 Server Commands)
            — ecmp max-ecmp-routes
            — no ecmp
            — encryption-keygroup keygroup-id direction {inbound | outbound}
            — no encryption-keygroup direction {inbound | outbound}
            — grt-lookup
                — [no] enable-grt-local-management-only
                — [no] export-grt-rib-only policy-name [policy-name...(up to 10 max)]
                — [no] export-limit num-routes
                — [no] export-v6-limit num-routes
            — interface ip-int-name
            — interface ip-int-name tunnel create      (see Service Interface Tunnel Commands)
            — no interface ip-int-name
            — ipsec      (see the IPSec Command Reference)
            — maximum-ipv6-routes number [log-only] [threshold percent]
            — no maximum-ipv6-routes
            — maximum-routes number [log-only] [threshold percent]
            — no maximum-routes
            — [no] ospf
            — route-distinguisher [rd]
            — no route-distinguisher
            — router-id ip-address
            — no router-id
            — service-name service-name
            — no service-name
            — sgt-qos
                — application dscp-app-name dscp {dscp-value | dscp-name}
                — application dot1p-app-name dot1p dot1p-priority
                — no application {dscp-app-name | dot1p-app-name}
                — dscp dscp-name fc fc-name
                — no dscp dscp-name
```

```
— [no] shutdown
— snmp-community community-name [hash | hash2] [version SNMP-version]
— no snmp-community community-name [hash | hash2]
— source-address
    — application {app [ping | ptp | ssh | telnet | traceroute]} | {[ip-int-name | ip-
      address]}
    — no application app {[ping | ptp | ssh | telnet | traceroute]}
    — application6 app {[ping | telnet | ssh | traceroute]} | ipv6-address
    — no application6 app {[ping | telnet | ssh | traceroute]}
— [no] spoke-sdp sdp-id
    — [no] shutdown
— [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] [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
— type hub
— no type
— vrf-export policy-name [policy-name...(up to 5 max)]
— no vrf-export
— vrf-import policy-name [policy-name...(up to 5 max)]
— no vrf-import
— vrf-target {ext-community | {[export ext-community] [import ext-community]}}
— no vrf-target
— zone zone-id [create]
```

BGP Commands

```
config
    — service
        — vprn
            — [no] bgp
                — [no] advertise-inactive
                — [no] aggregator-id-zero
                — [no] as-override
                — authentication-key [authentication-key | hash-key] [hash | hash2]
                — no authentication-key
                — [no] backup-path [ipv4]
                — best-path-selection
                    — always-compare-med {zero | infinity}
                    — no always-compare-med
                    — [no] as-path-ignore [ipv4] [ipv6]
                — [no] bfd-enable
                — connect-retry seconds
                — no connect-retry
                — [no] damping
                — description description-string
```

- **no description**
- **disable-communities** [**standard**] [**extended**]
- **no disable-communities**
- [**no**] **disable-fast-external-failover**
- [**no**] **enable-bgp-vpn-backup** [**ipv4**]
- [**no**] **enable-peer-tracking**
- **export** *policy-name* [*policy-name*...(up to 5 max)]
- **no export**
- [**no**] **family** [**ipv4**] [**ipv6**]
- [**no**] **graceful-restart**
  - **stale-routes-time** *time*
  - **no stale-routes-time**
- **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**
- **multipath** *max-paths*
- **no multipath**
- **preference** *preference*
- **no preference**
- [**no**] **rapid-withdrawal**
- [**no**] **remove-private** [**limited**]
- **router-id** *ip-address*
- **no router-id**
- [**no**] **shutdown**
- [**no**] **group** *name* [**dynamic-peer**]
  - [**no**] **advertise-inactive**
  - [**no**] **aggregator-id-zero**
  - [**no**] **as-override**
  - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
  - **no authentication-key**
  - [**no**] **bfd-enable**
  - **connect-retry** *seconds*
  - **no connect-retry**
  - [**no**] **damping**
  - **description** *description-string*
  - **no description**
  - **disable-communities** [**standard**] [**extended**]

- **no disable-communities**
- [**no**] **disable-fast-external-failover**
- [**no**] **enable-peer-tracking**
- **export** *policy-name* [*policy-name*...(up to 5 max)]
- **no export**
- [**no**] **family** [**ipv4**] [**ipv6**]
- [**no**] **graceful-restart**
  - **stale-routes-time** *time*
  - **no stale-routes-time**
- **hold-time** *seconds* [**strict**]
- **no hold-time**
- **import** *policy-name* [*policy-name*...(up to 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**] **next-hop-self**
- [**no**] **passive**
- **peer-as** *as-number*
- **no peer-as**
- **preference** *preference*
- **no preference**
- **prefix-limit** *limit*
- **no prefix-limit**
- [**no**] **remove-private** [**limited**]
- **ttl-security** *min-ttl-value*
- **no ttl-security**
- [**no**] **shutdown**
- [**no**] **neighbor** *ip-address*
  - [**no**] **advertise-inactive**
  - [**no**] **aggregator-id-zero**
  - [**no**] **as-override**
  - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
  - **no authentication-key**
  - [**no**] **bfd-enable**
  - **connect-retry** *seconds*
  - **no connect-retry**
  - **description** *description-string*

- [**no**] **damping**
- **no description**
- **disable-communities** [**standard**] [**extended**]
- **no disable-communities**
- [**no**] **disable-fast-external-failover**
- [**no**] **enable-peer-tracking**
- **export** *policy-name* [*policy-name...*(up to 5 max)]
- **no export**
- [**no**] **family** [**ipv4**] [**ipv6**]
- [**no**] **graceful-restart**
  - **stale-routes-time** *time*
  - **no stale-routes-time**
- **hold-time** *seconds* [**strict**]
- **no hold-time**
- **import** *policy-name* [*policy-name...*(up to 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**] **next-hop-self**
- [**no**] **passive**
- **peer-as** *as-number*
- **no peer-as**
- **preference** *preference*
- **no preference**
- **prefix-limit** *limit*
- **no prefix-limit**
- [**no**] **remove-private** [**limited**]
- **ttl-security** *min-ttl-value*
- **no ttl-security**
- [**no**] **shutdown**

OSPF Commands

```
config
    — service
        — vprn
            — [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
                        — [no] advertise-subnet
                        — 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
                        — interface-type {broadcast | point-to-point}
                        — no interface-type
                        — lfa-policy-map route-nh-template template-name
                        — no lfa-policy-map
                        — [no] loopfree-alternate-exclude
                        — message-digest-key key-id md5 {key | hash-key | hash2-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] loopfree-alternate-exclude
                    — [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] 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* | *hash2-key*} [**hash** | **hash2**]
  - **no message-digest-key** *key-id*
  - **retransmit-interval** *seconds*
  - **no retransmit-interval**
  - [**no**] **shutdown**
  - **transit-delay** *seconds*
  - **no transit-delay**
- **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**
- [**no**] **loopfree-alternate**
- **loopfree-alternate-exclude** **prefix-policy** *prefix-policy* [*prefix-policy...*(*up to 5 max*)]
- **no loopfree-alternate-exclude**
- **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*
- **reference-bandwidth** [**tbps** *Tera-bps*] [**gbps** *Giga-bps*] [**mbps** *Mega-bps*] [**kbps** *Kilo-bps*]
- **no reference-bandwidth**
- **router-id** *ip-address*
- **no router-id**
- [**no**] **shutdown**
- [**no**] **super-backbone**
- [**no**] **suppress-dn-bit**
- **timers**
  - **lsa-arrival** *lsa-arrival-time*
  - **no lsa-arrival**
  - **lsa-generate** *max-lsa-wait* [*lsa-initial-wait* [*lsa-second-wait*]]
  - **no lsa-generate**
  - **spf-wait** *max-spf-wait* [*spf-initial-wait* [*spf-second-wait*]]
  - **no spf-wait**
- **vpn-domain** *id* {**0005** | **0105** | **0205** | **8005**}
- **no vpn-domain**

- **vpn-tag** *vpn-tag*
- **no vpn-tag**

RIP Commands

Global RIP Commands

config
- **service**
  - **vprn**
    - [**no**] **rip**
      - **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
      - **no authentication-key**
      - **authentication-type** {**none** | **password** | **message-digest-20st**}
      - **no authentication-type**
      - **check-zero** {**enable** | **disable**}
      - **no check-zero**
      - **description** *description-string*
      - **no description**
      - **export** *policy-name* [*policy-name*...(**up to 5 max**)]
      - **no export**
      - [**no**] **group** *name*
      - **import** *policy-name* [*policy-name*...(**up to 5 max**)]
      - **no import**
      - **message-size** *max-num-of-routes*
      - **no message-size**
      - **metric-in** *metric*
      - **no metric-in**
      - **metric-out** *metric*
      - **no metric-out**
      - **preference** *preference*
      - **no preference**
      - [**no**] **propagate-metric**
      - **receive** *receive-type*
      - **no receive**
      - **send** *send-type*
      - **no send**
      - [**no**] **shutdown**
      - **split-horizon** {**enable** | **disable**}
      - **no split-horizon**
      - **timers** *update timeout flush*
      - **no timers**

Group RIP Commands

```
config
    — service
        — vprn
            — [no] rip
                — [no] group name
                    — authentication-key [authentication-key | hash-key] [hash | hash2]
                    — no authentication-key
                    — authentication-type {none | password | message-digest}
                    — no authentication-type
                    — check-zero {enable | disable}
                    — no check-zero
                    — description description-string
                    — no description
                    — export policy-name [policy-name...(up to 5 max)]
                    — no export
                    — import policy-name [policy-name...(up to 5 max)]
                    — no import
                    — message-size max-num-of-routes
                    — no message-size
                    — metric-in metric
                    — no metric-in
                    — metric-out metric
                    — no metric-out
                    — preference preference
                    — no preference
                    — receive receive-type
                    — no receive
                    — send send-type
                    — no send
                    — [no] shutdown
                    — split-horizon {enable | disable}
                    — no split-horizon
                    — timers update timeout flush
                    — no timers
                    — [no] neighbor ip-int-name
```

Neighbor RIP Commands

```
config
    — service
        — vprn
            — [no] rip
                — [no] group name
                    — [no] neighbor ip-int-name
                        — authentication-key [authentication-key | hash-key] [hash |
                            hash2]
                        — no authentication-key
                        — authentication-type {none | password | message-digest}
                        — no authentication-type
```

- **check-zero** {**enable** | **disable**}
- **no check-zero**
- **description** *description-string*
- **no description**
- **export** *policy-name* [*policy-name*...(**up to 5 max**)]
- **no export**
- **import** *policy-name* [*policy-name*...(**up to 5 max**)]
- **no import**
- **message-size** *max-num-of-routes*
- **no message-size**
- **metric-in** *metric*
- **no metric-in**
- **metric-out** *metric*
- **no metric-out**
- **preference** *preference*
- **no preference**
- **receive** *receive-type*
- **no receive**
- **send** *send-type*
- **no send**
- [**no**] **shutdown**
- **split-horizon** {**enable** | **disable**}
- **no split-horizon**
- **timers** *update timeout flush*
- **no timers**

VPRN Security Zone Configuration Commands

```
config
    — service
        — vprn service-id [customer customer-id] [create]
        — no vprn service-id
            — zone {zone-id | name} [create]
            — no zone zone-id
                — abort
                — begin
                — commit
                — description description-string
                — no description
                — [no] auto-bind
                — inbound
                    — limit
                        — concurrent-sessions {tcp | udp | icmp | other} sessions
                        — no concurrent-sessions {tcp | udp | icmp | other}
                — [no] interface ip-int-name
                    — [no] shutdown
                — log log-id
                — no log
                — name zone-name
                — no name
                — nat
                    — pool pool-id [create]
                    — no pool pool-id
                        — description description-string
                        — no description
                        — direction {zone-outbound | zone-inbound | both}
                        — no direction
                        — entry entry-id [create]
                        — no entry entry-id
                            — ip-address ip-address [to ip-address] interface ip-int-
                                name
                            — no ip-address
                            — port port [to port]
                            — no port
                        — name pool-name
                        — no name
                — outbound
                    — limit
                        — concurrent-sessions {tcp | udp | icmp | other} sessions
                        — no concurrent-sessions {tcp | udp | icmp | other}
                — policy {policy-id | name}
                — no policy
                — [no] shutdown
```

Router Advertisement Commands

```
config
    — service
        — vprn service-id [customer customer-id] [create]
        — no vprn service-id
            — [no] router-advertisement
                — [no] interface ip-int-name
                    — current-hop-limit number
                    — no current-hop-limit
                    — [no] managed-configuration
                    — max-advertisement-interval seconds
                    — no max-advertisement-interval
                    — min-advertisement-interval seconds
                    — no min-advertisement-interval
                    — mtu mtu-bytes
                    — no mtu
                    — [no] other-stateful-configuration
                    — prefix ipv6-prefix/prefix-length
                    — no prefix
                        — [no] autonomous
                        — [no] on-link
                        — preferred-lifetime {seconds | infinite}
                        — no preferred-lifetime
                        — valid-lifetime{seconds | infinite}
                        — no valid-lifetime
                    — reachable-time milli-seconds
                    — no reachable-time
                    — retransmit-time milli-seconds
                    — no retransmit-time
                    — router-lifetime seconds
                    — no router-lifetime
                    — [no] shutdown
```

Local DHCP and DHCPv6 Server Commands

For complete descriptions of all local DHCP and DHCPv6 server commands, refer to the Router Configuration Guide, “Local DHCP and DHCPv6 Server Commands”.

config
- **service**
  - **vprn**
    - **dhcp**
      - **local-dhcp-server** *server-name* [**create**]
      - **no local-dhcp-server** *server-name*
        - **description** *description-string*
        - **no description**
        - [**no**] **force-renew**
        - **lease-hold-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
        - **no lease-hold-time**
        - **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**
          - **offer-time** [**min** *minutes*] [**sec** *seconds*]
          - **no offer-time**
          - **options**
            - **custom-option** *option-number* **address** *ip-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* [*ipaddress*...(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**

- — **netbios-name-server** *ip-address* [*ip-address*...(up to 4 max)]
- — **no netbios-name-server**
- — **netbios-node-type** *netbios-node-type*
- — **no netbios-node-type**
- — **subnet** {*ip-address*/*mask* | *ip-address netmask*} [**create**]
- — **no subnet** {*ip-address*/*mask* | *ip-address netmask*}
- — **address-range** *start-ip-address end-ip-address*
- — **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* [**percent**] [**event-when-depleted**]
- — **no minimum-free**
- — **options**
- — **custom-option** *option-number* **address** *ip-address* [*ipaddress*...(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**
- — **subnet-mask** *ip-address*
- — **no subnet-mask**
- — [**no**] **shutdown**
- — [**no**] **use-gi-address**
- — [**no**] **use-pool-from-client**

— **dhcp6**

- — **local-dhcp-server** *server-name* [**create**]
- — **no local-dhcp-server** *server-name*
- — **description** *description-string*
- — **no description**
- — [**no**] **force-renew**
- — **lease-hold-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- — **no lease-hold-time**
- — **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**
- **offer-time** [**min** *minutes*] [**sec** *seconds*]
- **no offer-time**
- **options**
  - **custom-option** *option-number* **address** *ip-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* [*ipaddress*...(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**
  - **netbios-name-server** *ip-address* [*ip-address*...(up to 4 max)]
  - **no netbios-name-server**
  - **netbios-node-type** *netbios-node-type*
  - **no netbios-node-type**
- **prefix** *ipv6-address/prefix-length* [**pd**] [**wan-host**] [**create**]
- **no prefix** *ipv6-address/prefix-length*
- **subnet** {*ip-address*/*mask* | *ip-address netmask*} [**create**]
- **no subnet** {*ip-address*/*mask* | *ip-address netmask*}
  - **address-range** *start-ip-address end-ip-address*
  - **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* [**percent**] [**event-when-depleted**]
  - **no minimum-free**
  - **options**
    - **custom-option** *option-number* **address** *ip-address* [*ipaddress*...(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
                                    — subnet-mask ip-address
                                    — no subnet-mask
                                — subnet-binding key {sys-id-svc-id | sys-id | string} unbind-
                                    delay [hrs hours] [min mins] [sec secs]
                                — no subnet-binding
                            — [no] shutdown
                            — use-link-address [scope scope]
                            — no use-link-address
                            — [no] use-pool-from-client
                            — user-ident user-ident
                            — no user-ident
```

Interface Commands

```
config
    — service
        — vprn
            — interface ip-int-name
            — interface ip-int-name tunnel create      (see Service Interface Tunnel Commands)
            — no interface ip-int-name
                — address {ip-address/mask | ip-address netmask} [broadcast {all-ones |
                    host-ones}]
                — no address {ip-address/mask | ip-address netmask}
                — allow-directed-broadcasts
                — no allow-directed-broadcasts
                — arp-retry-timer ms-timer
                — no arp-retry-timer
                — arp-timeout seconds
                — no arp-timeout
                — bfd transmit-interval [receive receive-interval] [multiplier multiplier]
                    [echo-receive echo-interval] [type np]
                — no bfd
                — description description-string
                — no description
                — dhcp
                — icmp
                — ip-mtu octets
                — no ip-mtu
                — ipcp
                    — dns ip-address [secondary ip-address]
                    — dns secondary ip-address
                    — no dns [ip-address] [secondary ip-address]
                    — peer-ip-address ip-address
                    — no peer-ip-address
                — l4-load-balancing {includeL4 | excludeL4}
                — no l4-load-balancing
                — [no] local-dhcp-server local-server-name
                — [no] local-proxy-arp
```

- [**no**] **loopback**
- **mac** *ieee-address*
- **no mac** [*ieee-address*]
- **proxy-arp-policy** *policy-name* [*policy-name*...(up to 5 max)]
- [**no**] **remote-proxy-arp**
- [**no**] **sap** *sap-id* [**create**]
  - **accounting-policy** *acct-policy-id*
  - **no accounting-policy** [*acct-policy-id*]
  - [**no**] **collect-stats**
  - **description** *description-string*
  - **no description**
  - **egress**
    - **agg-rate-limit** *agg-rate* [**cir** *cir-rate*]
    - **no agg-rate-limit**
    - **filter ip** *ip-filter-id*
    - **no filter ip** [*ip-filter-id*]
    - **filter ipv6** *ipv6-filter-id*
    - **no filter ipv6** [*ipv6-filter-id*]
    - **filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
    - **no filter** [[**ip** [*ip-filter-id*]] [[**ipv6** [*ipv6-filter-id*]]
    - [**no**] **qinq-mark-top-only**
    - **qos** *policy-id*
    - **no qos** [*policy-id*]
    - **scheduler-mode** {**4-priority** | **16-priority**}
    - [**no**] **shaper-group** *shaper-group-name* [**create**]
  - **ingress**
    - **agg-rate-limit** *agg-rate* [**cir** *cir-rate*]
    - **no agg-rate-limit**
    - **filter ip** *ip-filter-id*
    - **no filter ip** [*ip-filter-id*]
    - **filter ipv6** *ipv6-filter-id*
    - **no filter ipv6** [*ipv6-filter-id*]
    - **filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
    - **no filter** [[**ip** [*ip-filter-id*]] [[**ipv6** [*ipv6-filter-id*]]
    - **match-qinq-dot1p** {**top** | **bottom**}
    - **no match-qinq-dot1p**
    - **qos** *policy-id*
    - **no qos** [*policy-id*]
    - **scheduler-mode** {**4-priority** | **16-priority**}
    - [**no**] **shaper-group** *shaper-group-name* [**create**]
  - [**no**] **shutdown**
- **secondary** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**igp-inhibit**]
- **no secondary** {*ip-address/mask* | *ip-address netmask*}
- [**no**] **shutdown**
- **spoke-sdp** *sdp-id:vc-id*
- **no spoke-sdp** *sdp-id:vc-id*
  - **egress**
    - **vc-label** *egress-vc-label*
    - **no vc-label** [*egress-vc-label*]
  - **ingress**
    - **filter ip** *ip-filter-id*
    - **filter ipv6** *ipv6-filter-id*
    - **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]

```
                        — vc-label ingress-vc-label
                        — no vc-label [ingress-vc-label]
                    — [no] shutdown
                — static-arp ip-address ieee-address
                — no static-arp ip-address [ieee-address]
                — static-arp ieee-address unnumbered
                — no static-arp [ieee-address] unnumbered
                — [no] teid-load-balancing
                — unnumbered {ip-int-name | ip-address}
                — no unnumbered
                — vpls service-name
                — no vpls
                — vrrp virtual-router-id [owner]
                — no vrrp virtual-router-id
```

IPv6 Interface Commands

```
config
    — service
        — vprn service-id [customer customer-id] [create]
        — no vprn service-id
            — [no] interface ip-int-name
                — [no] ipv6
                    — address ipv6-address/prefix-length [eui-64][preferred]
                    — no address ipv6-address/prefix-length
                    — bfd transmit-interval [receive receive-interval] [multiplier multiplier]
                    — no bfd
                    — [no] dhcp6-relay
                        — description description-string
                        — no description
                        — [no] option
                            — interface-id
                            — interface-id ascii-tuple
                            — interface-id ifindex
                            — interface-id sap-id
                            — interface-id string
                            — no interface-id
                            — [no] remote-id
                        — [no] server ipv6z-address
                        — [no] shutdown
                        — [no] source-address ipv6-address
                    — [no] dhcp6-server
                        — max-nbr-of-leases max-nbr-of-leases
                        — no max-nbr-of-leases
                        — [no] prefix-delegation
                            — [no] prefix ipv6-address/prefix-length
                                — duid duid [iaid iaid]
                                — no duid
                                — preferred-lifetime seconds
                                — preferred-lifetime infinite
                                — no preferred-lifetime
```

```
                                    — valid-lifetime seconds
                                    — valid-lifetime infinite
                                    — no valid-lifetime
                                — [no] shutdown
                        — icmp6
                            — packet-too-big [number seconds]
                            — no packet-too-big
                            — param-problem [number seconds]
                            — no param-problem
                            — time-exceeded [number seconds]
                            — no time-exceeded
                            — unreachables [number seconds]
                            — no unreachables
                        — link-local-address ipv6-address [preferred]
                        — local-dhcp-server server-name [create] (see Local DHCP and
                                DHCPv6 Server Commands)
                        — no local-dhcp-server server-name
                        — neighbor ipv6-address mac-address
                        — no neighbor ipv6-address
                        — [no] reachable-time seconds
                        — no reachable-time
                        — [no] stale-time seconds
                        — no stale-time
```

Interface DHCP Commands

```
config
    — service
        — vprn
            — interface
                — 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 remote-id
                        — [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
```

Interface ICMP Commands

```
config
    — service
        — vprn
            — interface
                — icmp
                    — [no] mask-reply
                    — ttl-expired number seconds
                    — no ttl-expired [number seconds]
                    — unreachables number seconds
                    — no unreachables [number seconds]
```

Interface SAP IPSec-Tunnel Commands

For complete descriptions of all VPRN IPSec commands, refer to the IPSec Command Reference.

```
config
    — service
        — vprn
            — interface
                — sap
                    — ipsec-tunnel ipsec-tunnel-name [create]
```

Routed VPLS Commands

```
config
    — service
        — vprn service-id
            — interface ip-interface-name [create]
            — no interface ip-interface-name
                — vpls service-name
                — no vpls
                    — ingress
                        — v4-routed-override-filter ipv4-filter-id
                        — no v4-routed-override-filter
                        — v6-routed-override-filter ipv6-filter-id
                        — no v6-routed-override-filter
```

Interface VRRP Commands

```
config
    — service
        — vprn
            — interface ip-int-name
                — vrrp virtual-router-id [owner]
                — no vrrp virtual-router-id
                — authentication-key [authentication-key | hash-key] [hash | hash2]
                — no authentication-key
                — [no] backup ip-address
                — [no] bfd-enable [base | service-id] interface interface-name dst-ip
                    ip-address
                — init-delay seconds
                — no init-delay
                — mac ieee-address
                — no mac
                — [no] master-int-inherit
                — message-interval {[seconds] [milliseconds milliseconds]}
                — no message-interval
                — [no] ping-reply
                — policy vrrp-policy-id
                — no policy
                — [no] preempt
                — priority priority
                — no priority
                — [no] shutdown
                — [no] ssh-reply
                — [no] standby-forwarding
                — [no] telnet-reply
                — [no] traceroute-reply
```

TWAMP Light Commands

```
config
    — service
        — vprn
            — twamp-light
                — [no] reflector
                    — description description-string
                    — [no] prefix ip-prefix/prefix-length [create]
                        — description description-string
                    — udp-port number
                    — no udp-port
                    — [no] shutdown
```

Show Commands

show
— **service**
— **egress-label** *start-label* [*end-label*]
— **id** *service-id*
— **all**
— **arp** [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*]
— **base**
— **dhcp**
— **statistics** [**interface** *ip-int-name* | *ip-address*]
— **summary** [**interface** *interface-name* | **saps**]
— **interface** [{[*ip-address* | *ip-int-name*] [*interface-type*] [**detail**] [**family**]} | **summary**]
— **sap** [*sap-id* [**detail**]]
— **sap-using description (see the VLLServices Command Reference, Show Commands)**
— **sdp** {[*sdp-id*[**:***vc-id*] | **far-end** *ip-address*]} [**detail**]
— **sdp** [*sdp-id*[**:***vc-id*]
— **twamp-light**
— **ingress-label** *start-label* [*end-label*]
— **service-using** **vprn** [**sdp** *sdp-id*] [**customer** *customer-id*]

show
— **router** [*service-id*]
— **aggregate** [**active**]
— **arp** [*ip-address* | *ip-int-name* | **mac** *ieee-mac-address*] [**sdp** *sdp-id:vc-id*] [**summary**]
— **bgp**
— **damping** [*ip-prefix*[*/ip-prefix-length*]] [*damp-type*] [**detail**] [**ipv4**]
— **damping** [*ip-prefix*[*/ip-prefix-length*]] [*damp-type*] [**detail**] **ipv6**
— **damping** [*ip-prefix*[*/ip-prefix-length*]] [*damp-type*] [**detail**] **vpn-ipv4**
— **damping** [*ip-prefix*[*/ip-prefix-length*]] [*damp-type*] [**detail**] **vpn-ipv6**
— **group** [*name*] [**detail**]
— **inter-as-label** (see **Note:**)
— **neighbor** [*ip-address* [[*family*] *filter1* [**brief**]]]
— **neighbor** [*as-number* [[*family*] *filter2*]]
— **neighbor** *ip-address* **orf** [*filter3*]
— **neighbor** *ip-address* **graceful-restart**
— **next-hop** [*family*] [*ip-address*] [**detail**]
— **paths**
— **routes** [*family*] [**received**] [**url** *file-url*]
— **routes** [*family*] [**brief**]
— **routes** [*family*] *prefix* [**detail** | **longer** | **hunt** [**brief**]]
— **routes** [*family*] **community** *comm-id*
— **routes** [*family*] **aspath-regex** *reg-ex*
— **summary** [**all**]
— **summary** [**family** *family*] [**neighbor** *ip-address*]
— **dhcp**
— **statistics** [**interface** *ip-int-name* | *ip-address*]
— **summary**
— **interface** [{[*ip-address* | *ip-int-name*] [**detail**]} | **summary** | **exclude-services**]
— **route-table** [*ip-address* | *mask*] [**longer** | **best**] [**protocol** *protocol*] [**summary**]
— **sgt-qos** [*service-id*]
— **application** [*app-name*] [**dscp** | **dot1p**]

```
            — dscp-map dscp-map [dscp-name]
        — 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]
```

Note: The **inter-as-label** command appears in the **show>router>bgp** command hierarchy; however, it is not applicable in the VPRN BGP context and, if executed, will return empty output.

Clear Commands

```
clear
    — service
        — id service-id
            — arp
            — dhcp
                — statistics [sap sap-id | sdp sdp-id:vc-id | interface {ip-int-name | ip-address}]
            — dhcp6
                — statistics [interface {ip-int-name | ipv6-address}]
            — fdb {all | mac ieee-address | sap sap-id | mesh-sdp sdp-id[:vc-id] | spoke-sdp
              sdp-id:vc-id}
            — mesh-sdp sdp-id[:vc-id] ingress-vc-label
            — mfib
                — statistics {all | ip | mac | group grp-address}
            — spoke-sdp sdp-id:vc-id ingress-vc-label
        — statistics
            — id service-id
                — cem
                — counters
                — mesh-sdp sdp-id[:vc-id] {all | counterss | stp}
                — spoke-sdp sdp-id:vc-id {all | counters}
            — sap sap-id {all | cem | counters }
            — sap-aggregation-group svcId-id:groupName {all | counters}
            — sdp sdp-id keep-alive
```

Debug Commands

```
debug
    — service
        — id service-id
        — [no] id service-id
            — [no] dhcp
                — detail-level {low | medium | high}
                — no detail-level
                — mac ieee-address
                — no mac ieee-address
                — mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}
                — no mode
                — sap sap-id
                — no sap sap-id
                — sdp sdp-id:vc-id
                — no sdp sdp-id:vc-id
            — [no] dhcp6
                — detail-level {low | medium | high}
                — no detail-level
                — mac ieee-address
                — no mac ieee-address
                — mode {all | dropped-only}
                — no mode
                — sap sap-id
                — no sap sap-id
            — event-type {config-change | svc-oper-status-change | sap-oper-status-change |
                sdpbind-oper-status-change}
            — no event-type
            — sap sap-id
            — no sap sap-id
            — sdp sdp-id:vc-id
            — no sdp sdp-id:vc-id
```

Command Descriptions

- VPRN Service Configuration Commands
- Show Service Commands
- Show Router Commands
- Clear Service Commands
- Debug Service Commands

VPRN Service Configuration Commands

- Generic Commands
- Global Commands
- BGP Commands
- OSPF Commands
- RIP Commands
- VPRN Security Configuration Commands
- Router Advertisement Commands
- Local DHCP and DHCPv6 Server Commands
- Interface Commands
- IPv6 Interface Commands
- Interface DHCP Commands
- Interface ICMP Commands
- Interface SAP Commands
- Interface Spoke SDP Commands
- Routed VPLS Commands
- Interface VRRP Commands

Generic Commands

description

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

Context config>service>vprn
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
config>service>vprn>dhcp>local-dhcp-server
config>service>vprn>dhcp>local-dhcp-server>pool
config>service>vprn>interface
config>service>vprn>if>dhcp
config>service>vprn>if>ipv6>dhcp6-relay
config>service>vprn>if>sap
config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor
config>service>vprn>twamp-light>reflector
config>service>vprn>twamp-light>reflector>prefix
config>service>vprn>zone
config>service>vprn>zone>nat>pool

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

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

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

Default no description

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

shutdown

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

Context config>service>vprn
config>service>vprn>bgp
config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor
config>service>vprn>interface
config>service>vprn>if>dhcp
config>service>vprn>if>ipv6>dhcp6-relay
config>service>vprn>if>ipv6>dhcp6-server>prefix-delegation
config>service>vprn>if>spoke-sdp
config>service>vprn>if>sap
config>service>vprn>if>vrrp
config>service>vprn>ospf
config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link
config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor
config>service>vprn>router-advertisement>interface
config>service>vprn>spoke-sdp
config>service>vprn>twamp-light>reflector
config>service>vprn>zone
config>service>vprn>zone>interface

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 can 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 are described below in **Special Cases**.

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

Special Cases **Service Admin State —** bindings to an SDP within the service will be put into the out-of-service state when the service is shut down. While the service is shut down, all customer packets are dropped and counted as discards for billing and debugging purposes.

A service is regarded as operational providing that one IP Interface SAP and one SDP is operational.

Global Commands

vprn

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

Context config>service

Description 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 a separate routing instance from the one used for service network core routing connectivity. VPRN services allow the IP addressing scheme used by the subscriber to overlap with other addressing schemes used by other VPRN services or by the provider and, potentially, the entire Internet.

IP interfaces defined within the context of a 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, which 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. To change the association between service and customer, the service must be deleted and recreated with a new customer association.

Once a service is created, the use of **customer** *customer-id* is optional to navigate into the service configuration context. Attempting to edit a service with an incorrect *customer-id* results in an error.

Multiple VPRN services are created in order 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 a VPRN service ID belong 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 shut down and deleted.

Default n/a

Parameters *service-id* — the unique service identification number or name that identifies the service in the service domain. The ID must be unique to this service and cannot be used for any other service of any type (such as Epipe, Cpipe, IES). However, a VPRN instance in the service provider network can include different *service-ids* on the routers in the network.

Values 1 to 2147483647 or *service-name*

customer-id — an existing customer identification number to be associated with the service. This parameter is required during service creation and is optional for service editing or deleting.

Values 1 to 2147483647

create — keyword is mandatory when creating a VPRN service

aggregate

Syntax **aggregate** *ip-prefix/ip-prefix-length* [**summary-only**]
no aggregate *ip-prefix/ip-prefix-length*

Context config>service>vprn

Description This command creates an aggregate route.

Use this command to group a number of routes with common prefixes into a single entry in the routing table. This reduces the number of routes that need to be advertised by this router and reduces the number of routes in the routing tables of downstream routers.

Both the original components and the aggregated route (source protocol aggregate) are offered to the Routing Table Manager (RTM). Subsequent policies can be configured to assign protocol-specific characteristics, such as the OSPF tag, to aggregate routes.

Multiple entries with the same prefix but a different mask can be configured; routes are aggregated to the longest mask. If one aggregate is configured as 10.0/16 and another as 10.0.0/24, then route 10.0.128/17 would be aggregated into 10.0/16 and route 10.0.0.128/25 would be aggregated into 10.0.0/24. If multiple entries are made with the same prefix and the same mask, the previous entry is overwritten.

The **no** form of the command removes the aggregate.

Default no aggregate

Parameters *ip-prefix/ip-prefix-length* — the destination address of the aggregate route

Values		
	ipv4-prefix	a.b.c.d (host bits must be 0)
	ipv4-prefix-length	0 to 30
	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 to FFFF]H
		d: [0 to 255]D
	ipv6-prefix-length	0 to 128

summary-only — 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.

auto-bind

Syntax **auto-bind** {**ldp** | **gre** | **rsvp-te** | **mpls**}
no auto-bind

Context config>service>vprn

Description This command specifies the type of automatic binding for the SDP assigned to this service. When **auto-bind** is used, it is not required that a spoke-SDP be configured for the service.

Default no auto-bind

Parameters **ldp —** specifies LDP as the automatic binding for the SDP assigned to the service

gre — specifies GRE as the automatic binding for the SDP assigned to the service

rsvp-te — specifies RSVP-TE as the automatic binding for the SDP assigned to the service

mpls — specifies that both LDP and RSVP-TE can be used to resolve the BGP next-hop for VPRN routes in an associated VPRN instance.

autonomous-system

Syntax **autonomous-system** *as-number*
no autonomous-system

Context config>service>vprn

Description This command defines the autonomous system (AS) to be used by this VPN virtual routing/forwarding table (VRF).

The **no** form of the command removes the defined AS from the given VPRN context.

Default no autonomous-system

Parameters *as-number* — specifies the AS number for the VPRN service

Values 1 to 4294967295

ecmp

Syntax **ecmp** *max-ecmp-routes*
no ecmp

Context config>service>vprn

Description This command enables ECMP in the VPRN service context and configures the number of routes for path sharing; for example, the value 2 means two equal-cost routes will be used for cost sharing.

ECMP (Equal-Cost Multipath Protocol) refers to the distribution of packets over two or more outgoing links that share the same routing cost. ECMP provides a fast local reaction to route failures. ECMP is supported on static routes and dynamic (OSPF, IS-IS, and BGP) routes.

ECMP can only be used for routes with the same preference and same protocol. See the static-route command for information on preferences.

When more ECMP routes are available at the best preference than configured in *max-ecmp-routes,* then the lowest next-hop IP address algorithm is used to select the number of routes configured in *max-ecmp-routes*.

The **no** form of the command disables ECMP path sharing. If ECMP is disabled and multiple routes are available at the best preference and equal cost, the route with the lowest next-hop IP address is used.

The **no** form of the command disables ECMP path sharing.

Default no ecmp

Parameters *max-ecmp-routes* — the maximum number of equal-cost routes allowed on this VPRN instance, expressed as a decimal integer. Setting ECMP *max-ecmp-routes* to 1 yields the same result as entering **no ecmp**.

Values 0 to 8

encryption-keygroup

Syntax **encryption-keygroup** *keygroup-id* **direction** {**inbound** | **outbound**}
no encryption-keygroup direction {**inbound** | **outbound**}

Context config>service>vprn

Description This command is used to bind a key group to a VPRN service for inbound or outbound packet processing. When configured in the outbound direction, packets egressing the node use the **active-outbound-sa** associated with the key group configured. When configured in the inbound direction, received packets must be encrypted using one of the valid security associations configured for the key group.

The encryption (enabled or disabled) configured on an SDP used to terminate a Layer 3 spoke SDP of a VPRN always overrides any VPRN-level configuration for encryption.

Encryption is enabled once the outbound direction is configured.

The **no** form of the command removes the key group from the service in the specified direction (inbound or outbound).

Default n/a

Parameters *keygroup-id* — the number of the key group being configured

Values 1 to 15 or *keygroup-name* (up to 64 characters)

direction {**inbound** | **outbound**} — mandatory keywords when binding a key group to a service for a particular direction

grt-lookup

Syntax **grt-lookup**

Context config>service>vprn

Description This command enables the context for global routing table lookup within a VPRN, for in-band node management on the 7705 SAR. The GRT lookup occurs only on the local 7705 SAR system IP address within a VPRN.

enable-grt-local-management-only

Syntax [**no**] **enable-grt-local-management-only**

Context config>service>vprn>grt-lookup

Description This command enables global routing table lookup of the 7705 SAR system IP address on network ingress traffic in a VPRN, so that only management traffic is transported using the VPRN.

On network ingress, when a packet arrives from the transport tunnel to the VPRN, a lookup is performed within the VPRN on the inner customer packet IP header. If the destination IP address in the header matches the local 7705 SAR system IP address, and the **enable-grt-local-management-only** command is configured, then the packet is extracted to the CSM for processing as management traffic. If **enable-grt-local-management-only** is not enabled, the packet is routed using the 7705 SAR VRF FIB.

The **no** form of the command disables the global routing table lookup function for the 7705 SAR system IP address.

Default no enable-grt-local-management-only

export-grt-rib-only

Syntax [**no**] **export-grt-rib-only** *policy-name* [*policy-name...*(up to 10 max)]

Context config>service>vprn>grt-lookup

Description This command enables specific route policies to be exported to the GRT RIB.

On network egress, packets generated from the CSM with a source IP address that matches the 7705 SAR system IP address and the destination IP address of either the far-end 5620 SAM or other management entity must perform a GRT lookup in order to be resolved. A route policy can be configured with the IP address prefix of the far-end management entity and with the action to accept. This policy can be configured under the **config>router> policy-options** context, and can be installed in the GRT RIB using the **export-grt-rib-only** command. The route installed in the GRT RIB will have a next hop of the corresponding VRF tunnel.

This prevents any user data traffic in the GRT data path from leaking into the VPRN, and ensures that only the management traffic originating from the system IP address and the CSM gets transported through the VPRN. The management packets get routed by the corresponding VPRN transport tunnel, which means the VPRN route is leaked into the GRT so the GRT resolves the route using the corresponding VPRN.

Up to ten policies can be exported to the GRT RIB.

The **no** form of the command restores the default of not exporting routes to the GRT RIB.

Default no export-grt-rib-only

Parameters *policy-name* — the name of the route policy to be exported to the GRT RIB

export-limit

Syntax [**no**] **export-limit** *num-routes*

Context config>service>vprn>grt-lookup

Description This command limits the number of IPv4 routes that can be placed in a route policy to be exported to the GRT RIB.

The **no** form of the command reverts to the default of allowing five routes to be placed in a route policy to be exported from the VPRN to the GRT RIB.

Default no export-limit

Parameters *num-routes* — the number of routes per policy to be exported to the GRT RIB (up to 5 routes)

Values 1 to 5

export-v6-limit

Syntax [**no**] **export-v6-limit** *num-routes*

Context config>service>vprn>grt-lookup

Description This command limits the number of IPv6 routes that can be placed in a route policy to be exported to the GRT RIB.

The **no** form of the command reverts to the default of allowing five IPv6 routes to be placed in a route policy to be exported from the VPRN to the GRT RIB.

Default no export-v6-limit

Parameters *num-routes* — the number of IPv6 routes per policy to be exported to the GRT RIB

Values 1 to 5

maximum-ipv6-routes

Syntax **maximum-ipv6-routes** *number* [**log-only**] [**threshold** *percent*]
no maximum-ipv6routes

Context config>service>vprn

Description This command specifies the maximum number of IPv6 routes that can be held within a VPN virtual routing /forwarding (VRF) context. Local, host, and aggregate routes are not counted.

The VPRN service ID must be in a shutdown state before **maximum-ipv6-routes** command parameters can be modified.

If the **log-only** parameter is not specified and the **maximum-ipv6-routes** value is set to a value below the existing number of IPv6 routes in a VRF, then the extra IPv6 routes will not be added to the VRF.

The maximum IPv6 route threshold can dynamically change to increase the number of supported IPv6 routes even when the maximum has already been reached. Protocols will resubmit the IPv6 routes that were initially rejected.

The **no** form of the command disables any limit on the number of IPv6 routes within a VRF context. Issue the **no** form of the command only when the VPRN instance is shut down.

Default no maximum-ipv6routes (0 or disabled)

Parameters *number* — the maximum number of IPv6 routes to be held in a VRF context

Values 1 to 2147483647

log-only — specifies that if the maximum limit is reached, the event only will be logged. The **log-only** parameter does not disable the learning of new IPv6 routes.

percent — the percentage at which a warning log message and SNMP will be used. There are two warning levels: mid-level and high-level. A mid-level warning occurs when the **threshold** *percent* value is reached, and a high-level warning occurs at the halfway level between the maximum number of IPv6 routes and the *percent* value ([max + mid] / 2). For example, if the **maximum-ipv6-routes** number is 100, and *percent* is 60, then the mid-level warning occurs at 60 IPv6 routes, and the high-level warning occurs at 80 IPv6 routes.

Values 0 to 100

maximum-routes

Syntax **maximum-routes** *number* [**log-only**] [**threshold** *percent*]
no maximum-routes

Context config>service>vprn

Description This command specifies the maximum number of IPv4 routes that can be held within a VPN virtual routing /forwarding (VRF) context. Local, host, and aggregate routes are not counted.

The VPRN service ID must be in a shutdown state before **maximum-routes** command parameters can be modified.

If the **log-only** parameter is not specified and the **maximum-routes** value is set to a value below the existing number of IPv4 routes in a VRF, then the extra IPv4 routes will not be added to the VRF.

The maximum IPv4 route threshold can dynamically change to increase the number of supported IPv4 routes even when the maximum has already been reached. Protocols will resubmit the IPv4 routes that were initially rejected.

The **no** form of the command disables any limit on the number of IPv4 routes within a VRF context. Issue the **no** form of the command only when the VPRN instance is shut down.

Default no maximum-routes (0 or disabled)

Parameters *number* — the maximum IPv4 number of routes to be held in a VRF context

Values 1 to 2147483647

log-only — specifies that if the maximum limit is reached, the event only will be logged. The **log-only** parameter does not disable the learning of new routes.

percent — the percentage at which a warning log message and SNMP will be used. There are two warning levels: mid-level and high-level. A mid-level warning occurs when the **threshold** *percent* value is reached, and a high-level warning occurs at the halfway level between the maximum number of IPv4 routes and the *percent* value ([max + mid] / 2). For example, if the **maximum-routes** number is 100, and *percent* is 60, then the mid-level warning occurs at 60 IPv4 routes, and the high-level warning occurs at 80 IPv6 routes.

Values 0 to 100

route-distinguisher

Syntax **route-distinguisher** [*rd*]
no route-distinguisher

Context config>service>vprn

Description This command sets the identifier that gets attached to routes to which the VPN belongs. 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.

AS numbers can be either 2-byte or 4-byte values.

Default no route-distinguisher

Parameters *rd* — the route distinguisher value

Values *ip-addr*:*comm-val* | *2byte-asnumber*:*ext-comm-val* | *4-byte-asnumber*:*comm-val*

where

ip-addr: a.b.c.d

comm-val: 0 to 65535

2-byte-asnumber: 1 to 65535

ext-comm-val: 0 to 4294967295

4-byte-asnumber: 1 to 4294967295

router-id

Syntax **router-id** *ip-address*
no router-id

Context config>service>vprn
config>service>vprn>bgp

Description 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, in dotted-decimal notation

Values a.b.c.d

service-name

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

Context config>service>vprn

Description This command configures a service name that can be used for reference in configuration and show commands.

Parameters *service-name* — up to 64 characters

sgt-qos

Syntax **sgt-qos**

Context config>service>vprn

Description This command enables the context to configure DSCP/dot1p re-marking for self-generated traffic.

application

Syntax **application** *dscp-app-name* **dscp** {*dscp-value* | *dscp-name*}
application *dot1p-app-name* **dot1p** *dot1p-priority*
no application {*dscp-app-name* | *dot1p-app-name*}

Context config>service>vprn>sgt-qos

Description This command configures DSCP or dot1p re-marking for self-generated traffic. When an application is configured using this command, then the specified DSCP name or DSCP value is used for all packets generated by this application within the router instance that it is configured.

Using the value configured in this command has the following effects:

- sets the DSCP bits in the IP packet
- maps the value to the FC. This value will be signaled from the CSM to the egress forwarding complex.
- the dot1p and the LSP EXP bits are set by the egress complex for all packets based on the signaled FC. This includes ARP and IS-IS packets that, due to their nature, do not carry DSCP bits.
- the DSCP value in the egress IP header will be as configured in this command. The egress QoS policy will not overwrite this value.

Only one DSCP name or DSCP value can be configured per application. If multiple entries are configured, then a subsequent entry overrides the previous entry.

The **no** form of this command reverts the DSCP value for the application back to its default value.

Default n/a (that is, sgt-qos does not enforce a DSCP value and the application uses its default value, as shown in Table 107)

Table 107: Applications and Default Values for DSCP or dot1p Markings

Application	Supported Marking	Default DSCP/dot1p	Application	Supported Marking	Default DSCP/dot1p
IS-IS	dot1p	7	SSH (SCP)	DSCP	AF41
ARP	dot1p	7	SNMP (get, set, etc.)	DSCP	AF41
BGP	DSCP	NC1	SNMP trap/ log	DSCP	AF41
DHCP	DSCP	NC1	syslog	DSCP	AF41
1588 PTP [1, 2]	DSCP	NC1	ICMP (ping)	DSCP	BE
LDP (T-LDP)	DSCP	NC1	Traceroute	DSCP	BE
RSVP	DSCP	NC1	TACACS+	DSCP	AF41
OSPF	DSCP	NC1	DNS	DSCP	AF41
Telnet	DSCP	AF41	NTP	DSCP	NC1
TFTP	DSCP	AF41	SNTP	DSCP	AF41
FTP	DSCP	AF41	RADIUS	DSCP	AF41

Notes:

1. PTP in the context of SGT QoS is defined as Precision Timing Protocol and is an application in the 7705 SAR. The PTP application name is also used in areas such as event-control and logging. Precision Timing Protocol is defined in IEEE 1588-2008.
2. PTP in the context of IP filters is defined as Performance Transparency Protocol. IP protocols can be used as IP filter match criteria; the match is made on the 8-bit protocol field in the IP header.

Parameters *dscp-app-name* — the DSCP application name

Values bgp, dhcp, dns, ftp, icmp, ldp, ntp, ospf, ptp, radius, rsvp, snmp, snmp-notification, ssh, syslog, tacplus, telnet, tftp, traceroute

dscp-value — the DSCP value when this packet egresses. The respective egress policy should provide the mapping for the DSCP value to either LSP-EXP bits or IEEE 802.1p (dot1p) bits as appropriate; otherwise, the default mapping applies.

Values 0 to 63

dscp-name — the DSCP name

Values none, be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

dot1p-app-name — the dot1p application name

Values arp, isis

dot1p-priority — the dot1p priority

Values none, 0 to 7

dscp

Syntax **dscp** *dscp-name* **fc** *fc-name*
no dscp *dscp-name*

Context config>service>vprn>sgt-qos

Description This command creates a mapping between the DSCP of the self-generated traffic and the forwarding class.

Self-generated traffic that matches the specified DSCP will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all of the 64 DSCP values to the forwarding class. For undefined code points, packets are assigned to the forwarding class specified under the **default-action** command.

All DSCP names that define a DSCP value must be explicitly defined.

The **no** form of this command removes the DSCP-to-forwarding class association. The default action then applies to that code point value.

Default n/a

Parameters *dscp-name* — the DSCP name to be associated with the forwarding class. DSCP can only be specified by its name and only an existing value can be specified. The software provides names for the well-known code points.

Values be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

fc-name — the forwarding class name. All packets with a DSCP value or MPLS EXP bits that are not defined will be placed in this forwarding class.

Values be, l2, af, l1, h2, ef, h1, nc

Default n/a (the fc name must be specified)

snmp-community

Syntax **snmp-community** *community-name* [**hash** | **hash2**] [**version** *SNMP-version*]
no snmp-community *community-name* [**hash** | **hash2**]

Context config>service>vprn

Description This command sets the SNMP community name to be used with the associated VPRN instance.

If an SNMP community name is not specified, SNMP access is not allowed.

The **no** form of the command removes the SNMP community name from the given VPRN context.

Default n/a

Parameters *community-name* — one or more SNMP community names

Values community-name: 32 characters (max)
hash-key: 33 characters (max)
hash2-key: 96 characters (max)

hash, hash2 — the hashing scheme for the community name

SNMP-version **—** the SNMP version

Values v1, v2c, both

source-address

Syntax **source-address**

Context config>service>vprn

Description This command enters the context to specify the source address and application that should be used in all unsolicited packets.

application

Syntax **application app** {[ping | ptp | ssh | telnet | traceroute]} | {[*ip-int-name* | *ip-address*]}
no application app {[ping | ptp | ssh | telnet | traceroute]}

Context config>service>vprn>source-address

Description This command configures the application to use the IPv4 source address.

The no form of the command removes the application name from using the IPv4 source address.

Parameters **app —** the application name

Values ping, ptp, ssh, telnet, traceroute

ip-int-name | *ip-address* — the name of the IPv4 interface or IPv4 address. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Values *ip-int-name*: up to 32 alphanumeric characters
ip-address: a.b.c.d

application6

Syntax **application6 app** {[ping | telnet | ssh | traceroute]} | *ipv6-address*
no application6 app {[ping | telnet | ssh | traceroute]}

Context config>service>vprn>source-address

Description This command configures the application to use the IPv6 source address.

The **no** form of the command removes the application name from using the IPv6 source address.

Parameters **app —** the application name

Values ping, telnet, ssh, traceroute

ipv6-address — the IPv6 address

Values

ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x - [0..FFFF]H d - [0..255] D

spoke-sdp

Syntax [**no**] **spoke-sdp** *sdp-id*

Context config>service>vprn

Description 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 n/a

Special Cases **VPRN —** several SDPs can be bound to a VPRN service. Each SDP must be destined for a different 7705 SAR, 7750 SR, or 7710 SR router. If two *sdp-id* bindings terminate on the same 7705 SAR, an error occurs and the second SDP binding is rejected.

Parameters *sdp-id* — the SDP identifier

Values 1 to 17407

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**] [**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 This command creates static route entries within the associated router instance. When configuring an IPv4 static route, either **next-hop**, **ipsec-tunnel**, or **black-hole** must be configured. For an IPv6 static route, **ipsec-tunnel** must be used.

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 enter as many parameters as necessary to uniquely identify that static route.

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 match, 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**, then the **cpe-check** test will be removed from the associated static route.

Default no static-route

Parameters *ip-prefix/prefix-length* — the destination address of the static route

Values
ip-prefix (IPv4 static route) a.b.c.d (host bits must be 0)
prefix-length (IPv4 static route) 0 to 30
ip-prefix (IPv6 static route) x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255] D
prefix-length (IPv6 static route) 0 to 126

netmask — the subnet mask in dotted-decimal notation

Values a.b.c.d (network bits all 1 and host bits all 0)

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, the metric will not change unless specified.

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 ecmp command.

Values 1 to 255

Default 5

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 or IS-IS. This value is also used to determine the static route to install in the forwarding table: When modifying the *metric* values of an existing static route, the preference will not change unless specified.

If there are multiple static routes with the same preference but unequal *metric* values, the lower-cost (metric) route is installed. If there are multiple static routes with equal preference and metrics, then the 7705 SAR OS chooses the route with the lowest next-hop IP address as the best route. If there are multiple routes with unequal preferences, then the lower preference route is installed.

Values 0 to 65535

Default 1

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 to 4294967295

enable | disable — static routes can be administratively enabled or disabled. Use the **enable** parameter to re-enable a disabled static route. Use the **disable** parameter to disable a static route while maintaining the static route in the configuration. In order to disable or 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

ip-int-name — the interface name of the directly connected next hop. Interface names must be unique within the group of defined IP interfaces for **config router 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 within double quotes.

Values 1 to 32 characters (must start with a letter)

ip-address **—** the directly connected next-hop IP address used to reach the destination

The **next-hop**, **ipsec-tunnel**, and **black-hole** keywords are mutually exclusive. If a static route is configured with the same destination address, subnet mask, and next-hop IP address as a previously configured static route, the newly configured route replaces the previous one, and unless specified, the respective defaults for **preference** and **metric** will be applied.

The **next-hop** *ip-address* 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.

Values *ip-int-name:* 32 chars max
ip-address (IPv4 static route): a.b.c.d
ip-address (IPv6 static route): x:x:x:x:x:x:x:x [-interface]
x:x:x:x:x:x:d.d.d.d [-interface]
interface - 32 chars max,
mandatory for link
local addresses
x - [0..FFFF]H
d - [0..255] D

ipsec-tunnel-name — the name of the IPSec tunnel that specifies the local and peer gateway addresses for the tunnel

The **ipsec-tunnel**, **black-hole**, and **next-hop** keywords are mutually exclusive. If a static route is configured with the same destination address, subnet mask, and IPSec tunnel name as a previously configured static route, the newly configured route replaces the previous one, and unless specified, the respective defaults for **preference** and **metric** will be applied.

Values 32 chars max

bfd-enable — associates the state of the static route to a BFD session between the local system and the configured next hop. This keyword is only available when the **next-hop** keyword is used.

cpe-ip-address — 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 *cpe-ip-address* cannot be in the same subnet as the static route subnet itself to avoid possible circular references. This option and BFD support are mutually exclusive on a given static route.

Values a.b.c.d

Default no cpe-check enabled

seconds — the interval between ICMP pings to the target IP address

Values 1 to 255 seconds

Default 1 second

count **—** the number of consecutive ping-replies that must be missed in order to declare the CPE down and to deactivate the associated static route

Values 1 to 255

Default 3

log — enables the logging of transitions between active and inactive routes based on the CPE connectivity check. Events should be sent to the system log, syslog and SNMP traps.

black-hole — specifies a blackhole route, meaning that if the destination address on a packet matches this static route, it will be silently discarded.

The **black-hole**, **ipsec-tunnel**, and **next-hop** keywords are mutually exclusive. If a static route is configured with the same destination address and subnet mask as a previously configured static route, the newly configured route replaces the previous one, and unless specified, the respective defaults for **preference** and **metric** will be applied.

type

Syntax **type hub**
no type

Context config>service>vprn

Description This command designates the type of VPRN instance being configured for hub and spoke topologies.

The **no** form of the command resets to the default of a fully meshed VPRN.

Default no type

Parameters **hub —** a hub VPRN, which allows all traffic from the hub SAP to be routed directly to the destination, while all traffic from spoke VPRNs or network interfaces can only be routed to a hub SAP

vrf-export

Syntax **vrf-export** *policy-name* [*policy-name...*(up to 5 max)]
no vrf-export

Context config>service>vprn

Description This command specifies the export policies to control routes exported from the local VPN virtual routing/ forwarding table (VRF) to other VRFs on the same or remote PE routers (via MP-BGP). The policy (and *policy-name*) are defined under the **config>router>policy-options>policy-statement** command.

The **no** form of the command removes all route policy names from the export list.

Default n/a

Parameters *policy-name* — the route policy statement name (up to 32 characters)

vrf-import

Syntax **vrf-import** *policy-name* [*policy-name...*(up to 5 max)]
no vrf-import

Context config>service>vprn

Description This command sets the import policies to control routes imported to the local VPN virtual routing/ forwarding table (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 n/a

Parameters *policy-name* — the route policy statement name (up to 32 characters)

vrf-target

Syntax **vrf-target** {*ext-community* | {[**export** *ext-community*] [**import** *ext-community*]}}
no vrf-target

Context config>service>vprn

Description This command facilitates a simplified method to configure the route target to be added to advertised routes or compared against received routes from other VRFs on the same or remote PE routers (via MP-BGP).

BGP-VPN routes imported with a **vrf-target** statement 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 in the same router.

Specified **vrf-import** or **vrf-export** policies override the **vrf-target** policy.

The **no** form of the command removes the route target from the VRF.

Default no vrf-target

Parameters *ext-community* — an extended BGP community in the *type***:***x***:***y* format.

Values *ip-addr*:*comm-val* | *2byte-asnumber*:*ext-comm-val* | *4-byte-asnumber*: *comm-val*

where

ip-addr: a.b.c.d

comm-val: 0 to 65535

2-byte-asnumber: 0 to 65535

ext-comm-val: 0 to 4294967295

4-byte-asnumber: 0 to 4294967295

export *ext-community* **—** communities allowed to be sent to remote PE neighbors

import *ext-community* **—** communities allowed to be accepted from remote PE neighbors

BGP Commands

bgp

Syntax [**no**] **bgp**

Context config>service>vprn

Description This command enables the BGP protocol on the VPRN service.

The **no** form of this command disables the BGP protocol on the 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 This command enables the advertising of inactive BGP routes 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.

The **no** form of this command disables the advertising of inactive BGP routes to other BGP peers.

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 This command is used to set the router ID in the BGP aggregator path attribute to 0 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 only the router ID (set to 0) to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level, and 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 the 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 global 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

backup-path

Syntax [**no**] **backup-path** [**ipv4**]

Context config>service>vprn>bgp

Description This command enables BGP Fast Reroute (FRR) with Prefix-Independent Convergence (PIC), allowing for the creation of a backup path for IPv4 BGP learned prefixes belonging to a VPRN. Multiple paths must be received for a prefix in order to take advantage of this feature.

When a prefix has a backup path, and its primary paths fail, the affected traffic is rapidly diverted to the backup path without waiting for control plane reconvergence to occur. The time to reroute the traffic is independent of the number of prefixes sharing the primary or backup paths.

Default no backup-path

Parameters **ipv4 —** enables a backup path for IPv4 BGP learned prefixes

best-path-selection

Syntax **best-path-selection**

Context config>service>vprn>bgp

Description This command enables path selection configuration.

always-compare-med

Syntax **always-compare-med** {**zero** | **infinity**}
no always-compare-med

Context config>service>vprn>bgp>path-selection

Description This command specifies how the Multi-Exit Discriminator (MED) path attribute is used in the BGP route selection process. If this command is set to zero or infinity, the MED attribute is always used in the route selection process regardless of the peer AS that advertised the route.

This command determines what MED value is inserted in the RIB-IN.

The **no** form of the command means that only the MEDs of routes that have the same peer ASs are compared.

Default no always-compare-med

Parameters **zero —** specifies that for routes learned without a MED attribute, a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.

infinity — specifies that for routes learned without a MED attribute, 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** [**ipv4**] [**ipv6**]

Context config>service>vprn>bgp>path-selection

Description This command determines whether the AS path is used to determine the best BGP route.

If this command is enabled, the AS paths of incoming routes are not used in the route selection process.

When **as-path-ignore** is used without specifying one or more keywords, then all keywords are configured. When one or more keywords are specified, then only those keywords are configured.

The **no** form of the command means that the AS paths of incoming routes are used to determine the best BGP route.

Default no as-path-ignore

Parameters **ipv4 —** specifies support for IPv4 routes

ipv6 — specifies support for IPv6 routes

as-override

Syntax [**no**] **as-override**

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description 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 the BGP loop detection mechanism. It should be used carefully.

Default no as-override

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 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 MD5 message-based digest.

The authentication key can be any combination of ASCII characters up to 255 characters long.

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, the entire string must be enclosed 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, the entire string must be enclosed 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 that 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 that the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command enables the use of bidirectional 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 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 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), 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 s

Parameters *seconds* — the BGP connect retry timer value, in seconds, expressed as a decimal integer

Values 1 to 65535

damping

Syntax [**no**] **damping**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command enables BGP damping for learned routes that are defined within the VPRN service. Damping parameters are set at the route policy level. See 7705 SAR OS Router Configuration Guide, 'Route Policy Command Reference".

The **no** form of the command disables learned 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.

Default no damping

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 This command configures BGP to disable sending communities.

Default no disable-communities

Parameters **standard —** specifies standard communities that existed before VPRNs or RFC 2547

extended — specifies BGP communities that were expanded after the concept of RFC 2547 was introduced, to include handling the route target in the VRF

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 This command configures BGP fast external failover.

For EBGP neighbors, fast external failover controls whether the router should drop an EBGP session immediately upon an interface-down event, or whether the BGP session is kept up until the hold-time expires.

When fast external failover is disabled, the EBGP session stays up until the hold-time expires or the interface comes back up again. If the BGP routes become unreachable as a result of the interface going down, they are immediately withdrawn from other peers.

Default no disable-fast-external-failover

enable-bgp-vpn-backup

Syntax [**no**] **enable-bgp-vpn-backup** [**ipv4**]

Context config>service>vprn>bgp

Description This command allows BGP-VPN routes imported into the VPRN to be used as backup paths for IPv4 BGP learned prefixes.

Parameters **ipv4 —** allow BGP-VPN routes to be used as backup paths for IPv4 prefixes

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 This command enables BGP peer tracking. BGP peer tracking allows a BGP peer to be dropped immediately if the route used to resolve the BGP peer address is removed from the IP routing table and there is no alternative available. The BGP peer will not wait for the hold timer to expire; therefore, the BGP reconvergence process is accelerated.

The **no** form of the command disables peer tracking.

Default no enable-peer-tracking

export

Syntax **export** *policy-name* [*policy-name*…(up to 5 max)]
no export

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command specifies the export policies used to control routes advertised to BGP neighbors. Route policies are configured in the **config>router>policy-options** context. Refer to the section on "Route Policy" in the 7705 SAR OS Router Configuration Guide.

When multiple policy names are specified, the policies are evaluated in the order in which they are specified. A maximum of five (5) policy names can be configured. The first policy that matches is applied.

If a non-existent route policy is applied to a VPRN instance, the CLI generates a warning message. This message is only generated during an interactive CLI session. 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 routes are advertised and non-BGP routes are not advertised

Parameters *policy-name* — the 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.

family

Syntax **family** [**ipv4**] [**ipv6**]
no family

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command specifies the address families to be negotiated with one or more multiprotocol BGP peers of the VPRN.

The **no** form of this command removes the specified address family from the associated BGP sessions.

Default ipv4

Parameters **ipv4** — provisions IPv4 support

ipv6 — provisions IPv6 support

graceful-restart

Syntax [**no**] **graceful-restart**

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command enables graceful restart for BGP in the VPRN context. If the control plane of a GR-capable router fails, the VPRN BGP peers (GR helpers) temporarily preserve neighbor information, so packets continue to be forwarded through the failed GR router using the last known routes. The helper state remains until the peer completes its restart or exits if the GR timer value is exceeded.

The 7705 SAR acts as a GR helper; it does not request graceful restart but agrees to graceful restart requests from a peer.

The **no** form of the command disables graceful restart and removes all graceful restart configurations in the VPRN BGP instance.

Default no graceful-restart

stale-routes-time

Syntax **stale-routes-time** *time*
no stale-routes-time

Context config>service>vprn>bgp>graceful-restart
config>service>vprn>bgp>group>graceful-restart
config>service>vprn>bgp>group>neighbor>graceful-restart

Description This command configures the maximum amount of time in seconds that stale routes should be maintained after a graceful restart is initiated.

The **no** form of the command resets the stale routes time back to the default value.

Default 360 s

Parameters *time* — the amount of time that stale routes should be maintained after a graceful restart is initiated

Values 1 to 3600 s

group

Syntax [**no**] **group** *name* [**dynamic-peer**]

Context config>service>vprn>bgp

Description 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 shut down before it can be deleted.

Default no group—no peer groups are defined

Parameters *name* — the peer group 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.

dynamic-peer — this option specifies 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 only be one BGP group with this option set in any given VPRN. No BGP neighbors can be manually configured in a BGP group with this option set.

Default disabled

neighbor

Syntax [**no**] **neighbor** *ip-address*

Context config>service>vprn>bgp>group

Description 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 configurations.

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 shut down before it can be deleted. If the neighbor is not shut down, the command will not result in any action except a warning message on the CLI indicating that the neighbor is still administratively up.

Default no neighbor—no neighbors are defined

Parameters *ip-address* — the IP address of the BGP peer router

Values *ip-address (IPv4 address):* a.b.c.d
ip-address (IPv6 address): x:x:x:x:x:x:x:x [-interface]
x:x:x:x:x:x:d.d.d.d [-interface]
interface - 32 chars max,
mandatory for link
local addresses
x - [0..FFFF]H
d - [0..255]D

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 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.

The **strict** option ensures that the negotiated hold time value is not set to a value less than the configured value.

Even though the 7705 SAR OS implementation allows setting the **keepalive** time separately, the configured **keepalive** timer is overridden by the **hold-time** value under the following circumstances.

- 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.
- If the **hold-time** is set to 0, then the operational value of the **keepalive** time is set to 0; 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 s

Parameters *seconds* — the hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is permanently up.

Values 0, 3 to 65535

strict — when used, the advertised BGP hold time from the far-end BGP peer must be greater than or equal to the specified hold-time value

import

Syntax **import** *policy-name* [*policy-name*…(up to 5 max)]
no import

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command specifies the import route policy to be used to determine which routes are accepted from peers. Route policies are configured in the **config>router>policy-options** context. Refer to the section on "Route Policy" in the 7705 SAR OS Router Configuration Guide.

When multiple policy names are specified, the policies are evaluated in the order in which they are specified. A maximum of five (5) policy names can be specified. The first policy that matches is applied.

When multiple **import** commands are issued, the last command entered will override the previous command.

The **no** form of the command removes all route policy names from the import list.

Default no import—BGP routes are accepted by default

Parameters *policy-name* — the 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.

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the BGP keepalive timer. A Keepalive message is sent every time this timer expires.

The **keepalive** 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. The **keepalive** value is generally one-third of the **hold-time** interval. Even though the 7705 SAR 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 value** 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 0, 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 s

Parameters *seconds* — the keepalive timer, in seconds, expressed as a decimal integer

Values 0 to 21845

local-address

Syntax **local-address** *ip-address*
no local-address

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command 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 7705 SAR OS 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

Parameters *ip-address* — the local address. The allowed value is a valid routable IP address on the router, either an interface or system IP address.

Values *ip-address (IPv4 address):* a.b.c.d

ip-address (IPv6 address): x:x:x:x:x:x:x:x [-interface]
x:x:x:x:x:x:d.d.d.d [-interface]
interface - 32 chars max,
mandatory for link
local addresses
x - [0..FFFF]H
d - [0..255]D

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 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 attribute 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 group) or neighbor level (only applies to specified peer). By specifying this parameter 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 group 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 example:

Example: 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 to 4294967295

private — specifies that 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 This command configures the default value of the BGP local preference attribute if it is not already specified in incoming routes.

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 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 a 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

Parameters *local-preference* — the local preference value to be used as the override value, expressed as a decimal integer

Values 0 to 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 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 group) or neighbor level (only applies to specified peer). The most specific value is used.

When applied to an ongoing BGP peer session, this command does not take effect until the BGP peer session is re-established.

The **no** form of the command used at the global level reverts to the default (**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 ignore-loop

Parameters **drop-peer** — sends a notification to the remote peer and drops the session

discard-route — discards routes received from a peer with the same AS number as the router itself. This option prevents routes looped back to the router from being added to the routing information base and consuming memory. When this option is changed, the change will not be active for an established peer until the connection is re-established for the peer.

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 This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the advertised MED 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 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 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 to 4294967295

igp-cost — the MED is set to the IGP cost of the IP prefix that is defined via a route policy

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 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 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.

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 s

Parameters *seconds* — the minimum path attribute advertising interval, in seconds, expressed as a decimal integer

Values 2 to 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 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 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.

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 s

Parameters *seconds* — the minimum route advertising interval, in seconds, expressed as a decimal integer

Values 1 to 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 This command configures the time to live (TTL) value at an originating EBGP peer. The TTL value is entered in the IP header of packets that are sent to a terminating EBGP peer that is multiple hops away.

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 1 — EBGP peers are directly connected

Parameters *ttl-value* — the TTL value that will be entered in the IP header of packets that are sent to a terminating EBGP peer that is multiple hops away

Values 1 to 255

multipath

Syntax **multipath** *max-paths*
no multipath

Context config>service>vprn>bgp

Description This command enables BGP multipath.

When multipath is enabled, BGP load-shares traffic across multiple links. Multipath can be configured to load-share traffic across a maximum of 16 routes. If the equal-cost routes available are more than the configured value, then routes with the lowest next-hop IP address value are chosen.

This configuration parameter is set at the global level (applies to all peers).

Multipath is disabled if the value is set to 1. When multipath is disabled and multiple equal-cost routes are available, the route with the lowest next-hop IP address will be used.

The **no** form of the command reverts to the default where **multipath** is disabled.

Default no multipath

Parameters *max-paths* — the number of equal-cost routes to use for multipath routing

Values 1 to 16

next-hop-self

Syntax [**no**] **next-hop-self**

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the group or neighbor to always set the next-hop path attribute to its own physical interface when advertising to a peer.

This command is primarily used to avoid third-party route advertisements when connected to a multi-access network.

The **no** form of the command used at the group level allows third-party route advertisements in a multi-access network.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default no next-hop-self

passive

Syntax [**no**] **passive**

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command enables and disables 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, and 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

peer-as

Syntax **peer-as** *as-number*
no peer-as

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description 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. This requirement is necessary since the peer will be in a different autonomous system than that of this router.

This command may be configured under the group level for all neighbors in a particular group.

Default no AS numbers defined

Parameters *as-number* — the autonomous system number, expressed as a decimal integer

Values 1 to 4294967295

preference

Syntax **preference** *preference*
no preference

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description 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 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 7705 SAR OS assigns the highest default preference to BGP routes as 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 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 170

Parameters *preference* — the route preference, expressed as a decimal integer

Values 1 to 255

prefix-limit

Syntax **prefix-limit** *limit*
no prefix-limit

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures the maximum number of routes that BGP can learn from a peer.

When the number of routes reaches 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 to 4294967295

rapid-withdrawal

Syntax [**no**] **rapid-withdrawal**

Context config>service>vprn>bgp

Description This command disables the delay on issuing BGP withdrawals.

By default, BGP withdrawals (messages containing the routes that are no longer valid) are delayed up to the **min-route-advertisement** to allow for efficient packing of BGP Update messages. However, when the **rapid-withdrawal** command is enabled, the delay on sending BGP withdrawals is disabled.

The **no** form of the command returns BGP withdrawal processing to its default behavior.

Default no rapid-withdrawal

remove-private

Syntax [**no**] **remove-private** [**limited**]

Context config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command allows all private AS numbers to be removed from the AS path before advertising them to BGP peers. The **no** form of the command includes private AS numbers in the AS path attribute.

If the **limited** keyword is included, only the leading private ASNs up to the first public ASN are removed.

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 7705 SAR OS 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 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 no remove-private

ttl-security

Syntax **ttl-security** *min-ttl-value*
no ttl-security

Context config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor

Description This command configures TTL security parameters for incoming packets. When the feature is enabled, BGP accepts incoming IP packets from a peer only if the TTL value in the packet is greater than or equal to the minimum TTL value configured for that peer.

The **no** form of the command disables TTL security.

Default no ttl-security

Parameters *min-ttl-value* — the minimum TTL value for an incoming packet

Values 1 to 255

Default 1

OSPF Commands

ospf

Syntax [**no**] **ospf**

Context config>service>vprn

Description This command enables access to the context to define OSPF parameters for VPRN.

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 and removes all associated configuration parameters.

Default no ospf

area

Syntax [**no**] **area** *area-id*

Context config>service>vprn>ospf

Description 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, address ranges, and so on, that are currently assigned to this area.

The 7705 SAR supports a maximum of four areas.

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 to 255.255.255.255 (dotted-decimal)
0 to 4294967295 (decimal integer)

area-range

Syntax **area-range** *ip-prefix/mask* [**advertise** | **not-advertise**]
no area-range *ip-prefix/mask*

Context config>service>vprn>ospf>area
config>service>vprn>ospf>area>nssa

Description 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 to other areas. Multiple range commands can be used to summarize or hide 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 advertisement or non-advertisement.

Default no area-range — no range of addresses is 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 *ip-prefix* — the IP prefix in dotted-decimal notation for the range

Values a.b.c.d (host bits must be 0)

mask — the subnet mask for the range, expressed as a decimal integer

Values 0 to 32

advertise | **not-advertise —** specifies whether to advertise the summarized range of addresses to other areas

Default advertise

blackhole-aggregate

Syntax [**no**] **blackhole-aggregate**

Context config>service>vprn>ospf>area

Description This command installs a low-priority blackhole route for the entire aggregate. Existing routes that make up the aggregate will have a higher priority and only the components of the range for which no route exists will be blackholed.

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*

Context config>service>vprn>ospf>area

Description 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 can be used to disable an interface without removing the configuration.

Default no interface

Parameters *ip-int-name* — the IP interface name. Interface names must be unique within the group of defined IP interfaces for the **config>service>vprn>interface** and **config>router>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.

advertise-subnet

Syntax [**no**] **advertise-subnet**

Context config>service>vprn>ospf>area>interface

Description 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.

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

authentication-key

Syntax **authentication-key** {*authentication-key* | *hash-key*} [**hash** | **hash2**]
no authentication-key

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link

Description 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.

All neighboring routers must use the same type of authentication and password for correct protocol communication. If the **authentication-type** is configured as password, the authentication 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

Parameters *authentication-key* — the authentication 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 can be any combination of ASCII characters up to 22 characters in length (**hash** parameter is used) or 121 characters in length (if the **hash2** parameter is used). 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 that 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 that 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>interface
config>service>vprn>ospf>area>virtual-link

Description This command enables authentication and specifies the type of authentication to be used on the OSPF interface or virtual link.

Both simple **password** and **message-digest** authentication are supported.

By default, authentication is not enabled on an interface or link.

The **no** form of the command disables authentication on the interface or link.

Default no authentication-type

Parameters **password —** 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 — enables message digest MD5 authentication in accordance with RFC 1321. If this option is configured, at least one message-digest-key must be configured.

dead-interval

Syntax **dead-interval** *seconds*
no dead-interval

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link

Description This command configures the time, in seconds, that OSPF waits before declaring a neighbor router or virtual-link neighbor down. If no Hello packets are received from a neighbor for the duration of the dead interval, the router or link 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, all nodes on the subnet must have the same dead interval

Virtual Link — if the **dead-interval** configured applies to a virtual link, the interval on both endpoints of the virtual link must have the same dead interval

Parameters *seconds* — the dead interval in seconds, expressed as a decimal integer

Values 1 to 65535

hello-interval

Syntax **hello-interval** *seconds*
no hello-interval

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link

Description This command configures the interval between OSPF hellos issued on the interface or virtual 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 but results in higher processing costs.

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

Default 10

Special Cases **OSPF Interface —** if the **hello-interval** configured applies to an interface, all nodes on the subnet must have the same hello interval

Virtual Link — if the **hello-interval** configured applies to a virtual link, the interval on both endpoints of the virtual link must have the same hello interval

Parameters *seconds* — the hello interval in seconds, expressed as a decimal integer

Values 1 to 65535

interface-type

Syntax **interface-type** {**broadcast** | **point-to-point**}
no interface-type

Context config>service>vprn>ospf>area>interface

Description 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 of the link, provided that the link is used as a point-to-point link.

If the interface type is not known when the interface is added to OSPF, and the IP interface is subsequently 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 broadcast – if the physical interface is Ethernet or unknown

point-to-point – if the physical interface is T1, E1, or SONET/SDH

Special Cases **Virtual Link —** a virtual link is always regarded as a point-to-point interface and is 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

lfa-policy-map

Syntax **lfa-policy-map route-nh-template** *template-name*
no lfa-policy-map

Context config>service>vprn>ospf>area>interface

Description This command applies a route next-hop policy template to an OSPF interface. When a route next-hop policy template is applied to an interface in OSPF, it is applied in all areas.

If the interface has been excluded from LFA with the loopfree-alternate-exclude command, the LFA policy has no effect on the interface.

If the route next-hop policy template is applied to a loopback interface or to the system interface, the command will not be rejected, but the policy will have no effect on the interface.

The **no** form of the command deletes the mapping of a route next-hop policy template to an OSPF interface.

Default no lfa-policy-map

Parameters *template-name* — the name of an existing template

message-digest-key

Syntax **message-digest-key** *key-id* **md5** {*key* | *hash-key* | *hash2-key*} [**hash** | **hash2**]
no message-digest-key *key-id*

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link

Description This command configures a message digest key when MD5 authentication is enabled on the interface or virtual link. Multiple message digest keys can be configured.

The **no** form of the command removes the message digest key identified by the *key-id*.

Default no message-digest-key

Parameters *key-id* — the *key-id* is expressed as a decimal integer

Values 1 to 255

key — the MD5 key, any alphanumeric string up to 16 characters in length

hash-key — the MD5 hash key, any combination of ASCII characters up to 33 characters in length (**hash** parameter is used). 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.

hash2-key — the MD5 hash key, any combination of ASCII characters up to 132 characters in length (**hash2** parameter is used). 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 that the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — specifies that the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted hash form is assumed.

metric

Syntax **metric** *metric*
no metric

Context config>service>vprn>ospf>area>interface

Description 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

Parameters *metric* — the metric to be applied to the interface, expressed as a decimal integer

Values 1 to 65535

mtu

Syntax **mtu** *bytes*
no mtu

Context config>service>vprn>ospf>area>interface

Description This command configures the MTU value used when negotiating an OSPF adjacency. 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>tdm>t1-e1>channel-group**

If this parameter is configured, the lesser of the values between the value configured with this command and the MTU configured (default or explicitly) in the preceding listed contexts is used.

To determine the actual packet size, add 14 bytes for a null Ethernet packet, 18 bytes for a dot1q Ethernet packet, and 22 bytes for a qinq Ethernet packet, to the size of the OSPF (IP) packet MTU configured with this command.

Use the **no** form of this command to revert to the default.

Note: This command is only used for negotiating adjacencies with OSPF neighbors. After the adjacency is established, the maximum OSPF packet size is always based on the values derived from the MTU configured in the contexts listed above.

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 to 9710 (9724 – 14) (depends on the physical media)

passive

Syntax **[no] passive**

Context config>service>vprn>ospf>area>interface

Description 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 with the **config**>**router**>**service-prefix** command are passive. All other interfaces are not passive.

priority

Syntax **priority** *number*
no priority

Context config>service>vprn>ospf>area>interface

Description This command configures the priority of the OSPF interface that is used in an election of the designated router 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 a designated router or backup designated router.

The **no** form of the command reverts the interface priority to the default value.

Default 1

Parameters *number* — the interface priority expressed as a decimal integer

Values 0 to 255

retransmit-interval

Syntax **retransmit-interval** *seconds*
no retransmit-interval

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link

Description 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. If the retransmit interval expires and no acknowledgment has been received, the LSA will be retransmitted.

The **no** form of this command reverts to the default interval.

Default 5

Parameters *seconds* — the retransmit interval in seconds, expressed as a decimal integer

Values 1 to 1800

transit-delay

Syntax **transit-delay** *seconds*
no transit-delay

Context config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>virtual-link

Description This command configures the estimated time, in seconds, that it takes to transmit a link-state advertisement (LSA) on the interface or virtual link.

The **no** form of this command reverts to the default delay time.

Default 1

Parameters *seconds* — the transit delay in seconds, expressed as a decimal integer

Values 1 to 1800

loopfree-alternate-exclude

Syntax [**no**] **loopfree-alternate-exclude**

Context config>service>vprn>ospf>area
config>service>vprn>ospf>area>interface

Description This command instructs OSPF to exclude a specific interface or all interfaces participating in a specific OSPF area from the LFA SPF calculation. The LFA SPF calculation can therefore be run only where it is needed.

If an interface is excluded from the LFA SPF in OSPF, it is excluded in all areas.

Default no loopfree-alternate-exclude

nssa

Syntax [**no**] **nssa**

Context config>service>vprn>ospf>area

Description This command creates the context to configure an OSPF Not So Stubby Area (NSSA) and adds or 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 that 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 area or NSSA are 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

originate-default-route

Syntax **originate-default-route** [**type-7**]
no originate-default-route

Context config>service>vprn>ospf>area>nssa

Description 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) or Autonomous System Border Router (ASBR).

Include the **type-7** 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 the **originate-default-route** command without the **type-7** parameter.

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

Parameters **type-7 —** specifies that a type 7 LSA should be used for the default route

Default type 3 LSA for the default route

redistribute-external

Syntax [**no**] **redistribute-external**

Context config>service>vprn>ospf>area>nssa

Description This command enables the redistribution of external routes into the Not So Stubby Area (NSSA) on an NSSA area border router (ABR) that is exporting the routes into non-NSSA areas.

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 that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an ABR 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

summaries

Syntax [**no**] **summaries**

Context config>service>vprn>ospf>area>nssa
config>service>vprn>ospf>area>stub

Description This command enables sending summary (type 3) advertisements into a stub area or NSSA on an 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

stub

Syntax [**no**] **stub**

Context config>service>vprn>ospf>area

Description This command enables access to the context to configure an OSPF stub area and adds or 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.

Existing virtual links of a non-stub area or NSSA are removed when its designation is changed to NSSA or stub.

An OSPF area cannot be both an NSSA and a stub area at the same time.

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

default-metric

Syntax **default-metric** *metric*
no default-metric

Context config>service>vprn>ospf>area>stub

Description This command configures the metric used by the 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 to the stub area

Values 1 to 16777215

virtual-link

Syntax [**no**] **virtual-link** *router-id* **transit-area** *area-id*

Context config>service>vprn>ospf>area

Description This command configures a virtual link to connect ABRs to the backbone.

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 or possible to connect an area to the backbone, the ABRs must be connected via a virtual link. The two ABRs form a point-to-point-like adjacency across the transit area. 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 an NSSA.

The **no** form of the command deletes the virtual link.

Default no virtual-link

Parameters *router-id* — the router ID of the virtual neighbor in IP address dotted-decimal notation

area-id — the area ID specified identifies the transit area that links the backbone area to the area that has no physical connection with the backbone, expressed in dotted-decimal notation or as a 32-bit decimal integer

Values 0.0.0.0 to 255.255.255.255 (dotted-decimal)
0 to 4294967295 (decimal integer)

export

Syntax **export** *policy-name* [*policy-name*…(up to 5 max)]
no export

Context config>service>vprn>ospf

Description This command specifies export route policies to determine which routes are exported from the routing table manager to OSPF. Export policies 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.

Refer to the 7705 SAR OS Router Configuration Guide for information on defining route policies.

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 names must already be defined.

external-db-overflow

Syntax **external-db-overflow** *limit seconds*
no external-db-overflow

Context config>service>vprn>ospf

Description This command enables limits on the number of non-default, AS-external-LSA entries that can be stored in the link-state database (LSDB) and specifies a wait timer before processing these entries after the limit is exceeded.

The *limit* value specifies the maximum number of entries that can be stored in the LSDB. Placing a limit on these 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 reaches or exceeds the *limit*, the table is in an overflow state. When in an overflow state, the router will not originate any new AS-external LSAs and will withdraw all the self-originated non-default external LSAs.

The *seconds* value specifies the 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

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 to 2147483647

seconds — 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 to 2147483647

external-preference

Syntax **external-preference** *preference*
no external-preference

Context config>service>vprn>ospf

Description This command configures the preference for OSPF external routes. The preference for internal routes is set with the preference command.

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 based on the default preferences as defined in Table 108.

Table 108: Route Preference Defaults by Route Type

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes

If multiple routes are learned with the same preference using the same protocol, the lowest-cost route is used. If multiple routes are learned with the same preference using the same protocol and the costs (metrics) are equal, the decision of which route to use is determined by the configuration of ECMP in the **config>router** context. Refer to the 7705 SAR OS Router Configuration Guide for information on ECMP.

Note: To configure a preference for static routes, use the **config>router>static-route** command. Refer to the 7705 SAR OS Router Configuration Guide for information.

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

Values 1 to 255

ignore-dn-bit

Syntax [**no**] **ignore-dn-bit**

Context config>service>vprn>ospf

Description This command specifies whether to ignore the DN (down) 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.

Default no ignore-dn-bit

loopfree-alternate

Syntax [**no**] **loopfree-alternate**

Context config>service>vprn>ospf

Description This command enables Loop-Free Alternate (LFA) computation by SPF under the OSPF routing protocol instance.

When this command is enabled, the OSPF SPF attempts to precalculate both a primary next hop and a LFA backup next hop for every learned prefix. When found, the LFA next hop is populated into the routing table along with the primary next hop for the prefix.

The **no** form of this command disables the LFA SPF calculation.

Default no loop-free alternate

loopfree-alternate-exclude

Syntax **loopfree-alternate-exclude prefix-policy** *prefix-policy* [*prefix-policy...*(*up to 5 max*)]
no loop-free-alternate-exclude

Context config>service>vprn>ospf

Description This command excludes from the LFA SPF calculation any prefixes that match a prefix entry in a prefix policy. If a prefix is excluded, it is not included in the LFA SPF calculation, regardless of its priority. Prefix policies are created with the command **config>router> policy-options>prefix-list** (for information on prefix lists, refer to the 7705 SAR OS Router Configuration Guide, "Route Policies").

The default action of the **loopfree-alternate-exclude** command, when not explicitly specified in the prefix policy, is to "reject". Therefore, even if the **default-action reject** statement was not explicitly stated for the prefix policy, a prefix that does not match any entry in the policy will be used in the LFA SPF calculation.

The **no** form of the command deletes the excluded prefix policy.

Default no loop-free-alternate-exclude

Parameters *prefix-policy* — the name of an existing prefix policy

overload

Syntax **overload** [**timeout** *seconds*]
no overload

Context config>service>vprn>ospf

Description 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 for directly attached interfaces continues 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 no timeout is specified, the overload state is maintained indefinitely.

If the **overload** command is encountered during the execution of an overload-on-boot command, the **overload** command takes precedence. This situation 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 of the reason the protocol entered the overload state.

Default no overload

Parameters *seconds* — the number of seconds to reset overloading

Values 60 to 1800

overload-include-stub

Syntax [**no**] **overload-include-stub**

Context config>service>vprn>ospf

Description This command is used to determine if the OSPF stub networks should be advertised with a maximum metric value when the system goes into an 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-on-boot

Context config>service>vprn>ospf

Description When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures OSPF upon bootup in the overload state until one of the following events occurs:

- the timeout timer expires (if a timeout has been specified)
- a manual override of the current overload state is entered with the **no overload** command

If no timeout is specified, the overload state is maintained indefinitely.

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 *seconds* — the number of seconds to reset overloading

Values 60 to 1800

preference

Syntax **preference** *preference*
no preference

Context config>service>vprn>ospf

Description 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 which route will be used.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is based on the default preferences as defined in Table 108. If multiple routes are learned with the same preference using the same protocol and the costs (metrics) are equal, the decision of which route to use is determined by the configuration of ECMP in the **config**>**router** context. Refer to the 7705 SAR OS Router Configuration Guide for information on ECMP.

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

Values 1 to 255

reference-bandwidth

Syntax **reference-bandwidth** *bandwidth-in-kbps*
reference-bandwidth [**tbps** *Tera-bps*] [**gbps** *Giga-bps*] [**mbps** *Mega-bps*] [**kbps** *Kilo-bps*]
no reference-bandwidth

Context config>service>vprn>ospf

Description This command configures the reference bandwidth used to calculate the default costs 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 kb/s or 100 Gb/s; therefore, the default auto-cost metrics for various link speeds are as follows:

- 10 Mb/s link: default cost of 10000
- 100 Mb/s link: default cost of 1000
- 1 Gb/s link: default cost of 100

The **reference-bandwidth** command assigns a default cost to the interface based on the interface speed. To override this default cost on an interface, use the **metric** *metric* command in the **config**>**router**>**ospf**>**area**>**interface** *ip-int-name* context.

The **no** form of the command resets the reference bandwidth to the default value.

Default reference-bandwidth 100000000

Parameters *bandwidth-in-kbps* — the reference bandwidth in kilobits per second, expressed as a decimal integer

Values 1 to 400000000

Tera-bps — the reference bandwidth in terabits per second, expressed as a decimal integer

Values 1 to 4

Giga-bps — the reference bandwidth in gigabits per second, expressed as a decimal integer

Values 1 to 999

Mega-bps — the reference bandwidth in megabits per second, expressed as a decimal integer

Values 1 to 999

Kilo-bps — the reference bandwidth in kilobits per second, expressed as a decimal integer

Values 1 to 999

router-id

Syntax **router-id** *ip-address*
no router-id

Context config>service>vprn>ospf

Description This command configures the router ID for a specific VPRN context. If the router ID is not defined under VPRN, the router ID from the base router context is inherited.

When configuring the router ID in the base instance of OSPF, the value overrides the router ID configured in the **config**>**router** context.

The default value for the base instance is inherited from the configuration in the **config**>**router** context. If the router ID in the **config**>**router** context is not configured, the following applies:

- the system uses the system interface address (which is also the loopback address)
- if a system interface address is not configured, the last 4 bytes of the chassis MAC address are used

When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for the instance, or reboot the entire router.

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

Default 0.0.0.0 (base OSPF)

Parameters *ip-address* — a 32-bit, unsigned integer uniquely identifying the router in the Autonomous System

super-backbone

Syntax [**no**] **super-backbone**

Context config>service>vprn>ospf

Description This command specifies whether CE-PE functionality is required. 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

Description 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 will not be set. When disabled, this instance of the OSPF router will follow the usual procedure to determine whether to set the DN bit.

Default no suppress-dn-bit

timers

Syntax **timers**

Context config>service>vprn>ospf

Description 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 an SPF calculation and the minimum time between successive SPF calculations.

Changing the timers affects CPU usage and network reconvergence times. Lower values reduce reconvergence time but increase CPU usage. Higher values reduce CPU usage but increase reconvergence time.

Default n/a

lsa-arrival

Syntax **lsa-arrival** *lsa-arrival-time*
no lsa-arrival

Context config>service>vprn>ospf>timers

Description This command defines the minimum delay that must pass between receipt of the same link-state advertisements (LSAs) arriving from neighbors.

It is recommended that the configured **lsa-generate** *lsa-second-wait* interval for the neighbors be equal to or greater than the *lsa-arrival-time*.

Use the **no** form of this command to return to the default.

Default no lsa-arrival

Parameters *lsa-arrival-time* — the timer in milliseconds. Values entered that do not match this requirement will be rejected.

Values 0 to 600000

lsa-generate

Syntax **lsa-generate** *max-lsa-wait* [*lsa-initial-wait* [*lsa-second-wait*]]
no lsa-generate

Context config>service>vprn>ospf>timers

Description This command 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.

It is recommended that the *lsa-arrival-time* be equal to or less than the *lsa-second-wait* interval.

Use the **no** form of this command to return to the default.

Default no lsa-generate

Parameters *max-lsa-wait* — the maximum interval, in milliseconds, between two consecutive occurrences of an LSA being generated

Values 10 to 600000

Default 5000

lsa-initial-wait — the first waiting period between LSAs generated, in milliseconds. When the LSA exceeds the *lsa-initial-wait* timer value and the topology changes, there is no wait period and the LSA is immediately generated.

When an LSA is generated, the initial wait period commences. If, within the specified *lsa-initial-wait* period, another topology change occurs, the *lsa-initial-wait* timer applies.

Values 10 to 600000

Default 5000

lsa-second-wait — the hold time, in milliseconds, between the first and second LSA generation. The next topology change is subject to this second wait period. With each subsequent topology change, the wait time doubles (that is, two times the previous wait time). This assumes that each failure occurs within the relevant wait period.

Values 10 to 600000

Default 5000

spf-wait

Syntax **spf-wait** *max-spf-wait* [*spf-initial-wait* [*spf-second-wait*]]
no spf-wait

Context config>service>vprn>ospf>timers

Description This command defines the maximum interval between two consecutive SPF calculations in milliseconds. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, the next SPF will run after 2000 ms, and the next SPF will run after 4000 ms, and so on, 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 ms. 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* — the maximum interval, in milliseconds, between two consecutive SPF calculations

Values 10 to 120000

Default 1000

spf-initial-wait — the initial SPF calculation delay, in milliseconds, after a topology change

Values 10 to 100000

Default 1000

spf-second-wait — the hold time, in milliseconds, between the first and second SPF calculation

Values 10 to 100000

Default 1000

vpn-domain

Syntax **vpn-domain** *id* {**0005** | **0105** | **0205** | **8005**}
no vpn-domain

Context config>service>vprn>ospf

Description This command specifies the type of extended community attribute exchanged using BGP to carry the OSPF VPN domain ID. The command applies to VPRN instances of OSPF only. An attempt to modify the value of this attribute will result in an inconsistent value error when the instance is not a VPRN instance. The parameters are mandatory and can be entered in any order.

Default no vpn-domain

Parameters *id* — specifies the 6-octet OSPF VPN domain identifier in the format "xxxx.xxxx.xxxx". This ID is exchanged using BGP in the extended community attribute associated with a prefix. This parameter applies to VPRN instances of OSPF only.

0005 | **0105** | **0205** | **8005 —** specifies the type of extended community attribute exchanged using BGP to carry the OSPF VPN domain ID

vpn-tag

Syntax **vpn-tag** *vpn-tag*
no vpn-tag

Context config>service>vprn>ospf

Description This command specifies the route tag for an OSPF VPN on a PE router and is used mainly to prevent routing loops. This field is set in the tag field of the OSPF external LSAs generated by the PE. The command applies to VPRN instances of OSPF only. An attempt to modify the value of this tag will result in an inconsistent value error when the instance is not a VPRN instance.

Default vpn-tag 0

Parameters *vpn-tag* — specifies the route tag for an OSPF VPN

Values 0 to 4294967295

RIP Commands

rip

Syntax **[no] rip**

Context config>service>vprn

Description This command enables the RIP protocol on a VPRN interface.

The **no** form of the command disables the RIP protocol on a VPRN interface.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command sets the authentication password to be passed between RIP neighbors. The authentication type and authentication key must match exactly in order for the RIP message to be considered authentic.

The authentication key can be any combination of ASCII characters up to 16 characters long. The hash-key can be any combination of ASCII characters up to 33 characters long.

The **no** form of the command removes the authentication password from the configuration and disables authentication.

Default no authentication-key

Parameters *authentication-key* — the authentication key. The key can be any combination of ASCII characters up to 16 characters in length (unencrypted). If spaces are used in the string, the entire string must be enclosed in double quotes.

hash-key — the hash key. The key can be any combination of ASCII characters up to 33 characters in length (encrypted). If spaces are used in the string, the entire string must be enclosed in double quotes.

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 unencrypted, 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** {**none** | **password** | **message-digest-20**}

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command sets the type of authentication to be used between RIP neighbors. Authentication type can be specified regardless of the configured send and receive parameters, but will only apply to RIPv2 packets.

The type and password must match exactly for the RIP message to be considered authentic and processed.

The **no** form of the command removes the authentication type from the configuration and disables authentication.

Default no authentication-type

Parameters **none** — disables authentication

password — 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-20 — configures 16-byte message digest for MD5 authentication. If this option is configured, then at least one message-digest key must be configured.

check-zero

Syntax **check-zero** {**enable** | **disable**}
no check-zero

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command enables checking for zero values in fields specified to be zero by the RIPv1 and RIPv2 specifications.

The **check-zero enable** command enables checking of the mandatory zero fields in the RIPv1 and RIPv2 specifications and rejecting of non-compliant RIP messages.

The **check-zero disable** command disables this check and allows the receipt of RIP messages even if the mandatory zero fields are non-zero.

The **check-zero** command can be enabled at all three RIP levels. The most specific value is used. If no check-zero value is set (no check-zero), the setting from the less-specific level is inherited by the lower level.

The **no** form of the command disables check-zero on the configuration.

Default no check-zero

Parameters **enable —** configures the router to reject RIP messages that do not have zero in the mandatory fields

disable — configures the router to accept RIP messages that do not have zero in the mandatory fields

export

Syntax **export** *policy-name* [*policy-name*... (up to 5 max)]

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command specifies the export policies to be used to control routes advertised to RIP neighbors.

By default, when no export policies are specified, RIP routes are advertised and non-RIP routes are not advertised.

The **no** form of the command removes all route policy names from the export list.

Default no export

Parameters *policy-name* — the 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.

export-limit

Syntax **export-limit** *number* [**log** *percentage*]
no export-limit

Context config>service>vprn>rip

Description This command configures the maximum number of routes (prefixes) that can be exported into RIP from the route table.

The **no** form of the command removes the configured parameter values.

Default no export-limit

Parameters *number* — specifies the maximum number of routes (prefixes) that can be exported into RIP from the route table

Values 1 to 4294967295

percentage — specifies the percentage of the export-limit, that when reached, causes a warning log message and SNMP notification to be sent

Values 1 to 100

group

Syntax [**no**] **group** *group-name*

Context config>service>vprn>rip

Description This command creates a context for configuring a RIP group of neighbors.

RIP groups logically associate RIP neighbor interfaces to facilitate a common configuration for RIP interfaces.

The **no** form of the command deletes the RIP neighbor interface group. Deleting the group will also remove the RIP configuration of all the neighbor interfaces currently assigned to this group.

Default no group

Parameters *group-name* — the RIP group 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.

import

Syntax **import** *policy-name* [*policy-name*... (up to 5 max)]

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command specifies the import policy to be used to control routes advertised from RIP neighbors.

By default, RIP accepts all routes from RIP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

The **no** form of the command removes all route policy names from the import list.

Default no import

Parameters *policy-name* — the 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.

message-size

Syntax **message-size** *max-num-of-routes*
no message-size

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command configures the maximum number of routes per RIP update message.

By default, each update can contain a maximum of 25 route advertisements. This limit is imposed by RIP specifications. RIP can be configured to send as many as 255 routes per update.

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

Default no message-size

Parameters *max-num-of-routes* — an integer value

Default 25

Values 25 to 255

metric-in

Syntax **metric-in** *metric*
no metric-in

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command configures the metric added to routes received from a RIP neighbor. The specified metric value is added to the hop count and shortens the maximum distance of the route.

When applying an export policy to a RIP configuration, the policy overrides the metric values determined through calculations involving the **metric-in** and **metric-out** values.

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

Default no metric-in

Parameters *metric* — the value added to the metric of routes received from a RIP neighbor, expressed as a decimal integer

Values 1 to 16

metric-out

Syntax **metric-out** *metric*
no metric-out

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command configures the metric added to routes exported into RIP and advertised to RIP neighbors. The specified metric value is added to the hop count and shortens the maximum distance of the route.

When applying an export policy to a RIP configuration, the policy overrides the metric values determined through calculations involving the **metric-in** and **metric-out** values.

The **no** form of the command removes the command from the configuration and reverts the metric-in value to the default.

Default no metric-out

Parameters *metric* — the value added to the metric of routes exported into RIP and advertised to RIP neighbors, expressed as a decimal integer

Values 1 to 16

neighbor

Syntax [**no**] **neighbor** *ip-int-name*

Context config>service>vprn>rip>group

Description This command creates a context for configuring a RIP neighbor interface.

By default, interfaces are not activated unless explicitly configured.

The **no** form of the command deletes the RIP interface configuration for this interface. The **shutdown** command in the **config>router>rip>group**>**neighbor** context can be used to disable an interface without removing the configuration for the interface.

Default no neighbor

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.

preference

Syntax **preference** *preference*
no preference

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command configures the route preference assigned to RIP routes. This value can be overridden by route policies.

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

Default no preference

Parameters *preference* — the route preference, expressed as an integer value

Default 100

Values 1 to 255

propagate-metric

Syntax [**no**] **propagate-metric**

Context config>service>vprn>rip

Description This command allows the RIP metric to be used to set the MP-BGP MED attribute when RIP is used as the CE-PE routing protocol for VPRNs. This is similar to the way the OSPF metric can be used to set the MP-BGP metric when OSPF is used as the CE-PE protocol.

MP-BGP uses the RIP metric to set the MED attribute, which is flooded throughout the MP-BGP peers and is then used to set the RIP metric at the other end and re-advertise the RIP metric to the far-end RIP neighbors.

receive

Syntax **receive** *receive-type*
no receive

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command configures the types of RIP updates that will be accepted and processed.

If you specify both or version-2, the RIP instance listens for, and accepts, packets sent to the broadcast (255.255.255.255) and multicast (224.0.0.9) addresses.

If version-1 is specified, the router only listens for and accepts packets sent to the broadcast address.

The default behavior is to accept and process both RIPv1 and RIPv2 messages.

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

Default both

Parameters *receive-type* — configures the type of RIP updates that will be accepted and processed

Values *receiver-type* values are **both**, **none**, **version-1**, and **version-2**, where:

both	specifies that RIP updates in either version 1 or version 2 format will be accepted
none	specifies that RIP updates will not be accepted
version-1	specifies that only RIP updates in version 1 format will be accepted
version-2	specifies that only RIP updates in version 2 format will be accepted

send

Syntax **send** *send-type*
no send

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command specifies the type of RIP messages sent to RIP neighbors.

If **multicast** is specified, the router sends RIPv2 messages to the multicast (224.0.0.9) destination address.

If **broadcast**, or **version-1** is specified, the router only listens for and accepts packets sent to the broadcast address.

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

Default broadcast

Parameters *send-type* — configures the type of RIP messages that will be sent to RIP neighbors

Values *send-type* values are **broadcast**, **multicast**, **none**, and **version-1**, where:

broadcast	sends RIPv2 formatted messages to the broadcast address

multicast	sends RIPv2 formatted messages to the multicast address
none	does not to send any RIP messages (silent listener)
version-1	sends RIPv1 formatted messages to the broadcast address

split-horizon

Syntax **split-horizon** {**enable** | **disable**}
no split-horizon

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command enables the use of split-horizon. RIP uses split-horizon with poison-reverse to protect from such problems as "counting to infinity". Split-horizon with poison reverse means that routes learned from a neighbor through an interface are advertised in updates out of the same interface but with a metric of 16 (infinity).

The split-horizon disable command enables split-horizon without poison-reverse. This allows the routes to be re-advertised on interfaces other than the interface that learned the route, with the advertised metric equaling an increment of the metric-in value.

This parameter can be set at three levels: global level (applies to all groups and neighbor interfaces), group level (applies to all neighbor interfaces in the group), or neighbor level (only applies to the specified neighbor interface). The most specific value is used. If no value is set (**no split-horizon**), the setting from the less-specific level is inherited by the lower level.

The **no** form of the command disables split-horizon.

Default enable

Parameters **enable —** enables split-horizon and poison-reverse

disable — disables poison-reverse but leaves split-horizon enabled

timers

Syntax **timers** *update timeout flush*

Context config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor

Description This command configures values for the update, timeout, and flush RIP timers.

The RIP update timer determines how often RIP updates are sent.

If the route is not updated by the time the RIP timeout timer expires, the route is declared invalid but is maintained in the RIP database.

The RIP flush timer determines how long a route is maintained in the RIP database after it has been declared invalid. After the flush timer expires, the route is removed from the RIP database.

The **no** form of the command reverts all timers to their default values.

Default no timers

Parameters *update* — the RIP update timer value, in seconds, expressed as a decimal integer

Values 1 to 600

Default 30

timeout — the RIP timeout value, in seconds, expressed as a decimal integer

Values 1 to 1200

Default 180

flush — the RIP flush timer value, in seconds, expressed as a decimal integer

Values 1 to 1200

Default 120

VPRN Security Configuration Commands

zone

Syntax **zone** {*zone-id* | *zone-name*} [**create**]
no zone *zone-id*

Context config>service>vprn

Description This command creates or specifies a security zone within a VPRN context. Each zone must have a unique ID.

All zones must be explicitly created with the **create** keyword. If no zones are created within a service or router context, a zone will not exist on that object.

Enter an existing zone without the **create** keyword to edit zone parameters.

The **no** form of this command deletes the zone. When a zone is deleted, all configuration parameters for the zone are also deleted.

Parameters *zone-id* — the zone ID number. The zone ID must be unique within the system.

Values 1 to 65534

abort

Syntax **abort**

Context config>service>vprn>zone

Description This command discards changes made to a security feature.

Default n/a

begin

Syntax **begin**

Context config>service>vprn>zone

Description This command enters the mode to create or edit security features.

Default n/a

commit

Syntax **commit**

Context config>service>vprn>zone

Description This command saves changes made to security features.

Default n/a

auto-bind

Syntax **auto-bind**
no auto-bind

Context config>service>vprn>zone

Description This command creates a security zone on automatically bound GRE, MPLE, or LDP transport tunnels configured for this service. Depending on how the security policy is configured, any traffic entering or exiting the zone is firewalled; traffic traveling between auto-bind LSPs in the zone is not firewalled.

Default n/a

inbound

Syntax **inbound**

Context config>service>vprn>zone

Description This command enables the context to configure limit parameters on inbound security sessions.

Default n/a

outbound

Syntax **outbound**

Context config>service>vprn>zone

Description This command enables the context to configure limit parameters for outbound security sessions on the CSM.

Default n/a

limit

Syntax **limit**

Context config>service>vprn>zone>inbound
config>service>vprn>zone>outbound

Description This command enables the context to configure limits on concurrent sessions for inbound or outbound firewall sessions on the CSM.

Default n/a

concurrent-sessions

Syntax **concurrent-sessions** {**tcp** | **udp** | **icmp** | **other**} *sessions*
no concurrent-sessions {**tcp** | **udp** | **icmp** | **other**}

Context config>service>vprn>zone>inbound>limit
config>service>vprn>zone>outbound>limit

Description This command configures the maximum number of concurrent firewall sessions that can be established per zone, in either the inbound or outbound direction.

Default n/a

Parameters **tcp —** specifies that TCP connection traffic is to be firewalled

udp — specifies that UDP connection traffic is to be firewalled

icmp — specifies that ICMP connection traffic is to be firewalled

other — specifies that the traffic to be firewalled is other than TCP, UDP, or ICMP

sessions — the maximum number of concurrent firewall sessions that can be created in a zone for the configured direction

Values 1 to 16383

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>service>vprn>zone

Description This command creates a logical IP routing interface for a zone. Once created, attributes such as an IP address can be associated with the IP interface. Multiple interfaces can be configured on a zone.

The **no** form of this command removes the IP interface and all the associated configurations.

Parameters *ip-int-name* — the name of the interface to be configured within the zone

Values 1 to 32 characters (must start with a letter)

log

Syntax **log** {*log-id* | *name*}
no log

Context config>service>vprn>zone

Description This command configures a log identifier for the specified zone. A log identifier can be configured in the `config>router>zone` context and the `config>security>policy` context.

The **no** form of this command removes logging for the zone.

Parameters *log-id* — specifies the identifier for the log

Values 1 to 32 characters

name

Syntax **name** *zone-name*
no name

Context config>service>vprn>zone

Description This command configures a zone name. The zone name is unique within the system. It can be used to refer to the zone under configure, show, and clear commands.

Parameters *zone-name* — specifies the name of the zone

Values 1 to 32 characters (must start with a letter). If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

nat

Syntax **nat**

Context config>service>vprn>zone

Description This command enters the context to configure NAT security parameters for a zone.

pool

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

Context config>service>vprn>zone>nat

Description This command configures the NAT pool for the security zone within a VPRN service. Each pool must have a unique ID.

All pools must be explicitly created with the **create** keyword.

Enter an existing pool without the **create** keyword to edit pool parameters.

The **no** form of this command deletes the specified NAT pool. When a pool is deleted, all configuration parameters for the pool will also be deleted.

Parameters *pool-id* — the pool ID number

Values 1 to 100

direction

Syntax **direction** {**zone-outbound** | **zone-inbound** | **both**}
no direction

Context config>service>vprn>zone>nat>pool

Description This command configures the NAT pool direction for the security zone. A specific NAT pool can be configured for different directions while using the same policy. For example, if the security policy entry direction is set to **both**, separate inbound and outbound pools can be created for that policy.

The **no** form of this command deletes the direction.

Parameters **zone-outbound —** configures a pool for the policy outbound traffic

zone-inbound — configures a pool for the policy inbound traffic

both — configures a pool for policy inbound and outbound traffic

entry

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

Context config>service>vprn>zone>nat>pool

Description This command configures a NAT pool entry within a VPRN service.

The **no** form of this command deletes the entry with the specified ID. When an entry is deleted, all configuration parameters for the entry will also be deleted.

Parameters *entry-id* — the entry ID number

Values 1 to 65535

ip-address

Syntax **ip-address** *ip-address* [**to** *ip-address*] **interface** *ip-int-name*
no ip-address

Context config>service>vprn>zone>nat>pool>entry

Description This command configures the source IP address or IP address range to which packets that match NAT policy are routed using NAT. An interface can also be configured, in which case all packets that match NAT policy are routed to the interface IP address. If the interface IP address is changed dynamically, NAT is updated accordingly. Only one IP address can be associated with an IP interface. Source IP addresses and interfaces cannot be used together in a single NAT pool.

The IP address for the interface must be entered in dotted-decimal notation.

The **no** form of the command removes the IP address assignment. The **no** form of this command can only be performed when the IP interface is administratively shut down. Shutting down the IP interface brings the interface operationally down.

Parameters *ip-address* — the source IP address to be used by NAT. The *ip-address* portion of the **ip-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 1.0.0.0 to 223.255.255.255

ip-int-name — the name of the interface to be used by NAT

port

Syntax **port** *port* [**to** *port*]
no port

Context config>service>vprn>zone>nat>pool>entry

Description This command configures the UDP/TCP port or port range. Packets that match NAT policy undergo network port address translation (NPAT) and are routed to their source UDP/TCP port. Configuring a UDP/TCP port pool requires an IP-address pool because the 7705 SAR does not support port address translation (PAT) alone.

The **no** form of this command deletes the port or port range.

Parameters *port* — the UDP/TCP port or range of ports to which NPAT is applied

name

Syntax **name** *pool-name*
no name

Context config>service>vprn>zone>nat>pool

Description This command configures a zone pool name. Pool names must be unique within the group of pools defined for a zone. It can be used to refer to the pool under configure, show, and clear commands.

Parameters *pool-name* — specifies the name of the pool. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

Values 1 to 32 characters (must start with a letter).

policy

Syntax **policy** {*policy-id* | *policy-name*}
no policy *policy-id*

Context config>service>vprn>zone

Description This command sets the policy to be used by the security zone to build its matching criteria for incoming packets.

The **no** form of this command deletes the specified policy.

Parameters *policy-id* — the number of the referenced policy

Values 1 to 65535

Router Advertisement Commands

router-advertisement

Syntax [**no**] **router-advertisement**

Context config>service>vprn

Description This command enables the context to configure router advertisement properties for all VPRN IPv6-enabled interfaces. By default, the command is disabled for all IPv6-enabled interfaces.

The **no** form of the command disables router advertisement on all IPv6 interfaces.

Default no router-advertisement

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>service>vprn>router-advertisement

Description This command configures router advertisement properties on a specified interface. The interface name must already exist in the **config>service>vprn>interface** context.

The **no** form of the command disables router advertisement on the specified router interface.

Default n/a

Parameters *ip-int-name* — a 1 to 32 character name (must start with a letter) of the IP interface. Interface names must be unique within the group of defined IP interfaces for the **config>service>vprn>interface** command. 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 within double quotes.

current-hop-limit

Syntax **current-hop-limit** *number*
no current-hop-limit

Context config>service>vprn>router-advertisement>interface

Description This command configures the current hop limit in the router advertisement messages. It informs the nodes on the subnet about the hop limit when originating IPv6 packets.

Default 64

Parameters *number* — the hop limit

Values 0 to 255 (a value of 0 means that there are an unspecified number of hops)

managed-configuration

Syntax [**no**] **managed-configuration**

Context config>service>vprn>router-advertisement>interface

Description This command sets the managed address configuration flag. This flag indicates that DHCPv6 is available for address configuration in addition to any address autoconfigured using stateless address autoconfiguration.

Default no managed-configuration

max-advertisement-interval

Syntax **max-advertisement-interval** *seconds*
no max-advertisement-interval

Context config>service>vprn>router-advertisement>interface

Description This command configures the maximum interval between sending router advertisement messages.

Default 600

Parameters *seconds* — the maximum interval, in seconds, between sending router advertisement messages

Values 4 to 1800

min-advertisement-interval

Syntax **min-advertisement-interval** *seconds*
no min-advertisement-interval

Context config>service>vprn>router-advertisement>interface

Description This command configures the minimum interval between sending ICMPv6 router advertisement messages.

Default 200

Parameters *seconds* — the minimum interval, in seconds, between sending ICMPv6 router advertisement messages

Values 3 to 1350

mtu

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

Context config>service>vprn>router-advertisement>interface

Description This command configures the MTU for the nodes to use when sending packets on the link.

The **no** form of the command means that the MTU option is not sent in the router advertisement messages.

Default no mtu

Parameters *mtu-bytes* — the MTU for the nodes to use when sending packets

Values 1280 to 9212

other-stateful-configuration

Syntax [**no**] **other-stateful-configuration**

Context config>router>vprn>router-advertisement>interface

Description This command sets the "Other configuration" flag. This flag indicates that DHCPv6lite is available for autoconfiguration of other (non-address) information such as DNS-related information or information on other servers in the network.

Default no other-stateful configuration

prefix

Syntax **prefix** *ipv6-prefix/prefix-length*
no prefix

Context config>service>vprn>router-advertisement>interface

Description This command configures an IPv6 prefix in the router advertisement messages. To support multiple IPv6 prefixes, use multiple prefix statements. No prefix is advertised until it is explicitly configured using prefix statements.

Default n/a

Parameters *ipv6-prefix/prefix-length* — the IPv6 prefix

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 to FFFF]H d: [0 to 255]D
prefix-length	0 to 128

autonomous

Syntax [**no**] **autonomous**

Context config>service>vprn>router-advertisement>if>prefix

Description This command specifies whether the prefix can be used for stateless address autoconfiguration.

Default autonomous

on-link

Syntax [**no**] **on-link**

Context config>service>vprn>router-advertisement>if>prefix

Description This command specifies whether the prefix can be used for on-link determination.

Default on-link

preferred-lifetime

Syntax **preferred-lifetime** {*seconds* | **infinite**}
no preferred-lifetime

Context config>service>vprn>router-advertisement>if>prefix

Description This command configures the time that this prefix will continue to be preferred (time until deprecation). The address generated from a deprecated prefix should not be used as a source address in new communications. However, packets received on such an interface are processed as expected.

Default 604800

Parameters *seconds* — the length of time, in seconds, that this prefix will be preferred

Values 1 to 4294967294

infinite — the prefix will always be preferred. A value of 4294967295 also represents infinity.

valid-lifetime

Syntax **valid-lifetime** {*seconds* | **infinite**}
no valid-lifetime

Context config>service>vprn>router-advertisement>if>prefix

Description This command specifies the length of time, in seconds, that the prefix is valid for the purpose of onlink determination. The address generated from an invalidated prefix should not appear as the destination or source address of a packet.

Default 2592000

Parameters *seconds* — the remaining length of time, in seconds, that this prefix will be valid

Values 1 to 4294967294

infinite — the prefix will always be valid. A value of 4294967295 also represents infinity.

reachable-time

Syntax **reachable-time** *milli-seconds*
no reachable-time

Context config>service>vprn>router-advertisement>interface

Description This command configures how long the router should be considered reachable by other nodes on the link after receiving a reachability confirmation.

Default no reachable-time

Parameters *milli-seconds* — the length of time, in milliseconds, that the router should be considered reachable

Values 0 to 3600000

retransmit-time

Syntax **retransmit-time** *milli-seconds*
no retransmit-time

Context config>service>vprn>router-advertisement>interface

Description This command configures the retransmission frequency of neighbor solicitation messages.

Default no retransmit-time

Parameters *milli-seconds* — the amount of time, in milliseconds, that a host should wait before retransmitting neighbor solicitation messages

Values 0 to 1800000

router-lifetime

Syntax **router-lifetime** *seconds*
no router-lifetime

Context config>service>vprn>router-advertisement>interface

Description This command configures the router lifetime.

Default no router-lifetime

Parameters *seconds* — the length of time, in seconds, that the prefix is valid for route determination

Values 0, 4 to 9000 (a value of 0 means that the router is not a default router on this link)

Local DHCP and DHCPv6 Server Commands

For complete descriptions of all local DHCP and DHCPv6 server commands, refer to the Router Configuration Guide, “Local DHCP and DHCPv6 Server Commands”.

local-dhcp-server

Syntax **local-dhcp-server** *server-name* [**create**]
no local-dhcp-server *server-name*

Context config>service>vprn>dhcp
config>service>vprn>dhcp6

Description This command creates a local DHCP or DHCPv6 server instance. A local DHCP or DHCPv6 server can serve multiple interfaces but is limited to the routing context in which it was created.

The **no** form of the command removes the local DHCP or DHCPv6 server instance.

Default n/a

Parameters *server-name* — the name of the local DHCP or DHCPv6 server

Values up to 32 alphanumeric characters

create — keyword is mandatory when creating a local DHCP or DHCPv6 server

Interface Commands

interface

Syntax **interface** *ip-int-name*
no interface *ip-int-name*

Context config>service>vprn

Description This command creates a logical IP routing interface for a Virtual Private Routed Network (VPRN). Once created, attributes such as an IP address and a 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 a VPRN service ID. The IP interface created is associated with the VPRN service routing instance and VPRN service routing table.

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.

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.

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 the interface and all the associated configurations. The interface must be administratively shut down before issuing the **no interface** command.

Parameters *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 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 and must start with a letter. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*} [**broadcast** {**all-ones** | **host-ones**}]
no address {*ip-address/mask* | *ip-address netmask*}

Context config>service>vprn>interface

Description This command assigns an IP address, IP subnet, and broadcast address format to a VPRN IP router 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 7705 SAR.

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, which 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, as shown in Table 109.

Table 109: VPRN Interface State and IP Address

Address	Administrative State	Operational 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 administrative states. The address and administrative states are independent and can be set independently. If an interface is in an administratively up state and an address is assigned, it becomes operationally up and the protocol interfaces and the MPLS LSPs associated with that IP interface will be reinitialized.

Parameters *ip-address* — the IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d (no multicast address)
(1.0.0.0 to 223.255.255.255 (with support of /31 subnets)

The "/" (forward slash) is a parameter delimiter that separates the *ip-address* portion of the IP address from the *mask*, which defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "/", and the *mask*. If a forward slash does not immediately follow the *ip-address*, a dotted-decimal *mask* must follow the prefix.

mask — 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*. The *mask* 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.

Values 0 to 32

netmask — the subnet mask, in dotted-decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted-decimal mask. The *netmask* parameter indicates the complete mask that will be used in a logical "AND" function to derive the local subnet of the IP address.

Values 128.0.0.0 to 255.255.255.252
(network bits all 1 and host bits all 0)
(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 indicates 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 — specifies that 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 — 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*, or the *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 negation 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 **allow-directed-broadcasts**
no allow-directed-broadcasts

Context config>service>vprn>interface

Description 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 for the subnet broadcast address of the egress IP interface.

When enabled, a frame destined for 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

arp-retry-timer

Syntax **arp-retry-timer** *ms-timer*
no arp-retry-timer

Context config>service>vprn>interface

Description This command specifies the length of time, in 100s of milliseconds, that the system waits before reissuing a failed ARP request.

The **no** form of the command resets the interval to the default value.



Note: The ARP retry default value of 5000 ms is intended to protect CPU cycles on the 7705 SAR, especially when it has a large number of interfaces. Configuring the ARP retry timer to a value shorter than the default should be done only on mission-critical links, such as uplinks or aggregate spoke SDPs transporting mobile traffic; otherwise, the retry interval should be left at the default value.

Default 50 (in 100s of ms)

Parameters *ms-timer* — the time interval, in 100s of milliseconds, the system waits before retrying a failed ARP request

Values 1 to 300

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>vprn>interface

Description This command configures the minimum time, in seconds, that 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 0 s, ARP aging is disabled.

The **no** form of this command restores **arp-timeout** to the default value.

Default 14400 s

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 0 specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 to 65535

bfd

Syntax **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*] [**echo-receive** *echo-interval*] [**type np**]
no bfd

Context config>service>vprn>interface
config>service>vprn>interface>ipv6

Description This command specifies the BFD parameters for the associated IP interface. If no parameters are defined, the default values are used.

The multiplier specifies the number of consecutive BFD messages that must be missed from the peer before the BFD session state is changed to down. In addition, the Route Table Manager (RTM) is notified and the static routes with BFD enabled will go down, based on BFD status.

The **no** form of the command removes BFD from the associated IGP protocol adjacency.

Default no bfd

Parameters *transmit-interval* — sets the transmit interval for the BFD session

Values 10 to 100000 in milliseconds

Default 100

receive-interval — sets the receive interval for the BFD session

Values 10 to 100000 milliseconds

Default 100

multiplier — sets the multiplier for the BFD session

Values 3 to 20

Default 3

echo-interval — (does not apply to IPv6 interfaces) sets the minimum echo receive interval for the BFD session

Values 100 to 100000 milliseconds

Default 100

type np — (does not apply to IPv6 interfaces) controls the value range of the *transmit-interval* and *receive-interval* parameters. If the **type np** option is not specified, the range of the *transmit-interval* and *receive-interval* parameter values is from 100 ms to 100000 ms. If the **type np** option is specified, the range of the *transmit-interval* and *receive-interval* parameter values is from 10 ms to 1000 ms, with the restriction that the maximum receiving detection time for the missing BFD packets must be less than or equal to 3000 ms. The maximum receiving detection time is the *receive-interval* parameter multiplied by the *multiplier* parameter.

Note: The BFD session must be disabled before the **type np** parameter can be changed. The **type np** parameter is only supported on VPRN services for SAPs.

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>vprn>interface

Description This command configures the IP maximum transmit unit (packet) for this interface.

The default value is derived from the port MTU.

The **no** form of the command returns the default value.

Default no ip-mtu

Parameters *octets* — specifies the MTU for this interface

Values 512 to 9732

ipcp

Syntax **ipcp**

Context config>service>vprn>interface

Description This command allows access to the Internet Protocol Control Protocol (IPCP) context within the interface configuration. Within this context, IPCP extensions can be configured to define such things as the remote IP address and DNS IP address to be signaled via IPCP on the associated PPP interface.

This command is only applicable if the associated SAP/port is a PPP/MLPPP interface.

Default n/a

dns

Syntax **dns** *ip-address* [**secondary** *ip-address*]
dns secondary *ip-address*
no dns [*ip-address*] [**secondary** *ip-address*]

Context config>service>vprn>if>ipcp

Description This command defines the DNS address(es) to be assigned to the far end of the associated PPP/MLPPP link via IPCP extensions.

This command is only applicable if the associated SAP/port is a PPP/MLPPP interface with an IPCP encapsulation.

The **no** form of the command deletes the specified primary DNS address, the secondary DNS address, or both addresses from the IPCP extension **peer-ip-address** configuration.

Default no dns

Parameters *ip-address* — a unicast IPv4 address for the primary DNS server to be signaled to the far end of the associated PPP/MLPPP link via IPCP extensions

Values a.b.c.d (unicast only)

secondary *ip-address* **—** a unicast IPv4 address for the secondary DNS server to be signaled to the far end of the associated PPP/MLPPP link via IPCP extensions

Values a.b.c.d (unicast only)

peer-ip-address

Syntax **peer-ip-address** *ip-address*
no peer-ip-address

Context config>service>vprn>if>ipcp

Description This command defines the remote IP address to be assigned to the far end of the associated PPP/ MLPPP link via IPCP extensions.

This command is only applicable if the associated SAP/port is a PPP/MLPPP interface with an IPCP encapsulation.

The interface must be shut down to modify the IPCP configuration.

The **no** form of the command deletes the IPCP extension **peer-ip-address** configuration.

Default no peer-ip-address (0.0.0.0)

Parameters *ip-address* — a unicast IPv4 address to be signaled to the far end of the associated PPP/ MLPPP link by IPCP extensions

Values a.b.c.d (unicast only)

l4-load-balancing

Syntax **l4-load-balancing** {**includeL4** | **excludeL4**}
no l4-load-balancing

Context config>service>vprn>interface

Description This command configures Layer 4 load balancing at the interface level. Configuration must be done on the ingress network interface (that is, the interface on the node that the packet is received on). When enabled, Layer 4 source and destination port fields of incoming TCP/UDP packets are included in the hashing calculation to randomly determine the distribution of packets.

You can add additional fields to generate more randomness and more equal distribution of packets with the teid-load-balancing command.

The default configuration on the interface is to match the Layer 4 load balancing configuration in the **config>system** context. Using this command to modify Layer 4 load-balancing configuration on an interface overrides the system-wide load-balancing settings for that interface.

Parameters **includeL4 —** layer 4 source and destination port fields are included in the hashing calculation for TCP/UDP packets

excludeL4 — layer 4 source and destination port fields are not included in the hashing calculation for TCP/UDP packets

local-dhcp-server

Syntax [**no**] **local-dhcp-server** *local-server-name*

Context config>service>vprn>interface

Description This command associates the interface with a local DHCP server configured on the system. A routed VPLS interface may not be associated with a local DHCP server.

The **no** form of the command removes the association of the interface with the local DHCP server.

Default none

Parameters *local-server-name* — the name of the local DHCP server

Values up to 32 alphanumeric characters

local-proxy-arp

Syntax [**no**] **local-proxy-arp**

Context config>service>vprn>interface

Description This command enables local proxy ARP on the interface.

Local proxy ARP allows the 7705 SAR to respond to ARP requests received on an interface for an IP address that is part of a subnet assigned to the interface. The router responds to all requests for IP addresses within the subnet with its own MAC address and forwards all traffic between the hosts in the subnet.

Local proxy ARP is used on subnets where hosts are prevented from communicating directly.

When **local-proxy-arp** is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

Default no local-proxy-arp

loopback

Syntax [**no**] **loopback**

Context config>service>vprn>interface

Description 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.

Default none

mac

Syntax **mac** *ieee-address*
no mac [*ieee-address*]

Context config>service>vprn>interface

Description 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-address* — the 48-bit MAC address for the static ARP. Allowed values are any non-broadcast, non-multicast MAC, and non-IEEE reserved MAC addresses.

Values *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff*
(where *aa, bb, cc, dd, ee,* and *ff* are hexadecimal numbers and cannot be all zeros)

proxy-arp-policy

Syntax **proxy-arp-policy** *policy-name* [*policy-name*...(up to 5 max)]
no proxy-arp-policy

Context config>service>vprn>interface

Description This command enables proxy ARP on the interface and specifies an existing policy statement that controls the flow of routing information by analyzing match and action criteria. The policy statement is configured in the **config>router>policy-options** context (refer to the 7705 SAR OS Router Configuration Guide, "Route Policy Command Reference, Route Policy Options"). When proxy ARP is enabled, the 7705 SAR responds to ARP requests on behalf of another device.

Default no proxy-arp-policy

Parameters *policy-name* — the route policy statement 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 policy statement must already be defined.

remote-proxy-arp

Syntax [**no**] **remote-proxy-arp**

Context config>service>vprn>interface

Description This command enables remote proxy ARP on the interface, allowing a router on one network to respond to ARP requests intended for another node that is physically located on another network. The router effectively pretends to be the destination node by sending an ARP response to the originating node that associates the router's MAC address with the destination node's IP address (acts as a proxy for the destination node). The router then takes responsibility for routing traffic to the real destination.

Default no remote-proxy-arp

secondary

Syntax **secondary** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**igp-inhibit**]
no secondary {*ip-address/mask* | *ip-address netmask*}

Context config>service>vprn>interface

Description This command assigns an secondary IP address, IP subnet, and broadcast address format to the interface.

Default no secondary

Parameters *ip-address* — the IP address of the IP interface. The *ip-address* portion of the **secondary** 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 a.b.c.d

(1.0.0.0 to 223.255.255.255 (with support of /31 subnets)

The "/" (forward slash) is a parameter delimiter that separates the *ip-address* portion of the IP address from the *mask*, which defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "/", and the *mask*. If a forward slash does not immediately follow the *ip-address*, a dotted-decimal *mask* must follow the prefix.

mask — 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*. The *mask* 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.

Values 0 to 32

netmask — the subnet mask, in dotted-decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted-decimal mask. The *netmask* parameter indicates the complete mask that will be used in a logical "AND" function to derive the local subnet of the IP address.

Values 128.0.0.0 to 255.255.255.252
(network bits all 1 and host bits all 0)
(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 indicates 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 — specifies that 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 — 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*, or the *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 **secondary** command does not have a negation 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 **secondary** command must be executed with the **broadcast** parameter defined.

igp-inhibit — specifies that this secondary IP interface should not be recognized as a local interface by the running IGP. For OSPF and IS-IS, this means that the secondary IP interface will not be injected and used as a passive interface and will not be advertised as an internal IP interface into the IGP link state database. For RIP, this means that the secondary IP interface will not source RIP updates.

static-arp

Syntax **static-arp** *ip-address ieee-address*
no static-arp *ip-address* [*ieee-address*]
static-arp *ieee-address* **unnumbered**
no static-arp [*ieee-address*] **unnumbered**

Context config>service>vprn>interface

Description This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. 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.

Static ARP is used when a 7705 SAR needs to know about a device on an interface that cannot or does not respond to ARP requests. Therefore, the 7705 SAR OS configuration can specify to send a packet with a particular IP address to the corresponding ARP address.

The **no** form of the command removes a static ARP entry.

Default n/a

Parameters *ip-address* — the IP address for the static ARP in dotted-decimal notation

Values a.b.c.d

ieee-address — the 48-bit MAC address for the static ARP. Allowed values are any non-broadcast, non-multicast MAC, and non-IEEE reserved MAC addresses.

Values *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff*
(where *aa, bb, cc, dd, ee*, and *ff* are hexadecimal numbers and cannot be all zeros)

unnumbered — specifies the static ARP MAC addresses for an unnumbered interface. Unnumbered interfaces also support dynamic ARP. If this parameter is configured, it overrides any dynamic ARP.

teid-load-balancing

Syntax **[no] teid-load-balancing**

Context config>service>vprn>interface

Description This command configures TEID load balancing at the interface level. Configuration must be done on the ingress network interface (that is, the interface on the node that the packet is received on). The TEID attribute is included in the header of GTP (general packet radio system tunneling protocol) packets. When TEID load balancing is enabled, the TEID field of incoming TCP/UDP packets is included in the hashing calculation to randomly determine the distribution of packets.

You can add additional fields to generate more randomness and more equal distribution of packets with the l4-load-balancing command.

Default no teid-load-balancing

unnumbered

Syntax **unnumbered** {*ip-int-name* | *ip-address*}
no unnumbered

Context config>service>vprn>interface

Description This command configures an IP interface as an unnumbered interface and specifies an IP address or interface name to be used for the interface. Unnumbered interfaces are point-to-point interfaces that are not explicitly configured with a dedicated IP address and subnet; instead, they borrow (or link to) an IP address from another interface on the system (the system IP address, another loopback interface, or any other numbered interface) and use it as the source IP address for packets originating from the interface.

By default, no IP address exists on an IP interface until it is explicitly created.

The **no** form of the command removes the IP address assignment from the IP interface.

Default no unnumbered

Parameters *ip-int-name* | *ip-address* — the IP interface name or address to associate with the unnumbered IP interface

Values *ip-int-name*: 1 to 32 characters (must start with a letter)
ip-address: a.b.c.d

IPv6 Interface Commands

ipv6

Syntax [**no**] **ipv6**

Context config>service>vprn>interface

Description This command enables the context to configure parameters for a VPRN IPv6 interface.

address

Syntax **address** *ipv6-address/prefix-length* [**eui-64**] [**preferred**]
no address *ipv6-address/prefix-length*

Context config>service>vprn>if>ipv6

Description This command assigns an address to the IPv6 interface.

Parameters *ipv6-address/prefix-length* — the address of the IPv6 interface

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
prefix-length	0 to 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 the MAC address on Ethernet interfaces. For interfaces without a MAC address, for example ATM interfaces, the base MAC address of the chassis is used

preferred — specifies that the IPv6 address is the preferred IPv6 address for this interface. A preferred address is an address assigned to an interface whose use by upper layer protocols is unrestricted. A preferred addresses maybe used as the source (or destination) address of packets sent from (or to) the interface

dhcp6-relay

Syntax [**no**] **dhcp6-relay**

Context config>service>vprn>if>ipv6

Description This command enables the context to configure DHCPv6 relay parameters for the interface.

The **no** form of the command disables DHCPv6 relay.

option

Syntax [**no**] **option**

Context config>service>vprn>if>ipv6>dhcp6-relay

Description This command enables the context to configure DHCPv6 relay information options.

The **no** form of the command disables DHCPv6 relay information options.

interface-id

Syntax **interface-id**
interface-id ascii-tuple
interface-id ifindex
interface-id sap-id
interface-id *string*
no interface-id

Context config>service>vprn>if>ipv6>dhcp6-relay>option

Description This command enables the sending of interface ID options in the DHCPv6 relay packet.

The **no** form of the command disables the sending of interface ID options in the DHCPv6 relay packet.

Parameters **ascii-tuple —** specifies that the ASCII-encoded concatenated tuple will be used (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>if>detail**.)

sap-id — specifies that the SAP identifier will be used

string — a string of 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-id

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

Context config>service>vprn>if>ipv6>dhcp6-relay>option

Description This command enables the sending of the remote ID option in the DHCPv6 relay packet.

The client DHCP Unique Identifier (DUID) is used as the remote ID.

The **no** form of the command disables the sending of remote ID option in the DHCPv6 relay packet.

server

Syntax [**no**] **server** *ipv6z-address*

Context config>service>vprn>if>ipv6>dhcp6-relay

Description This command configures an IPv6 address to the DHCPv6 server.

The **no** form of the command disables the specified IPv6 address.

Parameters *ipv6z-address* — the IPv6 address of the DHCPv6 server (a maximum of eight addresses can be configured)

Values *ipv6z-address* x:x:x:x:x:x:x:x [-interface]
x:x:x:x:x:x:d.d.d.d [-interface]
x: [0 to FFFF]H
d: [0 to 255]D
interface: 32 characters maximum,
mandatory for link local addresses

source-address

Syntax [**no**] **source-address** *ipv6-address*

Context config>service>vprn>if>ipv6>dhcp6-relay

Description This command assigns the source IPv6 address of the DHCPv6 relay messages.

The **no** form of the command disables the specified IPv6 address.

Parameters *ipv6-address* — the source IPv6 address of the DHCPv6 relay messages

Values *ipv6-address* x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D

dhcp6-server

Syntax [**no**] **dhcp6-server**

Context config>service>vprn>if>ipv6

Description This command enables the context to configure DHCPv6 server parameters for the VPRN interface.

The **no** form of the command disables the DHCPv6 server.

max-nbr-of-leases

Syntax **max-nbr-of-leases** *max-nbr-of-leases*
no max-nbr-of-leases

Context config>service>vprn>if>ipv6>dhcp6-server

Description This command configures the maximum number of lease states installed by the DHCPv6 server function allowed on this interface.

The **no** form of the command returns the value to the default.

Default 8000

Parameters *max-nbr-of-leases* — the maximum number of lease states installed by the DHCPv6 server function allowed on this interface

Values 0 to 8000

prefix-delegation

Syntax [**no**] **prefix-delegation**

Context config>service>vprn>if>ipv6>dhcp6-server

Description This command configures prefix delegation options for delegating a long-lived prefix from a delegating router to a requesting router, where the delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned.

The **no** form of the command disables prefix delegation.

prefix

Syntax [**no**] **prefix** *ipv6-address/prefix-length*

Context config>service>vprn>if>ipv6>dhcp6-server>pfx-delegate

Description This command specifies the IPv6 prefix that is delegated by the system.

Parameters *ipv6-address/prefix-length* — the address of the IPv6 interface

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
prefix-length	0 to 128

duid

Syntax **duid** *duid* [**iaid** *iaid*]
no duid

Context config>service>vprn>if>ipv6>dhcp6>pfx-delegate>prefix

Description This command configures the DHCP Unique Identifier (DUID) of the DHCPv6 server client.

Parameters *duid* — the ID of the requesting router. If set to a non-zero value, the prefix defined will only be delegated to this router. If set to 0, the prefix will be delegated to any requesting router.

iaid *iaid* — the identity association identification (IAID) from the requesting router that needs to match in order to delegate the defined prefix. If set to 0, no match on the received IAID is done.

preferred-lifetime

Syntax **preferred-lifetime** *seconds*
preferred-lifetime infinite
no preferred-lifetime

Context config>service>vprn>if>ipv6>dhcp6>pfx-delegate>prefix

Description This command configures the IPv6 prefix preferred lifetime. The preferred-lifetime value cannot be larger than the valid-lifetime value.

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

Default 604800 seconds (7 days)

Parameters *seconds* — the time, in seconds, that this prefix remains preferred

Values 1 to 4294967294

infinite — specifies that this prefix remains preferred infinitely

valid-lifetime

Syntax **valid-lifetime** *seconds*
valid-lifetime infinite
no valid-lifetime

Context config>service>vprn>if>ipv6>dhcp6>pfx-delegate>prefix

Description This command configures the time, in seconds, that the prefix is valid.

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

Default 2592000 seconds (30 days)

Parameters *seconds* — the time, in seconds, that this prefix remains valid

Values 1 to 4294967295

infinite — specifies that this prefix remains valid infinitely

icmp6

Syntax **icmp6**

Context config>service>vprn>if>ipv6

Description This command configures ICMPv6 parameters for the interface.

packet-too-big

Syntax **packet-too-big** [*number seconds*]
no packet-too-big

Context config>service>vprn>if>ipv6>icmp6

Description This command specifies whether, and how often, "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* — the number of "packet-too-big" ICMPv6 messages to send in the time frame specified by the *seconds* parameter

Values 10 to 1000

Default 100

seconds — the time frame, in seconds, that is used to limit the number of "packet-too-big" ICMPv6 messages issued

Values 1 to 60

Default 10

param-problem

Syntax **param-problem** [*number seconds*]
no packet-too-big

Context config>service>vprn>if>ipv6>icmp6

Description This command specifies whether, and how often, "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

Parameters *number* — the number of "parameter-problem" ICMPv6 messages to send in the time frame specified by the *seconds* parameter

Values 10 to 1000

Default 100

seconds — the time frame, in seconds, that is used to limit the number of "parameter-problem" ICMPv6 messages issued

Values 1 to 60

Default 10

time-exceeded

Syntax **time-exceeded** [*number seconds*]
no time-exceeded

Context config>service>vprn>if>ipv6>icmp6

Description This command specifies whether, and how often, "time-exceeded" ICMPv6 messages should be sent. When enabled, ICMPv6 "time-exceeded" messages are generated by this interface.

Default 100 10

Parameters *number* — the number of "time-exceeded" ICMPv6 messages are to be issued in the time frame specified by the *seconds* parameter

Values 10 to 1000

Default 100

seconds — the time frame, in seconds, that is used to limit the number of "time-exceeded" ICMPv6 messages to be issued

Values 1 to 60

Default 10

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>vprn>if>ipv6>icmp6

Description This command specifies whether, and how often, ICMPv6 host and network destination unreachable messages are generated by this interface.

Default 100 10

Parameters *number* — the number of destination unreachable ICMPv6 messages to send issued in the time frame specified by the *seconds* parameter

Values 10 to 1000

Default 100

seconds — the time frame, in seconds, that is used to limit the number of destination unreachable ICMPv6 messages to be sent

Values 1 to 60

Default 10

link-local-address

Syntax **link-local-address** *ipv6-address* [**preferred**]

Context config>service>vprn>if>ipv6

Description This command assigns the IPv6 link local address to the interface.

Default n/a

Parameters *ipv6-address* — the IPv6 link local address

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

preferred — specifies that the IPv6 address is the preferred IPv6 address for this interface. A preferred address is an address assigned to an interface whose use by upper layer protocols is unrestricted. A preferred addresses maybe used as the source (or destination) address of packets sent from (or to) the interface

neighbor

Syntax **neighbor** *ipv6-address mac-address*
no neighbor *ipv6-address*

Context config>service>vprn>if>ipv6

Description This command configures IPv6-to-MAC address mapping on the interface.

Default n/a

Parameters *ipv6-address* — the address of the IPv6 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 to 128]

mac-address — 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 unicast MAC addresses and non-IEEE reserved MAC addresses

reachable-time

Syntax **reachable-time** *seconds*
no reachable-time

Context config>service>vprn>if>ipv6

Description This command specifies the time that an IPv6 neighbor remains in a reachable state.

Default no reachable-time

Parameters *seconds* — the number of seconds that an IPv6 neighbor remains in a reachable state

Values 30 to 3600

Default 30

stale-time

Syntax **stale-time** *seconds*
no stale-time

Context config>service>vprn>if>ipv6

Description This command specifies the time that an IPv6 neighbor cache entry remains in a stale state. When the specified time elapses, the system removes the neighbor cache entry.

Default no stale-time

Parameters *seconds* — the number of seconds that an IPv6 neighbor remains in a stale state

Values 60 to 65535

Default 14400

Interface DHCP Commands

dhcp

Syntax **dhcp**

Context config>service>vprn>interface

Description This command enables the context to configure DHCP parameters.

gi-address

Syntax **gi-address** *ip-address* [**src-ip-addr**]
no gi-address

Context config>service>vprn>if>dhcp

Description This command configures the gateway interface address for the DHCP Relay Agent. By default, the GIADDR used in the relayed DHCP packet is the primary address of an interface. Specifying the GIADDR allows the user to choose a secondary address.

Default no gi-address

Parameters *ip-address* — the IP address of the gateway interface in dotted-decimal notation

Values a.b.c.d (host bits must be 0)

src-ip-addr — specifies that the GIADDR is to be used as the source IP address for DHCP relay packets

option

Syntax [**no**] **option**

Context config>service>vprn>if>dhcp

Description This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 suboptions.

The **no** form of this command returns the system to the default.

Default no option

action

Syntax **action** {**replace** | **drop** | **keep**}
no action

Context config>service>vprn>if>dhcp>option

Description This command configures the processing required when the 7705 SAR 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 keep—(as 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 occurs if the gi-addr (gateway interface address) of the received packet is the same as the ingress address on the router. In this 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 the case of Vendor Specific Options (VSOs). When the **keep** parameter is specified, the router will insert its own VSO into the Option 82 field. This will only be done when the incoming message has an Option 82 field already.

If no Option 82 field is present, the router will not create the Option 82 field. In this 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 This command sends either an ASCII tuple or the interface index (If Index) on the specified SAP ID in the **circuit-id** suboption 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** suboption of the DHCP packet is left empty.

The **no** form of this command returns the system to the default.

Default ascii-tuple

Parameters **ascii-tuple —** the ASCII-encoded concatenated "tuple" will be used, where the "tuple" consists of the *access-node-identifier*, *service-id*, and *interface-name*, separated by the syntax symbol "|"

ifindex — the interface index will be used

sap-id — the SAP ID will be used

vlan-ascii-tuple — specifies that the format will include VLAN-id and dot1p bits, in addition to the **ascii-tuple**. 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.

remote-id

Syntax **remote-id** [**mac** | **string** *string*]
no remote-id

Context config>service>vprn>if>dhcp>option

Description This command sends the MAC address of the remote end (typically, the DHCP client) in the **remote-id** suboption of the DHCP packet. This command identifies the host at the other end of the circuit. If disabled, the **remote-id** suboption 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 —** the MAC address of the remote end is encoded in the suboption

string **—** the remote ID

Values up to 32 alphanumeric characters

vendor-specific-option

Syntax [**no**] **vendor-specific-option**

Context config>service>vprn>if>dhcp>option

Description This command enables the Alcatel-Lucent vendor-specific suboption of the DHCP relay packet and enters the context for configuring the vendor-specific suboptions.

client-mac-address

Syntax [**no**] **client-mac-address**

Context config>service>vprn>if>dhcp>option>vendor

Description This command enables the sending of the MAC address in the Alcatel-Lucent vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the Alcatel-Lucent vendor specific suboption of the DHCP relay packet.

Default no client-mac-address

sap-id

Syntax [**no**] **sap-id**

Context config>service>vprn>if>dhcp>option>vendor

Description This command enables the sending of the SAP ID in the Alcatel-Lucentt vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the Alcatel-Lucent vendor-specific suboption of the DHCP relay packet.

Default no sap-id

service-id

Syntax [**no**] **service-id**

Context config>service>vprn>if>dhcp>option>vendor

Description This command enables the sending of the service ID in the Alcatel-Lucent vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the Alcatel-Lucent vendor-specific suboption of the DHCP relay packet.

Default no service-id

string

Syntax **string** *text*
no string

Context config>service>vprn>if>dhcp>option>vendor

Description This command specifies the vendor-specific suboption string of the DHCP relay packet.

The **no** form of the command returns the default value.

Default no string

Parameters *text* — any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, the entire string must be enclosed within double quotes.

system-id

Syntax [**no**] **system-id**

Context config>service>vprn>if>dhcp>option>vendor

Description This command specifies whether the system ID is encoded in the Alcatel-Lucent vendor-specific suboption of Option 82.

Default n/a

server

Syntax **server** *server1* [*server2...*(up to 8 max)]
no server

Context config>service>vprn>if>dhcp

Description This command specifies a list of servers where requests will be forwarded. The list of servers can be 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 specified, 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* — the DHCP server IP address

trusted

Syntax [**no**] **trusted**

Context config>service>vprn>if>dhcp

Description This command enables or disables trusted mode on an IP interface.

According to RFC 3046, *DHCP Relay Agent Information Option*, a DHCP request where the gi-addr (gateway interface address) is 0.0.0.0 and which contains an 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 7705 SAR) will modify the request's gi-addr to be equal to the ingress interface and forward the request.

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. Thus, in this case, there is no reason for enabling the **trusted** option.

The **no** form of this command returns the system to the default.

Default no trusted

Interface ICMP Commands

icmp

Syntax **icmp**

Context config>service>vprn>interface

Description This command configures Internet Control Message Protocol (ICMP) parameters on a VPRN service and enters the context for configuring ICMP.

mask-reply

Syntax [**no**] **mask-reply**

Context config>service>vprn>if>icmp

Description This command enables responses to 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

ttl-expired

Syntax **ttl-expired** *number seconds*
no ttl-expired [*number seconds*]

Context config>service>vprn>if>icmp

Description This command configures the rate at which 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-s time interval.

The **no** form of this command disables 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 along with the *seconds* parameter.

Values 10 to 100

seconds — the time, in seconds, used to limit the number of ICMP TTL-expired messages that can be issued, expressed as a decimal integer

Values 1 to 60

unreachables

Syntax **unreachables** *number seconds*
no unreachables [*number seconds*]

Context config>service>vprn>if>icmp

Description This command enables and configures the rate of 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 *number* and *seconds* parameters by indicating the maximum number of destination unreachable messages that 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 100 per 10-s 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 along with the *seconds* parameter.

Values 10 to 100

seconds — the time frame, in seconds, used to limit the number of ICMP unreachable messages that can be issued

Values 1 to 60

Interface SAP Commands

sap

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

Context config>service>vprn>interface

Description This command creates a Service Access Point (SAP) within a service when used with the **create** keyword. The **create** keyword is not needed when entering an existing SAP to edit SAP parameters.

A SAP is a combination of port and encapsulation parameters that identify the service access point on the interface and within the 7705 SAR. 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.

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. Channelized TDM ports are always access ports.

If a port is shut down with the **shutdown** command, all SAPs on that port become operationally down. When a service is shut down, 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.

To configure a VPRN interface SAP that is used for a private IPSec tunnel interface, see sap in Service Interface Tunnel Commands.

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

Parameters *sap-id* — the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

create — keyword used to create a SAP instance

accounting-policy

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

Context config>service>vprn>if>sap

Description This command creates the accounting policy context that can be applied to an interface SAP.

An accounting policy must be defined before it can be associated with a SAP. Accounting policies are configured in the **config log** context. A maximum of one accounting policy can be associated with a SAP at one time.

If the *acct-policy-id* does not exist, an error message is generated.

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

Default no accounting policy

Parameters *acct-policy-id* — the accounting *policy ID* as configured in the **config log accounting-policy** context

Values 1 to 99

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>vprn>if>sap

Description This command enables accounting and statistical data collection for either an interface SAP 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. 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 collect-stats

egress

Syntax **egress**

Context config>service>vprn>if>sap

Description This command enables the context to configure egress SAP 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 policy is defined, no filtering is performed.

ingress

Syntax **ingress**

Context config>service>vprn>if>sap

Description This command enables the context to configure ingress SAP 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 policy is defined, no filtering is performed.

agg-rate-limit

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

Context config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress

Description This command sets the aggregate rate limits (PIR and CIR) for the SAP. The *agg-rate* sets the PIR value. The *cir-rate* sets the CIR value. When aggregate rate limits are configured on a second-generation (Gen-2) adapter card, the scheduler mode must be set to 16-priority. On the third-generation (Gen-3) Ethernet adapter card, only the PIR component (*agg-rate*) of the **agg-rate-limit** can be set and the scheduler mode is always 4-priority-hqos. For information on adapter card generations, refer to the "Evolution of Ethernet Adapter Cards, Modules, and Platforms" section in the 7705 SAR OS Interface Configuration Guide.

Configuring the *cir-rate* is optional. If a *cir-rate* is not entered, then the *cir-rate* is set to its default value (0 kb/s). If a *cir-rate* has been set and the *agg-rate* is changed without re-entering the *cir-rate*, the *cir-rate* automatically resets to 0 kb/s. For example, to change the *agg-rate* from 2000 to 1500 while maintaining a *cir-rate* of 500, use the command **agg-rate-limit** 1500 **cir** 500.

The *cir-rate* for this command is blocked on Gen-3 Ethernet adapter cards and Gen-3 platforms with Ethernet ports, such as the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X. Gen-3 Ethernet adapter cards and Gen-3 platforms with Ethernet ports support 4-priority-hqos scheduling mode. Gen-3 platforms with a TDM MDA support 4-priority scheduling mode.

If the specified SAP is a LAG SAP, then *agg-rate* can be configured regardless of the scheduler mode setting, and *cir-rate* can be configured but is not applied to Gen-3 Ethernet adapter cards or platforms. If the active port is on a Gen-3 card or platform, then only agg-rate is applicable. If the active port is on a Gen-2 card, then agg-rate and cir-rate apply when the scheduler mode is set to 16-priority. If the active port is on a Gen-1 card, then *agg-rate* and *cir-rate* are not applicable.

The **no** form of the command sets the *agg-rate* to the maximum and the *cir-rate* to 0 kb/s.

Default no agg-rate-limit

Parameters *agg-rate* — sets the PIR for the aggregate of all the queues on the SAP. The **max** keyword applies the maximum physical port rate possible.

Values 1 to 10000000 kb/s, or **max**

Default max

cir-rate — sets the CIR for the aggregate of all the queues on the SAP

Values 0 to 10000000 kb/s, or **max**

Default 0 kb/s

filter

Syntax **filter ip** *ip-filter-id*
no filter ip [*ip-filter-id*]
filter ipv6 *ipv6-filter-id*
no filter ipv6 [*ipv6-filter-id*]
filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
no filter [[**ip** [*ip-filter-id*]] [[**ipv6** [*ipv6-filter-id*]]

Context config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress

Description This command associates an IPv4 or IPv6 filter policy with an ingress or egress 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* or *ipv6-filter-id* with an ingress or egress SAP. The *ip-filter-id* or *ipv6-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 will be returned.

Only one filter ID can be assigned to an interface unless the interface is dual-stack (supports both IPv4 and IPv6). A dual-stack interface can have one IPv4 and one IPv6 filter ID assigned to it.

In general, filters applied to SAPs apply to all packets on the SAP. One exception is that IP match criteria are not applied to non-IP packets, in which case 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 is not removed from the system unless the scope of the created filter is set to **local**. To avoid deletion of the filter ID and only break the association with the service object, use the **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Parameters *ip-filter-id* — the IPv4 filter policy. The filter ID or filter name must already exist within the created IPv4 filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

ipv6-filter-id — the IPv6 filter policy. The filter ID or filter name must already exist within the created IPv6 filters.

Values 1 to 65535 or *ipv6-filter-name* (up to 64 characters)

match-qinq-dot1p

Syntax **match-qinq-dot1p {top | bottom}**
no match-qinq-dot1p

Context config>service>vprn>if>sap>ingress

Description This command specifies which dot1q tag position (top or bottom) in a qinq-encapsulated packet should be used when QoS evaluates dot1p classification.

The **no** form of the command restores the default dot1p evaluation behavior for the SAP, which means that the inner (bottom) tag (second tag) dot1p bits are used for classification.

By default, the dot1p bits from the inner tag service-delineating dot1q tag are used.

Table 110 shows which set of dot1p bits are used for QoS purposes when **match-qinq-dot1p** is configured. To use the table, find the row that represents the settings for **Port/SAP Type** and **Match-QinQ-Dot1q Setting**. Use the **Existing Packet Tags** column to identify which dot1q tags are available in the packet. Then use the **P-bits Used for Match** column to identify which dot1q tag contains the dot1p bits that are used for QoS dot1p classification.

Default no match-qinq-dot1p

Parameters **top** — the **top** parameter and **bottom** parameter are mutually exclusive. When the **top** parameter is specified, the outer tag's dot1p bits (topmost P-bits) are used (if existing) to match any **dot1p** *dot1p-value* entries

bottom — the **bottom** parameter and **top** parameter are mutually exclusive. When the **bottom** parameter is specified, the bottommost P-bits (second tag's P-bits) are used (if existing) to match any **dot1p** *dot1p-value* entries.

Table 110: Match-QinQ-Dot1p Matching Behavior

Port/ SAP Type	Match-QinQ-Dot1p Setting [1]	Existing Packet Tags	P-bits Used for Match
Null	n/a	None	None
Null	n/a	Dot1p (VLAN ID 0)	None [2]
Null	n/a	Dot1q	None [2]
Null	n/a	TopQ BottomQ	None [2]
Dot1Q	n/a	None	None

Table 110: Match-QinQ-Dot1p Matching Behavior (Continued)

Port/ SAP Type	Match-QinQ-Dot1p Setting [1]	Existing Packet Tags	P-bits Used for Match
Dot1Q	n/a	Dot1p (default SAP VLAN ID 0)	Dot1p P-bits
Dot1Q	n/a	Dot1q	Dot1q P-bits
QinQ/ X.Y	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.Y	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ X.0	Top	TopQ	TopQ P-bits
QinQ/ X.0	Default or Bottom	TopQ	TopQ P-bits
QinQ/ X.0	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.0	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ X.*	Top	TopQ	TopQ P-bits
QinQ/ X.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ X.*	Top	TopQ BottomQ	TopQ P-bits
QinQ/ X.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ 0.*	Top	None	None
QinQ/ 0.*	Default or Bottom	None	None
QinQ/ 0.*	Top	TopQ	TopQ P-bits
QinQ/ 0.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ 0.*	Top	TopQ BottomQ	TopQ P-bits
QinQ/ 0.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits
QinQ/ *.*	Top	None	None
QinQ/ *.*	Default or Bottom	None	None
QinQ/ *.*	Top	TopQ	TopQ P-bits
QinQ/ *.*	Default or Bottom	TopQ	TopQ P-bits
QinQ/ *.*	Top	TopQ BottomQ	TopQ P-bits
QinQ/ *.*	Default or Bottom	TopQ BottomQ	BottomQ P-bits

Notes:

1. "Default" in this column refers to the **no** form of the **match-qinq-dot1p** command.
2. For null encapsulation, the 7705 SAR does not process dot1p bits.

qinq-mark-top-only

Syntax **[no] qinq-mark-top-only**

Context config>service>vprn>if>sap>egress

Description When enabled, the **qinq-mark-top-only** command specifies which P-bits to mark during packet egress. When disabled, both sets of P-bits are marked. When enabled, only the P-bits in the top Q-tag are marked. The **no** form of the command is the default state (disabled).

Table 111 shows the dot1p remarking behavior for different egress port type/SAP type combinations and **qinq-mark-top-only** state, where "False" represents the default (disabled) state.

If a new tag is pushed, the dot1p bits of the new tag will be zero (unless the new tag is re-marked by the egress policy. The dot1p bits are configured using the **dot1p** parameter under the **config**>**qos** context.

Table 111: Dot1P Re-marking Behavior for the QinQ-mark-top-only Command

Egress Port Type/SAP Type	QinQ-mark-top-only State	Egress P-Bits Marked or Re-marked
Null [1]	n/a	None
Dot1q/ X [1]	n/a	Outer tag
Dot1q/ * [2]	n/a	None
QinQ/ X.Y [1]	False	Two outer tags [3]
	True	Outer tag [3]
QinQ/ X.* [4]	True or False	Outer tag
QinQ/ X.0 [1]	True or False	Outer tag
QinQ/ 0.* [1]	True or False	None
QinQ/ *.* [2]	True or False	None

Notes:

1. This port type/SAP type is supported by the following services: Epipe, Ipipe, VPLS, IES, and VPRN.
2. This port type/SAP type is supported by the following services: Epipe and VPLS.
3. Normally, when a new tag is pushed, the dot1p bits of the new tag will be zero, unless the P-bits are remarked by the egress policy. However, an exception to this occurs when the egress SAP type is X.Y and only one new outer tag must be pushed. In this case, the new outer tag will have its dot1p bits set to the inner tag's dot1p bits.
4. This port type/SAP type is supported by the following services: Epipe, Ipipe, and VPLS.

Default no qinq-mark-top-only

qos

Syntax **qos** *policy-id*
no qos [*policy-id*]

Context config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress

Description This command associates a QoS policy with an ingress or egress SAP. QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the *policy-id* does not exist, an error will be returned.

The **qos** command is used to associate both ingress and egress QoS policies. The **qos** command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error; for example, trying to associate an egress policy on SAP ingress.

Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a given type returns an error.

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

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

Parameters *policy-id* — the ingress or egress policy ID to associate with the SAP on ingress or egress. The policy ID or name must already exist.

Values 1 to 65535, or *policy-name* (up to 64 characters)

scheduler-mode

Syntax **scheduler-mode** {**4-priority** | **16-priority**}

Context config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress

Description This command sets the scheduler mode for the SAP and is part of the hierarchical QoS (H-QoS) feature on the 7705 SAR.

If the mode is 4-priority, then the SAP is considered an unshaped 4-priority SAP and the agg-rate-limit cannot be changed from its default values.

If the mode is 16-priority and the **agg-rate limit** parameters are configured to be non-default values, then the SAP is considered a shaped SAP. If the **agg-rate limit** parameters are left in their default settings, the SAP is considered an unshaped, 16-priority SAP.

This command is blocked on third-generation (Gen-3) Ethernet adapter cards and platforms, such as the 6-port Ethernet 10Gbps Adapter card and the 7705 SAR-X with Ethernet ports, which use only 4-priority-hqos scheduling mode. It is also blocked on the 7705 SAR-X with a TDM MDA, which uses only 4-priority scheduling mode.

If the specified SAP is a LAG SAP, **scheduler-mode** can be configured but is not applied to Gen-3 adapter cards and platforms. If one of the ports in the LAG is on a Gen-1 adapter card, then **scheduler-mode** cannot be configured.

Default 4-priority

Parameters **4-priority —** sets the scheduler mode for the SAP to be 4-priority mode

16-priority — sets the scheduler mode for the SAP to be 16-priority mode

shaper-group

Syntax [**no**] **shaper-group** *shaper-group-name*

Context config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress

Description This command applies a shaper group to a SAP. The shaper group must already be created and must be within the shaper policy assigned to the Ethernet MDA (for ingress) or port (for egress). A shaper group is a dual-rate aggregate shaper used to shape aggregate access ingress or egress SAPs at a shaper group rate. Multiple aggregate shaper groups ensure fair sharing of available bandwidth among different aggregate shapers.

The default shaper group cannot be deleted.

The **no** form of this command removes the configured **shaper-group**.

Default shaper-group “default”

Parameters *shaper-group-name* — the name of the shaper group. To access the default shaper group, enter “default”.

create — keyword used to create a shaper group

Interface Spoke SDP Commands

spoke-sdp

Syntax **spoke-sdp** *sdp-id*:*vc-id* [**create**]
no spoke-sdp *sdp-id*:*vc-id*

Context config>service>vprn>interface

Description This command binds a service to an existing Service Distribution Point (SDP).

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

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 is down.

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

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

Class-based forwarding is not supported on a spoke SDP used for termination on an IES or VPRN service. All packets are forwarded over the default LSP.

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 the service. Once the binding is removed, no packets are forwarded to the far-end router. The spoke SDP must be shut down first before it can be deleted from the configuration.

Default n/a

Special Cases **VPRN —** several SDPs can be bound to a VPRN service. Each SDP must be destined for a different 7705 SAR router. If two *sdp-id* bindings terminate on the same 7705 SAR, an error occurs and the second SDP is binding is rejected.

Parameters *sdp-id* — the SDP identifier

Values 1 to 17407

vc-id — the virtual circuit identifier

Values 1 to 4294967295

egress

Syntax **egress**

Context config>service>vprn>if>spoke-sdp

Description This command enables the context to configure egress SDP parameters.

vc-label

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

Context config>service>vprn>if>spoke-sdp>egress

Description This command configures the static MPLS VC label used by the 7705 SAR to send packets to the far-end device in this service via this SDP.

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

Values 16 to 1048575

ingress

Syntax **ingress**

Context config>service>vprn>if>spoke-sdp

Description This command enables the context to configure ingress SDP parameters.

filter

Syntax **filter ip** *ip-filter-id*
filter ipv6 *ipv6-filter-id*
no filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]

Context config>service>vprn>if>spoke-sdp>ingress

Description This command associates an IPv4 or IPv6 filter policy with a spoke SDP. Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria.

The *ip-filter-id* or *ipv6-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 will be returned.

Only one filter ID can be assigned to an interface unless the interface is dual-stack (supports both IPv4 and IPv6). A dual-stack interface can have one IPv4 and one IPv6 filter ID assigned to it.

In general, filters applied to ingress spoke SDPs will apply to all packets on the spoke SDP. One exception is that 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 spoke SDP.

Parameters *ip-filter-id* — the IP filter policy. The filter ID or filter name must already exist within the created IP filters.

Values 1 to 65535 or IPv4 *filter-name* (up to 64 characters)

ipv6-filter-id — specifies the IPv6 filter policy. The filter ID or filter name must already exist within the created IPv6 filters.

Values 1 to 65535 or *ipv6-filter-name* (up to 64 characters)

vc-label

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

Context config>service>vprn>if>spoke-sdp>ingress

Description This command configures the static MPLS VC label used by the far-end device to send packets to the 7705 SAR in this service via this SDP.

Parameters *ingress-vc-label* — a VC ingress value that indicates a specific connection

Values 2048 to 18431

Routed VPLS Commands

vpls

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

Context config>service>vprn>interface

Description This command within the IP interface context binds the IP interface to the specified VPLS service name.

The system does not attempt to resolve the service name 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 service name. If the IP interface is already in the administratively up state, the system immediately attempts to resolve the given service name.

Parameters *service-name* — specifies the service name that the system attempts to resolve to an allow-ip-int-binding enabled VPLS service associated with the service name. The specified service name is an ASCII string of up to 32 characters.

ingress

Syntax **ingress**

Context config>service>vprn>if>vpls

Description This command within the VPLS binding context defines the routed IPv4 optional filter override.

v4-routed-override-filter

Syntax **v4-routed-override-filter** *ipv4-filter-id*
no v4-routed-override-filter

Context config>service>vprn>if>vpls>ingress

Description This command specifies an IPv4 filter ID applied to all ingress packets entering the VPLS service. The filter overrides the existing ingress IPv4 filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional, and if not defined or removed, the IPv4 routed packets use the existing ingress IPv4 filter on the VPLS virtual ports.

The **no** form of the command removes the IPv4 routed override filter from the ingress IP interface.

Default n/a

Parameters *ipv4-filter-id* — specifies the IPv4 filter policy. The filter ID or filter name must already exist within the created IP filters.

Values 1 to 65535 or *filter-name* (up to 64 characters)

v6-routed-override-filter

Syntax **v6-routed-override-filter** *ipv6-filter-id*
no v6-routed-override-filter

Context config>service>vprn>if>vpls>ingress

Description This command specifies an IPv6 filter ID applied to all ingress packets entering the VPLS service. The filter overrides the existing ingress IPv6 filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional, and if not defined or removed, the IPv6 routed packets use the existing ingress IPv6 filter on the VPLS virtual ports.

The **no** form of the command removes the IPv6 routed override filter from the ingress IP interface.

Default n/a

Parameters *ipv6-filter-id* — the IPv6 filter policy. The filter ID or filter name must already exist within the created IP filters.

Values 1 to 65535 or *ipv6-filter-name* (up to 64 characters)

Interface VRRP Commands

vrrp

Syntax **vrrp** *virtual-router-id* [**owner**]
no vrrp *virtual-router-id*

Context config>service>vprn>interface

Description This command creates or edits a virtual router ID on the service IP interface. A virtual router ID is internally represented in conjunction with the IP interface name. This allows the virtual router ID to be used on multiple IP interfaces while representing different virtual router instances.

Two VRIDs can be defined on an IP interface. One, both, or none may be defined as owner.

The **no** form of this command removes the specified virtual router ID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the virtual router ID. The virtual router ID does not need to be shutdown in order to remove the virtual router instance.

Default n/a

Parameters *virtual-router-id* — specifies a new virtual router ID or one that can be modified on the IP interface

Values 1 to 255

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>service>vprn>if>vrrp

Description This command assigns a simple text password authentication key to generate master VRRP advertisement messages and validate received VRRP advertisement messages.

If the command is re-executed with a different password key defined, the new key will be used immediately. If a **no authentication-key** command is executed, the password authentication key is restored to the default value. The **authentication-key** command can be executed at any time.

To change the current in-use password key on multiple virtual router instances:

- identify the current master
- shut down the virtual router instance on all backups
- execute the **authentication-key** command on the master to change the password key
- execute the **authentication-key** command and **no shutdown** command on each backup

The **no** form of this command restores the default value of the key.

Default The authentication data field contains the value 0 in all 16 octets.

Parameters *authentication-key* — identifies the simple text password used when VRRP Authentication Type 1 is enabled on the virtual router instance. Type 1 uses a string eight octets long that is inserted into all transmitted VRRP advertisement messages and compared against all received VRRP advertisement messages. The authentication data fields are used to transmit the key.

The *authentication-key* parameter is expressed as a string consisting up to eight alpha-numeric characters. Spaces must be contained in quotation marks (" "). The quotation marks are not considered part of the string.

The string is case-sensitive and is left-justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet containing the first character. The second field holds the fifth through eighth characters. Any unspecified portion of the authentication data field is padded with the value 0 in the corresponding octet.

Values any 7-bit printable ASCII character

exceptions:	double quote	ASCII 34
	carriage return	ASCII 13
	line feed	ASCII 10
	tab	ASCII 9
	backspace	ASCII 8

hash-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 option 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** keyword 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** keyword specified.

hash2 — specifies the key is entered in a more complex encrypted form. If the **hash2** keyword is not used, the less-encrypted hash form is assumed.

backup

Syntax [**no**] **backup** *ip-address*

Context config>service>vprn>if>vrrp

Description This command configures virtual router IP addresses for backup.

Parameters *ip-address* — specifies the destination address for backup

bfd-enable

Syntax [**no**] **bfd-enable** [**base** | *service-id*] **interface** *interface-name* **dst-ip** *ip-address*

Context config>service>vprn>if>vrrp

Description This commands assigns a BFD session that provides a heart-beat mechanism for the given VRRP instance. Only one BFD session can be assigned to any given VRRP instance, but multiple VRRP sessions can use the same BFD session.

BFD controls the state of the associated 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 session are set with the **bfd-enable** command under the IP interface.

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

Default n/a

Parameters *service-id* — specifies the service ID or name of the interface running BFD

Values 1 to 65535 or *service-name*

interface-name — specifies the name of the interface running BFD

ip-address — specifies the destination address to be used for the BFD session

init-delay

Syntax **init-delay** *seconds*
no init-delay

Context config>service>vprn>if>vrrp

Description This command configures a VRRP initialization delay timer.

Default no init-delay

Parameters *seconds* — specifies the amount of time in seconds for the initialization delay timer for VRRP

Values 1 to 65535

mac

Syntax **mac** *ieee-address*
no mac

Context config>service>vprn>if>vrrp

Description This command assigns a specific MAC address to a VPRN IP interface.

The **no** form of the 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 (the default MAC address assigned to the interface, assigned by the system)

Parameters *ieee-address* — specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bbcc-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.

master-int-inherit

Syntax [**no**] **master-int-inherit**

Context config>service>vprn>if>vrrp

Description This command allows the master instance to dictate the master down timer (non-owner context only).

The master down interval is the time that the master router can be down before backup takes over. The master down interval is used to specify the master down timer. If the master down timer expires, the backup virtual router enters the master state.

Default no master-int-inherit

message-interval

Syntax **message-interval** [{*seconds*] [**milliseconds** *milliseconds*]}
no message-interval

Context config>service>vprn>if>vrrp

Description This command sets the advertisement timer and indirectly sets the master down timer on the virtual router instance. The message-interval setting must be the same for all virtual routers with the same VRID. Any VRRP advertisement message received with an advertisement interval field different from the virtual router instance configured message-interval value will be silently discarded.

The **message-interval** command is available for both non-owner and owner virtual routers. If the **message-interval** command is not executed, the default message interval is 1 s.

The **no** form of this command restores the default message-interval value of 1 s to the virtual router instance.

Parameters *seconds* — specifies the number of seconds that will transpire before the advertisement timer expires

Values 1 to 255

Default 1

milliseconds — specifies the time interval, in milliseconds, between sending advertisement messages. This parameter is not supported on non-redundant chassis.

Values 100 to 900

ping-reply

Syntax [**no**] **ping-reply**

Context config>service>vprn>if>vrrp

Description This command enables the non-owner master to reply to ICMP echo requests directed to the virtual router instance IP addresses. The ping request can be received on any routed interface.

Ping must not have been disabled at the management security level (either on the parental IP interface or based on the ping source host address). When ping reply is not enabled, ICMP echo requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP echo requests regardless of the setting of the ping reply configuration.

The **ping-reply** command is only available for non-owner virtual routers. If the **ping-reply** command is not executed, ICMP echo requests to the virtual router instance IP addresses are silently discarded.

The **no** form of this command restores the default operation of discarding all ICMP echo request messages destined for the non-owner virtual router instance IP addresses.

Default no ping-reply

policy

Syntax **policy** *vrrp-policy-id*
no policy

Context config>service>vprn>if>vrrp

Description This command associates a VRRP priority control policy with the virtual router instance (non-owner context only).

Default n/a

Parameters *vrrp-policy-id* — specifies a VRRP priority control policy. The VRRP policy ID must already exist in the system for the **policy** command to be successful

Values 1 to 9999

preempt

Syntax **preempt**
no preempt

Context config>service>vprn>if>vrrp

Description This command provides the ability to override an existing non-owner master with a virtual router backup that has a higher priority. Enabling preempt mode enhances the operation of the base priority and VRRP policy ID definitions on the virtual router instance. If the virtual router cannot preempt an existing non-owner master, the affect of the dynamic changing of the in-use priority is greatly diminished.

The **preempt** command is only available for non-owner VRRP virtual routers. The owner cannot be preempted because the priority of non-owners can never be higher than the owner. The owner will always preempt all other virtual routers when it is available.

Non-owner backup virtual router instances will only preempt when preempt is set and the current master has an in-use message priority value less than the backup virtual router instance in-use priority.

A master non-owner virtual router will only allow itself to be preempted when the incoming VRRP advertisement message priority field value is one of the following:

- greater than its in-use priority value
- equal to the in-use priority value, and the source IP address (primary IP address) is greater than its primary IP address

The **no** form of this command prevents a non-owner virtual router instance from preempting another, less-desirable, virtual router.

Default preempt

priority

Syntax **priority** *priority*
no priority

Context config>service>vprn>if>vrrp

Description This command configures a specific priority value for the virtual router instance. In conjunction with the optional **policy** command, the base priority derives the in-use priority of the virtual router instance.

The **priority** command is only available for non-owner VRRP virtual routers. The priority of owner virtual router instances is permanently set to 255 and cannot be changed. For non-owner virtual router instances, if the priority command is not executed, the base priority is set to 100.

The **no** form of this command restores the default value of 100.

Parameters *priority* — specifies the base priority used by the virtual router instance. If a VRRP priority control policy is not defined, the base priority will be the in-use priority for the virtual router instance.

Values 1 to 254

ssh-reply

Syntax [**no**] **ssh-reply**

Context config>service>vprn>if>vrrp

Description This command enables the non-owner master to reply to SSH requests directed at the IP addresses of the virtual router instances. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parental IP interface or based on the SSH source host address). Proper login and CLI command authentication are enforced.

When the **ssh-reply** command is not enabled, SSH packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to SSH requests regardless of the SSH reply configuration.

The **ssh-reply** command is only available for non-owner VRRP virtual routers.

The **no** form of this command restores the default operation of discarding all SSH packets destined to the non-owner virtual router instance IP addresses.

Default no ssh-reply

standby-forwarding

Syntax [**no**] **standby-forwarding**

Context config>service>vprn>if>vrrp

Description This command allows the forwarding of packets by a standby router when sent to the virtual router MAC address.

The **no** form of the command specifies that a standby router should not forward traffic sent to the virtual router MAC address. The standby router should forward traffic sent to the real MAC address of the standby router.

Default no standby-forwarding

telnet-reply

Syntax [**no**] **telnet-reply**

Context config>service>vprn>if>vrrp

Description This command enables the non-owner master to reply to TCP port 23 Telnet requests directed at the IP addresses of the virtual router instance. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Proper login and CLI command authentication are enforced.

If the **telnet-reply** command is not enabled, TCP port 23 Telnet packets to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to Telnet requests regardless of the Telnet reply configuration.

The **telnet-reply** command is only available for non-owner VRRP virtual routers.

The **no** form of this command restores the default operation of discarding all Telnet packets destined for the non-owner virtual router instance IP addresses.

Default no telnet-reply

traceroute-reply

Syntax [**no**] **traceroute-reply**

Context config>service>vprn>if>vrrp

Description This command enables a non-owner master to reply to traceroute requests directed to the virtual router instance IP addresses. The command is valid only if the VRRP virtual router instance associated with this entry is a non-owner. A non-owner backup virtual router never responds to traceroute requests regardless of the traceroute reply status.

Default no traceroute-reply

TWAMP Light Commands

twamp-light

Syntax **twamp-light**

Context config>service>vprn

Description This command enables the context for configuring TWAMP Light functionality.

Default disabled

reflector

Syntax **reflector**

Context config>service>vprn>twamp-light

Description This command enables the context for configuring TWAMP Light session reflector functionality. The relector functionality is enabled using the **no shutdown** command.

Default disabled

prefix

Syntax [**no**] **prefix** *ip-prefix/prefix-length* [**create**]

Context config>service>vprn>twamp-light>reflector

Description This command configures an IP address prefix containing one or more TWAMP Light session controllers. It is used to define which TWAMP Light packet prefixes the reflector will process. Once the prefix is configured, the TWAMP Light session reflector only responds to TWAMP Light packets from source addresses that are part of the prefix list.

Default no prefix

Parameters *ip-prefix/ip-prefix-length* — the IPv4 or IPv6 address prefix

Values	*ipv4-prefix*	a.b.c.d (host bits must be 0)
	ipv4-prefix-length	0 to 32
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 to FFFF]H d: [0 to 255]D
	ipv6-prefix-length	{0 to 128} | {0 to 64 | 128}

udp-port

Syntax **udp-port** *number*
no udp-port

Context config>service>vprn>twamp-light>reflector

Description This command configures the specific UDP port that the session reflector listens to for TWAMP Light packets. The session controller launching the TWAMP Light packets must have the same UDP port configured as on the session reflector.

Default no udp-port

Parameters *number* — the UDP port number

Values 1024 to 65535

Show Service Commands

Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

egress-label

Syntax **egress-label** *start-label* [*end-label*]

Context show>service

Description This command displays service information 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 labels in the specified range are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *end-label* — the ending egress label value for which to display services using the label range

Values 2049 to 131071

Default the *start-label* value

start-label — the starting egress 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, or 2048 to 131071

Output The following output is an example of service egress label information, and Table 112 describes the fields.

Sample Output

In the example below, services 3, 5 and 6 are IES, and services 5000 and 5001 are VPLS services.

```
*A:ALU-12>show>service# egress-label 0 131071
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Binding        Type  I.Lbl                E.Lbl
-------------------------------------------------------------------------------
3          15:15              Spok  0                    0
5          5:5                Spok  0                    0
6          5:6                Spok  0                    0
5000       15:5000            Mesh  0                    0
5000       15:5001            Spok  0                    0
5001       5001:100           Spok  0                    0
-------------------------------------------------------------------------------
Number of Bindings Found : 6
-------------------------------------------------------------------------------
===============================================================================
```

Table 112: Show Service Egress Output Fields

Label	Description
Svc Id	The ID that identifies a service
Sdp Binding	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 label range

id

Syntax **id** *service-id*

Context show>service

Description This command accesses the context to display information for a particular service ID. Once the particular service context has been accessed, the filtering commands listed below are available.

Parameters *service-id* — the unique service identification number or name that identifies the service in the service domain

all — detailed information about the service

arp — ARP entries for the service

base — basic service information

dhcp — DHCP entries for the service

endpoint — service endpoint information

fdb — FDB entries for the service

interface — service interfaces

labels — labels being used by this service

mac-move — MAC move related information about this service

sap — SAPs associated to the service

sdp — SDPs associated with the service

split-horizon — service split horizon groups

static-host — static hosts configured on this service

all

Syntax **all**

Context show>service>id

Description This command displays detailed information for all aspects of the service.

Output The following output is an example of service ID all information, and Table 113 describes the fields.

Sample Output

```
*A:ALU-2#show>service>id# all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 6
Service Type      : VPRN
Name              : VPRN_global
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 02/03/2010 21:59:01
Last Mgmt Change  : 03/01/2010 18:55:58
Admin State       : Down                 Oper State        : Down
Route Dist.       : 12.34.56.78:655      VPRN Type         : regular
AS Number         : None                 Router Id         : 255.0.0.0
ECMP              : Enabled              ECMP Max Routes   : 1
Max IPv4 Routes   : No Limit             Auto Bind         : None
Max IPv6 Routes   : No Limit
Vrf Target        : None
Vrf Import        : None
Vrf Export        : None
SAP Count         : 1                    SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 1:6  -(10.10.10.10)
-------------------------------------------------------------------------------
SDP Id            : 1:6                      Type              : Spoke
VC Type           : n/a                      VC Tag            : n/a
Admin Path MTU    : 0                        Oper Path MTU     : 0
Far End           : 10.10.10.10              Delivery          : MPLS
Admin State       : Up                       Oper State        : Down
Acct. Pol         : None                     Collect Stats     : Disabled
Ingress Label     : n/a                      Egress Label      : n/a
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 : 02/03/2010 21:59:01     Signaling         : n/a
Last Mgmt Change  : 03/01/2010 18:55:58
Class Fwding State : Down
Flags             : SdpOperDown
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. Dro. Pkts.     : n/a
I. Fwd. Octs.      : n/a                      I. Dro. Octs.     : n/a
E. Fwd. Pkts.      : n/a                      E. Fwd. Octets    : n/a
Associated LSP LIST :
No LSPs Associated

-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
SAP 1/5/2
-------------------------------------------------------------------------------
Service Id         : 6
SAP                : 1/5/2                    Encap             : null
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Multi Svc Site     : None
Last Status Change : 02/03/2010 21:59:01
Last Mgmt Change   : 02/03/2010 21:59:02
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100

Admin MTU          : 1514                     Oper MTU          : 1514
Ingr IP Fltr-Id    : 2                        Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                     qinq-pbit-marking : both
Egr Agg Rate Limit : max
Q Frame-Based Acct : Disabled

Acct. Pol          : None                     Collect Stats     : Disabled

Anti Spoofing      : None                     Nbr Static Hosts  : 0

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 2                        Egress qos-policy : 2
Shared Q plcy      : n/a                      Multipoint shared : Disabled
-------------------------------------------------------------------------------
TWAMP Light Reflector
-------------------------------------------------------------------------------
Admin State          : Up
Up Time              : 0d 00:12:01
Configured UDP Port  : 65535
Test Packets Rx      : 0                  Test Packets Tx     : 0

TWAMP Light Controller Prefix List
 192.168.1.1/32
    10.1.1.2/32
 172.16.254.9/3
    10.1.1.0/24
```

```
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0

Queueing Stats(Ingress QoS Policy 2)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 2)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
-------------------------------------------------------------------------------
Service Interfaces
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Interface
-------------------------------------------------------------------------------
If Name           : vprn_interface
Admin State       : Up                  Oper (v4/v6)      : Down/Down
Protocols         : None

IP Addr/mask      : Not Assigned
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
If Index          : 2                   Virt. If Index    : 2
Last Oper Chg     : 02/03/2010 21:59:02 Global If Index   : 125
SAP Id            : 1/5/2
TOS Marking       : Trusted             If Type           : VPRN
SNTP B.Cast       : False
MAC Address       :                     Arp Timeout       : 14400
```

```
IP MTU            : 1500                ICMP Mask Reply   : True
Arp Populate      : Disabled            Host Conn Verify  : Disabled
LdpSyncTimer      : None

Proxy ARP Details
Rem Proxy ARP     : Disabled            Local Proxy ARP   : Disabled
Policies          : none

Proxy Neighbor Discovery Details
Local Pxy ND      : Disabled
Policies          : none

DHCP Details
Admin State       : Down                Lease Populate    : 0
Action            : Keep                Trusted           : Disabled
DHCP6 Relay Details
Admin State       : Down                Lease Populate    : 0
Oper State        : Down                Nbr Resolution    : Disabled
If-Id Option      : None                Remote Id         : Disabled
Src Addr          : Not configured

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

IPCP Address Extension Details
Peer IP Addr      : Not configured
Peer Pri DNS Addr : Not configured
Peer Sec DNS Addr : Not configured
===============================================================================
*A:ALU-2#>show>service>id#
```

Table 113: Show Service ID All Output Fields

Label	Description
Service Detailed Information	
Service Id	The service identifier
Service Type	The type of service (VPRN)
Name	The service name
Description	A description of the service
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 service

Table 113: Show Service ID All Output Fields (Continued)

Label	Description
Admin State	The desired state of the service
Oper State	The current operational state of the service
Route Dist.	The route distribution number
VPRN Type	Only valid in services that accept mesh SDP bindings. It validates the VC ID portion of each mesh SDP binding defined in the service.
AS Number	The autonomous system number
Router Id	The router ID for this service
ECMP	Displays equal cost multipath information
ECMP Max Routes	The maximum number of routes that can be received from the neighbors in the group or for the specific neighbor
Max IPv4 Routes	The maximum number of routes that can be used for path sharing
Max IPv6 Routes	Not applicable
Auto Bind	The automatic binding type for the SDP assigned to this service
Vrf Target	The route target in the VRF applied to this service
Vrf Import	The VRF import policy applied to this service
Vrf Export	The VRF export policy applied to this service
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
VC Type	The VC type: ether or vlan
VC Tag	The explicit dot1Q value used when encapsulating to the SDP far end
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

Table 113: Show Service ID All Output Fields (Continued)

Label	Description
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented
Far End	Specifies the IP address of the remote end of the GRE, MPLS, or IP tunnel defined by this SDP
Delivery	Specifies the type of delivery used by the SDP: GRE, MPLS, or IP
Admin State	The administrative state of this SDP
Oper State	The operational state of this SDP
Acct. Pol	The accounting policy applied to the SDP
Collect Stats	Specifies whether accounting statistics are collected on the SDP
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP
Ing mac Fltr	Not applicable
Egr mac Fltr	Not applicable
Ing ip Fltr	The SDP ingress filter policy ID
Egr ip Fltr	The SDP egress filter policy ID
Ing ipv6 Fltr	Not applicable
Egr ipv6 Fltr	Not applicable
Admin ControlWord	The administrative state of the control word: Preferred (control word enabled) or Not Preferred (control word disabled)
Oper ControlWord	The operational state of the control word: True (control word enabled) or False (control word disabled)
Last Status Change	The date and time of the most recent status change to this SDP
Signaling	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP
Last Mgmt Change	The date and time of the most recent management-initiated change to this SDP

Table 113: Show Service ID All Output Fields (Continued)

Label	Description
Class Fwding State	Not applicable
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
KeepAlive Information	
Admin State	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	The length of the SDP echo request messages transmitted on this SDP
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 time to wait before the keepalive operating status is eligible to enter the alive state
Statistics	
I. Fwd. Pkts.	The number of forwarded ingress packets
I. Dro. Pkts.	The number of dropped ingress packets
I. Fwd. Octs.	The number of forwarded ingress octets
I. Dro. Octs.	The number of dropped ingress octets
E. Fwd. Pkts.	The number of forwarded egress packets
E. Fwd. Octets	The number of forwarded egress octets

Table 113: Show Service ID All Output Fields (Continued)

Label	Description
Associated LSP LIST	If the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far-end field. If the SDP type is GRE, the following message displays: SDP delivery mechanism is not MPLS
Number of SDPs	The total number of SDPs applied to this service ID
Service Access Points	
Service Id	The service identifier
SAP	The SAP identifier
Encap	The encapsulation type of the SAP
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
Multi Svc Site	Indicates the multiservice site that the SAP is a member of
Last Status Change	The time of the most recent operating status change to this SAP
Last Mgmt Change	The time of the most recent management-initiated change to this SAP
Sub Type	The supported sub type: regular
Dot1Q Ethertype	The value of the dot1q Ethertype
QinQ Ethertype	The value of the qinq Ethertype
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

Table 113: Show Service ID All Output Fields (Continued)

Label	Description
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
Ingr IP Fltr-Id	The ingress filter policy ID assigned to the SAP
Egr IP Fltr-Id	The egress filter policy ID assigned to the SAP
Ingr Mac Fltr-Id	Not applicable
Ingr IPv6 Fltr-ID	Not applicable
Egr IPv6 Fltr-ID	Not applicable
tod-suite	Indicates whether a time-based policy is applied to a multiservice site
qinq-pbit-marking	Indicates the qinq P-bit marking for the service: both or top
Egr Agg Rate Limit	Not applicable
Q Frame-Based Acct	Not applicable
Acct. Pol	The accounting policy applied to the SAP
Collect Stats	Specifies whether accounting statistics are collected on the SAP
Anti Spoofing	Not applicable
Nbr Static Hosts	Not applicable
QOS	
Ingress qos-policy	The SAP ingress QoS policy ID
Egress qos-policy	The SAP egress QoS policy ID
Shared Q plcy	Not applicable
Multipoint shared	Not applicable

Table 113: Show Service ID All Output Fields (Continued)

Label	Description
TWAMP-Light Reflector	
Admin State	Displays one of the following: Up—the server or prefix is administratively enabled (no shutdown) in configuration Down—the server or prefix is administratively disabled (shutdown) in configuration
Up Time	The time since the server process was started, measured in days (d), hours, minutes, and seconds
Configured UDP Port	The UDP port number used
Test Packets Rx	The total number of test packets received from session senders
Test Packets Tx	The total number of test packets sent to session senders
TWAMP Light Controller Prefix List	The IP address prefixes of TWAMP Light clients
SAP Statistics	
Last Cleared time	The date and time that a clear command was issued on the statistics
Forwarding Engine Stats	
Dropped	The number of packets or octets dropped by the forwarding engine
Off. HiPrio	The number of high-priority packets or octets offered to the forwarding engine
Off. LowPrio	The number of low-priority packets offered to the forwarding engine
Queueing Stats (Ingress QoS Policy)	
Dro. HiPrio	The number of high-priority packets or octets discarded, as determined by the SAP ingress QoS policy
Dro. LowPrio	The number of low-priority packets discarded, as determined by the SAP ingress QoS policy
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP ingress QoS policy
For. OutProf	The number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP ingress QoS policy

Table 113: Show Service ID All Output Fields (Continued)

Label	Description
Queueing Stats (Egress QoS Policy)	
Dro. InProf	The number of in-profile packets or octets discarded, as determined by the SAP egress QoS policy
Dro. OutProf	The number of out-of-profile packets or octets discarded, as determined by the SAP egress QoS policy
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP egress QoS policy
For. OutProf	The number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP egress QoS policy
Sap per Queue stats	
Ingress Queue *n*	The index of the ingress QoS queue of this SAP, where *n* is the index number
Off. HiPrio	The number of packets or octets of high-priority traffic for the SAP (offered)
Off. LoPrio	The number of packets or octets count of low-priority traffic for the SAP (offered)
Dro. HiPrio	The number of high-priority traffic packets or octets dropped
Dro. LoPrio	The number of low-priority traffic packets or octets dropped
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	The number of out-of-profile packets or octets (rate above CIR) forwarded
Egress Queue *n*	The index of the egress QoS queue of the SAP, where *n* is the index number
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	Number of out-of-profile packets or octets (rate above CIR) forwarded
Dro. InProf	The number of in-profile packets or octets dropped for the SAP
Dro. OutProf	The number of out-of-profile packets or octets discarded

Table 113: Show Service ID All Output Fields (Continued)

Label	Description
Service Interfaces	
Interface	
If Name	The name used to refer to the interface
Admin State	The desired state of the interface
Oper (v4/v6)	The operating state of the interface
Protocols	The protocols supported on the interface
IP Addr/mask	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.
Virt. If Index	The virtual interface index of the VPRN interface
Last Oper Chg	The date and time of the last operating state change on the interface
Global If Index	The global interface index of the VPRN interface
TOS Marking	Specifies whether the ToS marking is trusted or untrusted for the interface
If Type	The interface type
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled
MAC Address	The 48-bit IEEE 802.3 MAC address
Arp Timeout	The timeout for an ARP entry learned on the interface
IP MTU	The IP maximum transmit unit for the interface
ICMP Mask Reply	Specifies whether ICMP mask reply is enabled or disabled
ARP Populate	Specifies if ARP is enabled or disabled
Host Conn Verify	Not applicable
LdpSyncTimer	Not applicable
Proxy ARP Details	
Rem Proxy ARP	Indicates whether remote proxy ARP is enabled or disabled

Table 113: Show Service ID All Output Fields (Continued)

Label	Description
Local Proxy ARP	Indicates whether local proxy ARP is enabled or disabled
Policies	Specifies the policy statement(s) applied to proxy ARP
DHCP Details	
Admin State	The desired state of DHCP
Lease Populate	Not applicable
Action	The processing required that occurs when the 7705 SAR receives a DHCP request that already has a Relay Agent Information Option (Option 82):
Trusted	Indicates whether trusted mode is enabled or disabled on the IP interface
ICMP Details	
Redirects	The rate for ICMP redirect messages
Unreachables	The rate for ICMP unreachable messages
TTL Expired	The rate for ICMP TTL messages
IPCP Address Extension Details	
Peer IP Addr	Specifies the remote IP address to be assigned to the far-end of the associated PPP/MLPPP link via IPCP extensions
Peer Pri DNS Addr	Specifies a unicast IPv4 address for the primary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions
Peer Sec DNS Addr	Specifies a unicast IPv4 address for the secondary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions

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 VPRN instance.

Parameters *ip-address* — the IP address for which ARP entries will be displayed

Values a.b.c.d (1.0.0.0 to 223.255.255.255)

Default all IP addresses

ieee-address — the 48-bit MAC address for which ARP entries will be displayed. 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.

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

Default all MAC addresses

sap-id — the SAP ID for which ARP entries will be displayed. See Table 39 for a full list of SAP IDs.

ip-int-name — the IP interface name for which to display matching ARPs

Output The following output is an example of service ID ARP information, and Table 114 describes the fields.

Sample Output

```
*A:ALU-12>show>service>id# arp
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Type    Expiry    Interface         SAP
-------------------------------------------------------------------------------
190.11.1.1      00:03:fa:00:08:22 Other   00:00:00  ies-100-190.11.1  1/1/11:0
===============================================================================
*A:ALU-12>show>service>id#
```

Table 114: Show Service ID ARP Output Fields

Label	Description
IP Address	The IP address of the ARP entry
MAC Address	The MAC address 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

Table 114: Show Service ID ARP Output Fields (Continued)

Label	Description
Expiry	The age of the ARP entry
Interface	The IP interface name associated with the ARP entry
SAP	The port identifier of the SAP

base

Syntax **base**

Context show>service>id

Description This command displays basic information about the service ID, including service type, description, SAPs and SDPs.

Output The following output is an example of service ID base information, and Table 115 describes the fields.

Sample Output

```
*A:ALU-12>show>service>id# base
===============================================================================
Service Basic Information
===============================================================================
Service Id        : 6
Service Type      : VPRN
Name              : VPRN_global
Description       : (Not Specified)
Customer Id       : 1
Last Status Change: 02/03/2010 21:59:01
Last Mgmt Change  : 03/01/2010 18:55:58
Admin State       : Down                 Oper State        : Down

Route Dist.       : 12.34.56.78:655      VPRN Type         : regular
AS Number         : None                 Router Id         : 255.0.0.0
ECMP              : Enabled              ECMP Max Routes   : 1
Max IPv4 Routes   : No Limit             Auto Bind         : None
Max IPv6 Routes   : No Limit
Vrf Target        : None
Vrf Import        : None
Vrf Export        : None

SAP Count         : 1                    SDP Bind Count    : 1
-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                               Type         AdmMTU  OprMTU  Adm   Opr
-------------------------------------------------------------------------------
sap:1/5/2                                null         1514    1514    Up    Down
sdp:1:6 S(10.10.10.10)                   n/a          0       0       Up    Down
===============================================================================
```

Table 115: Show Service ID Base Output Fields

Label	Description
Service Id	The service identifier
Service Type	The type of service: VPRN
Name	The service name
Description	Generic information about the service
Customer Id	The customer identifier
Last Status Change	The date and time of the most recent status change to this service
Last Mgmt Change	The date and time of the most recent management-initiated change to this service
Admin State	The desired state of the service
Oper State	The operating state of the service
Route Dist.	The largest frame size (in octets) that the service can handle
VPRN Type	Only valid in services that accept mesh SDP bindings. It validates the VC ID portion of each mesh SDP binding defined in the service.
AS Number	The autonomous system number
Router ID	The router ID for this service
ECMP	Displays equal cost multipath information
ECMP Max Routes	The maximum number of routes that can be received from the neighbors in the group or for the specific neighbor
Max IPv4 Routes	The maximum number of routes that can be used for path sharing
Auto Bind	The automatic binding type for the SDP assigned to this service
Max IPv6 Routes	Not applicable
Vrf Target	The route target in the VRF applied to this service
Vrf Import	The VRF import policy applied to this service
Vrf Export	The VRF export policy applied to this service
SAP Count	The number of SAPs defined on the service
SDP Bind Count	The number of SDPs bound to the service
Service Access and Destination Points	
Identifier	The service access (SAP) and destination (SDP) points

Table 115: Show Service ID Base Output Fields (Continued)

Label	Description
Type	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP
AdmMTU	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
OprMTU	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
Adm	The administrative state of the SAP or SDP
Opr	The operating state of the SAP or SDP

dhcp

Syntax **dhcp**

Context show>service>id

Description This command enables the context to display DHCP information for the specified service.

statistics

Syntax **statistics** [**interface** *ip-int-name* | *ip-address*]

Context show>service>id>dhcp

Description This command displays DHCP statistics information.

Parameters *ip-int-name* — the interface name for which DHCP statistics will be displayed

ip-address — the IP address of the interface for which to display information

Values a.b.c.d (host bits must be 0)

Output The following output is an example of server ID DHCP statistics information, and Table 116 describes the fields.

Sample Output

```
*A:ALU-12>show>service>id>dhcp# statistics
===============================================================
DHCP Global Statistics, service 6
===============================================================
Rx Packets                          : 0
```

```
Tx Packets                          : 0
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 0
Client Packets Discarded            : 0
Client Packets Relayed              : 0
Server Packets Discarded            : 0
Server Packets Relayed              : 0
===============================================================
*A:ALU-12>show>service>id>dhcp#
```

Table 116: Show Service ID DHCP Statistics Output Fields

Label	Description
DHCP Global Statistics, service x	
Rx Packets	The number of packets received from the DHCP clients
Tx Packets	The number of packets transmitted to the DHCP clients
Rx Malformed Packets	The number of corrupted/invalid packets received from the DHCP clients
Rx 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
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

summary

Syntax **summary** [**interface** *interface-name* | **saps**]

Context show>service>id>dhcp

Description This command displays DHCP configuration summary information.

Parameters *interface-name* — the interface name for which DHCP summary statistics will be displayed

saps — displays SAPs per interface

Output The following output is an example of service ID DHCP summary information, and Table 117 describes the fields.

Sample Output

```
*A:ALU-12>show>service>id>dhcp# summary
===============================================================================
DHCP Summary, service 6
===============================================================================
Interface Name                   Arp      Used/                Info    Admin
  SapId/Sdp                      Populate Provided             Option  State
-------------------------------------------------------------------------------
vprn_interface                   No       0/0                  Keep    Down
-------------------------------------------------------------------------------
Interfaces: 1
===============================================================================
*A:ALU-12>show>service>id>dhcp#
```

Table 117: Show Service ID DHCP Summary Output Fields

Label	Description
DHCP Summary, service x	
Interface Name Sap/Sdp	The name of the interface
Arp Populate	Specifies whether ARP populate is enabled or disabled
Used/Provided	Used—the number of lease-states that are currently in use on the specified interface, that is, the number of clients on the interface that got an IP address by DHCP. This value is always less than or equal to the "Provided" field.
	Provided—the lease-populate value that is configured for a specific interface
Info Option	Keep—the existing information is kept on the packet and the router does not add any additional information
	Replace—on ingress, the existing information-option is replaced with the information-option from the router
	Drop—the packet is dropped and an error is logged
Admin State	The administrative state

interface

Syntax **interface** [{[*ip-address* | *ip-int-name*] [*interface-type*] [**detail**] [**family**]} | **summary**]

Context show>service>id

Description This command 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 a.b.c.d (host bits must be 0)

ip-int-name — the IP interface name for which to display information

interface type — displays either group or subscriber interfaces

detail — displays detailed IP interface information

family — displays the IP interface family

Values ipv4, ipv6

summary — displays summary IP interface information

Output The following output is an example of service ID interface information, and Table 118 describes the fields.

Sample Output

```
*A:ALU-12>show>service>id# interface
===============================================================================
Interface Table
===============================================================================
Interface-Name                   Adm         Opr(v4/v6)  Type    Port/SapId
   IP-Address                                                    PfxState
-------------------------------------------------------------------------------
vprn_interface                   Up          Down/Down   VPRN    1/5/2
   -                                                             -
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================

*A:ALU-12>show>service>id# interface summary
===============================================================================
Service Interface Summary
===============================================================================
Service Id          Interfaces          Admin-Up            Oper-Up(v4/v6)
-------------------------------------------------------------------------------
6                   1                   1                   0/0
===============================================================================
*A:ALU-12>show>service>id#


*A:ALU-12>show>service>id# interface detail
===============================================================================
```

```
Interface Table
===============================================================================

-------------------------------------------------------------------------------
Interface
-------------------------------------------------------------------------------
If Name      : vprn_interface
Admin State  : Up                                Oper (v4/v6)     : Down/Down
Protocols    : None

IP Addr/mask : Not Assigned
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
If Index     : 2                                 Virt. If Index   : 2
Last Oper Chg: 02/03/2010 21:59:02               Global If Index  : 125
SAP Id       : 1/5/2
TOS Marking  : Trusted                           If Type          : VPRN
SNTP B.Cast  : False
MAC Address  :                                   Arp Timeout      : 14400
IP MTU       : 1500                              ICMP Mask Reply  : True
Arp Populate : Disabled                          Host Conn Verify : Disabled
LdpSyncTimer : None

Proxy ARP Details
Rem Proxy ARP: Disabled                          Local Proxy ARP  : Disabled
Policies     : none

Proxy Neighbor Discovery Details
Local Pxy ND : Disabled
Policies     : none

DHCP Details
Admin State  : Down                              Lease Populate   : 0
Action       : Keep                              Trusted          : Disabled

DHCP6 Relay Details
Admin State  : Down                              Lease Populate   : 0
Oper State   : Down                              Nbr Resolution   : Disabled
If-Id Option : None                              Remote Id        : Disabled
Src Addr     : Not configured

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

IPCP Address Extension Details
Peer IP Addr*: Not configured
Peer Pri DNS*: Not configured
Peer Sec DNS*: Not configured
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:ALU-12>show>service>id#
```

Table 118: Show Service ID Interface Detailed Output Fields

Label	Description
Interface	
If Name	The name used to refer to the interface
Admin State	The desired state of the interface
Oper (v4/v6)	The operating state of the interface
Protocols	The protocols supported on this interface
IP Addr/mask	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.
Virt. If Index	The virtual interface index of the VPRN interface
Last Oper Chg	The date and time of the last operating state change on the interface
Global If Index	The global interface index of the VPRN interface
TOS Marking	Specifies whether the ToS marking is trusted or untrusted for the interface
If Type	The interface type
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled
MAC Address	The 48-bit IEEE 802.3 MAC address
Arp Timeout	The timeout for an ARP entry learned on the interface
IP MTU	The IP maximum transmit unit for the interface
ICMP Mask Reply	Specifies whether ICMP mask reply is enabled or disabled
ARP Populate	Specifies if ARP is enabled or disabled
Host Conn Verify	Not applicable
LdpSyncTimer	Not applicable
Proxy ARP Details	
Rem Proxy ARP	Indicates whether remote proxy ARP is enabled or disabled
Local Proxy ARP	Indicates whether local proxy ARP is enabled or disabled
Policies	Specifies the policy statement(s) applied to proxy ARP

Table 118: Show Service ID Interface Detailed Output Fields (Continued)

Label	Description
DHCP Details	
Admin State	The administrative state of DHCP
Lease Populate	Not applicable
Action	The processing required that occurs when the 7705 SAR receives a DHCP request that already has a Relay Agent Information Option (Option 82):
Trusted	Indicates whether trusted mode is enabled or disabled on the IP interface
ICMP Details	
Redirects	The rate for ICMP redirect messages
Unreachables	The rate for ICMP unreachable messages
TTL Expired	The rate for ICMP TTL messages
IPCP Address Extension Details	
Peer IP Addr	The remote IP address to be assigned to the far-end of the associated PPP/MLPPP link via IPCP extensions
Peer Pri DNS	The unicast IPv4 address for the primary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions
Peer Sec DNS	The unicast IPv4 address for the secondary DNS server to be signaled to the far-end of the associate PPP/MLPPP link via IPCP extensions

sap

Syntax **sap** [*sap-id* [**detail**]]

Context show>service>id

Description This command displays information for the SAPs associated with the service.

If no optional parameters are specified, a summary of all associated SAPs is displayed.

Parameters *sap-id* — the SAP ID for which SAP information will be displayed. See Table 39 for a full list of SAP IDs.

detail — displays detailed information for the SAP

Output The following output is an example of service SAP information, and Table 119 describes the fields.

Sample Output

```
*A:ALU-12>show>service>id# sap
===============================================================================
SAP(Summary), Service 6
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
1/5/2                           6          2     ip4     2     none   Up   Down
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-12>show>service>id# sap 1/5/2
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 6
SAP                : 1/5/2                    Encap             : null
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Multi Svc Site     : None
Last Status Change : 02/03/2010 21:59:01
Last Mgmt Change   : 02/03/2010 21:59:02
===============================================================================
*A:ALU-12>show>service>id#

*A:ALU-12>show>service>id# sap 1/5/2 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 6
SAP                : 1/5/2                    Encap             : null
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Multi Svc Site     : None
Last Status Change : 02/03/2010 21:59:01
```

```
Last Mgmt Change   : 02/03/2010 21:59:02
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                    QinQ Ethertype    : 0x8100

Admin MTU          : 1514                      Oper MTU          : 1514
Ingr IP Fltr-Id    : 2                         Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                       Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                       Egr IPv6 Fltr-Id  : n/a
tod-suite          : None                      qinq-pbit-marking : both
Ing Scheduler Mode : 4-priority                Egr Scheduler Mode: 16-priority
                                               Egr Agg Rate Limit: 7000
                                               Egr Agg cir       : 700
Q Frame-Based Acct : Disabled

Acct. Pol          : None                      Collect Stats     : Disabled

Anti Spoofing      : None                      Nbr Static Hosts  : 0

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 2                         Egress qos-policy : 2
Shared Q plcy      : n/a                       Multipoint shared : Disabled
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0

Queueing Stats(Ingress QoS Policy 2)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 2)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

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

Egress Queue 1
```

```
For. InProf           : 0                        0
For. OutProf          : 0                        0
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
===============================================================================
*A:ALU-12>show>service>id#


*A:ALU-12>>show>service>id# sap 1/5/2 atm
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 6
SAP                : 1/5/2                    Encap             : null
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Multi Svc Site     : None
Last Status Change : 02/03/2010 21:59:01
Last Mgmt Change   : 02/03/2010 21:59:02
===============================================================================
*A:ALU-12>show>service>id#


*A:ALU-12>show>service>id# sap 1/5/2 qos
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 6
SAP                : 1/5/2                    Encap             : null
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Multi Svc Site     : None
Last Status Change : 02/03/2010 21:59:01
Last Mgmt Change   : 02/03/2010 21:59:02

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 2                        Egress qos-policy : 2
Shared Q plcy      : n/a                      Multipoint shared : Disabled
===============================================================================
*A:ALU-12>show>service>id#


*A:ALU-12>show>service>id# sap 1/5/2 sap-stats
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 6
SAP                : 1/5/2                    Encap             : null
Admin State        : Up                       Oper State        : Down

Flags              : ServiceAdminDown
                     PortOperDown
Multi Svc Site     : None
Last Status Change : 02/03/2010 21:59:01
Last Mgmt Change   : 02/03/2010 21:59:02
```

```
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
Forwarding Engine Stats
Dropped               : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0

Queueing Stats(Ingress QoS Policy 2)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 2)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
===============================================================================
*A:ALU-12>show>service>id# sap 1/5/2 stats

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 6
SAP                : 1/5/2                    Encap             : null
Admin State        : Up                       Oper State        : Down
Flags              : ServiceAdminDown
                     PortOperDown
Multi Svc Site     : None
Last Status Change : 02/03/2010 21:59:01
Last Mgmt Change   : 02/03/2010 21:59:02
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
===============================================================================
*A:ALU-12>show>service>id#
```

Table 119: Show Service ID SAP Detailed Output Fields

Label	Description
Service Id	The service identifier
SAP	The SAP identifier
Encap	The encapsulation type of the SAP
Admin State	The administrative state of the SAP
Oper State	The operating state of the SAP
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, NoSapIpipeCeIpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode
Last Status Change	The time of the most recent operating status change to this SAP
Last Mgmt Change	The time of the most recent management-initiated change to this SAP
Sub Type	The supported sub type: regular
Dot1Q Ethertype	The value of the dot1q Ethertype
QinQ Ethertype	The value of the qinq Ethertype
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented
Ingr IP Fltr-Id	The ingress filter policy ID assigned to the SAP
Egr IP Fltr-Id	The egress filter policy ID assigned to the SAP
Ingr Mac Fltr-Id	Not applicable
Egr Mac Fltr-Id	Not applicable
tod-suite	Indicates whether a time-based policy is applied to a multiservice site
qinq-pbit-marking	Indicates the qinq P-bit marking for the service: both or top

Table 119: Show Service ID SAP Detailed Output Fields (Continued)

Label	Description
Ing Scheduler Mode	The scheduler mode for the SAP in the access ingress direction: 4-priority or 16-priority
Egr Scheduler Mode	The scheduler mode for the SAP in the access egress direction: 4-priority or 16-priority
Ing Agg Rate Limit	The PIR rate limit in the access ingress direction for the aggregate of the SAP queues
Egr Agg Rate Limit	The PIR rate limit in the access egress direction for the aggregate of the SAP queues
Ing Agg cir	The CIR rate limit in the access ingress direction for the aggregate of the SAP queues
Egr Agg cir	The CIR rate limit in the access egress direction for the aggregate of the SAP queues
Acct. Pol	The accounting policy ID assigned to the SAP
Collect Stats	Specifies whether accounting statistics are collected on the SAP
Anti Spoofing	Not applicable
Nbr Static Hosts	Not applicable
QOS	
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP
Egress qos-policy	The egress QoS policy ID assigned to the SAP
Shared Q plcy	Not applicable
Multipoint shared	Not applicable
Sap Statistics	
Last Cleared Time	The date and time that a clear command was issued on statistics
Forwarding Engine Stats	
Dropped	The number of packets or octets dropped by the forwarding engine
Off. HiPrio	The number of high-priority packets or octets offered to the forwarding engine
Off. LowPrio	The number of low-priority packets offered to the forwarding engine
Queueing Stats (Ingress QoS Policy)	
Dro. HiPrio	The number of high-priority packets or octets discarded, as determined by the SAP ingress QoS policy

Table 119: Show Service ID SAP Detailed Output Fields (Continued)

Label	Description
Dro. LowPrio	The number of low-priority packets discarded, as determined by the SAP ingress QoS policy
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP ingress QoS policy
For. OutProf	The number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP ingress QoS policy
Queueing Stats (Egress QoS Policy)	
Dro. InProf	The number of in-profile packets or octets discarded, as determined by the SAP egress QoS policy
Dro. OutProf	The number of out-of-profile packets or octets discarded, as determined by the SAP egress QoS policy
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded, as determined by the SAP egress QoS policy
For. OutProf	The number of out-of-profile packets or octets (rate above CIR) forwarded, as determined by the SAP egress QoS policy
Sap per Queue stats	
Ingress Queue *n*	The index of the ingress QoS queue of this SAP, where *n* is the index number
Off. HiPrio	The number of packets or octets of high-priority traffic for the SAP (offered)
Off. LoPrio	The number of packets or octets count of low-priority traffic for the SAP (offered)
Dro. HiPrio	The number of high-priority traffic packets or octets dropped
Dro. LoPrio	The number of low-priority traffic packets or octets dropped
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	The number of out-of-profile packets or octets (rate above CIR) forwarded
Egress Queue *n*	The index of the egress QoS queue of the SAP, where *n* is the index number
For. InProf	The number of in-profile packets or octets (rate below CIR) forwarded
For. OutProf	The number of out-of-profile packets or octets (rate above CIR) forwarded
Dro. InProf	The number of in-profile packets or octets dropped for the SAP

Table 119: Show Service ID SAP Detailed Output Fields (Continued)

Label	Description
Dro. OutProf	The number of out-of-profile packets or octets discarded

ingress-label

Syntax **ingress-label** *start-label* [*end-label*]

Context show>service

Description This command displays service information 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 labels in the specified range are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *end-label* — the ending ingress label value for which to display services using the label range

Values 2049 to 131071

Default the *start-label* value

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, or 2048 to 131071

Output The following output is an example of service ingress label information, and Table 120 describes the fields.

Sample Output

In the example below, services 3, 5 and 6 are IES, and services 5000 and 5001 are VPLS services.

```
*A:ALU-12>show>service# ingress-label 0 131071
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Binding        Type  I.Lbl                E.Lbl
-------------------------------------------------------------------------------
3          15:15              Spok  0                    0
5          5:5                Spok  0                    0
6          5:6                Spok  0                    0
5000       15:5000            Mesh  0                    0
5000       15:5001            Spok  0                    0
5001       5001:100           Spok  0                    0
-------------------------------------------------------------------------------
```

```
Number of Bindings Found : 6
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-12#
```

Table 120: Show Service ingress Output Fields

Label	Description
Svc Id	The ID that identifies a service
Sdp Binding	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 label range

sdp

Syntax **sdp** {[*sdp-id*[*:vc-id*] | **far-end** *ip-addr*]} [**detail**]
sdp [*sdp-id*[*:vc-id*]]

Context show>service>id

Description This command displays information for the SDPs associated with the service. If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — the SDP ID for which SDP information will be displayed

Values 1 to 17407

Default all SDPs

vc-id — the virtual circuit ID on the SDP ID

Values 1 to 4294967295

ip-addr **—** displays only SDPs matching with the specified far-end IP address

detail — displays detailed SDP information

Output The following output is an example of service ID SDP information, and Table 121 describes the fields.

Sample Output

```
*A:ALU-12>show>service>id# sdp
===============================================================================
Services: Service Destination Points
===============================================================================
SdpId             Type IP address      Adm     Opr       I.Lbl       E.Lbl
-------------------------------------------------------------------------------
1:6               Spok 10.10.10.10     Up      Down      n/a         n/a
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-12>show>service>id#


*A:ALU-12>show>service>id# sdp 1
===============================================================================
Service Destination Point (Sdp Id : 1)
===============================================================================
SdpId             Type IP address      Adm     Opr       I.Lbl       E.Lbl
-------------------------------------------------------------------------------
1:6               Spok 10.10.10.10     Up      Down      n/a         n/a
===============================================================================
*A:ALU-12>show>service>id#


*A:ALU-12>show>service>id# sdp far-end 10.10.10.10
===============================================================================
Service Destination Point(Far-End : 10.10.10.10)
===============================================================================
SdpId             Type IP address      Adm     Opr       I.Lbl       E.Lbl
-------------------------------------------------------------------------------
1:6               Spok 10.10.10.10     Up      Down      n/a         n/a
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-12>show>service>id#


*A:ALU-12>show>service>id# sdp detail
===============================================================================
Services: Service Destination Points Details
===============================================================================
-------------------------------------------------------------------------------
 Sdp Id 1:6  -(10.10.10.10)
-------------------------------------------------------------------------------
SDP Id              : 1:6                       Type              : Spoke
VC Type             : n/a                       VC Tag            : n/a
Admin Path MTU      : 0                         Oper Path MTU     : 0
Far End             : 10.10.10.10               Delivery          : MPLS
Admin State         : Up                        Oper State        : Down
Acct. Pol           : None                      Collect Stats     : Disabled
Ingress Label       : n/a                       Egress Label      : n/a
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 : 02/03/2010 21:59:01      Signaling         : n/a
Last Mgmt Change   : 03/01/2010 18:55:58
Class Fwding State : Down
Flags              : SdpOperDown

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. Dro. Pkts.     : n/a
I. Fwd. Octs.      : n/a                      I. Dro. Octs.     : n/a
E. Fwd. Pkts.      : n/a                      E. Fwd. Octets    : n/a

Associated LSP LIST :
No LSPs Associated
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
===============================================================================
```

Table 121: Show Service ID SDP Detailed Output Fields

Label	Description
Sdp Id	The SDP identifier
Type	Indicates whether the SDP is a spoke or a mesh
VC Type	The VC type: ether or vlan
VC Tag	The explicit dot1Q value used when encapsulating to the SDP far end
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	The IP address of the remote end of the GRE, MPLS, or IP tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: GRE, MPLS, or IP
Admin State	The administrative state of this SDP
Oper State	The operational state of this SDP
Acct. Pol	The accounting policy applied to the SDP
Collect Stats	Specifies whether accounting statistics are collected on the SDP

Table 121: Show Service ID SDP Detailed Output Fields (Continued)

Label	Description
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
Ing mac Fltr	Not applicable
Egr mac Fltr	Not applicable
Ing ip Fltr	The ingress filter policy ID assigned to the SDP
Egr ip Fltr	The egress filter policy ID assigned to the SDP
Admin ControlWord	The administrative state of the control word: Preferred (control word enabled) or Not Preferred (control word disabled)
Oper ControlWord	The operational state of the control word: True (control word enabled) or False (control word disabled)
Last Status Change	The date and time of the most recent change to the SDP
Signaling	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP
Last Mgmt Change	The date and time of the most recent management-initiated change to this SDP
Class Fwding State	Not applicable
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
KeepAlive Information	
Admin State	The administrative state of the keepalive process
Oper State	The operational state of the keepalive process
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP

Table 121: Show Service ID SDP Detailed Output Fields (Continued)

Label	Description
Hell Msg Length	The length of the SDP echo request messages transmitted on this SDP
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 time to wait before the keepalive operating status is eligible to enter the alive state
Statistics	
I. Fwd. Pkts.	The number of forwarded ingress packets
I. Dro. Pkts.	The number of dropped ingress packets
I. Fwd. Octs.	The number of forwarded ingress octets
I. Dro. Octs.	The number of dropped ingress octets
E. Fwd. Pkts.	The number of forwarded egress packets
E. Fwd. Octets	The number of forwarded egress octets
Associated LSP LIST	If the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far-end field. If the SDP type is GRE, the following message displays: SDP delivery mechanism is not MPLS.
Number of SDPs	The total number of SDPs applied to this service ID

twamp-light

Syntax **twamp-light**

Context show>service>id

Description This command displays OAM TWAMP Light status information.

Output The following output is an example of TWAMP Light information, and Table 122 describes the fields.

Sample Output

```
*A:ALU-3# show service id 10 twamp-light
===============================================================================
TWAMP Light Reflector
-------------------------------------------------------------------------------
```

```
Admin State          : Up
Up Time              : 0d 00:12:01
Configured UDP Port  : 1025
Test Packets Rx      : 0                   Test Packets Tx     : 0

TWAMP Light Controller Prefix List
 192.168.1.1/32
    10.1.1.2/32
 172.16.254.9/3
    10.1.1.0/24
===============================================================================
*A:ALU-3#
```

Table 122: Show TWAMP Light Output Fields

Label	Description
TWAMP Light Reflector	
Admin State	Displays one of the following: Up—the server or prefix is administratively enabled (no shutdown) in configuration Down—the server or prefix is administratively disabled (shutdown) in configuration
Up Time	The time since the server process was started, measured in days (d), hours, minutes, and seconds
Configured UDP Port	The UDP port number used
Test Packets Rx	The total number of test packets received from session senders
Test Packets Tx	The total number of test packets sent to session senders
TWAMP Light Controller Prefix List	The IP address prefixes of TWAMP Light clients

service-using

Syntax **service-using vprn** [**sdp** *sdp-id*] [**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 **vprn —** displays matching VPRN services

sdp-id **—** displays only services bound to the specified SDP ID

Values 1 to 17407

customer-id — displays only those services associated with the specified customer ID

Values 1 to 2147483647

Output The following output is an example of service-using information, and Table 123 describes the fields.

Sample Output

```
*A:ALU-12>show>service# service-using vprn
===============================================================================
Services [vprn]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
6            VPRN      Down   Down       1                 03/01/2010 18:55:58
-------------------------------------------------------------------------------
Matching Services : 1
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-12>show>service#


*A:ALU-12>show>service# service-using customer 1
===============================================================================
Services  Customer 1
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
1            Cpipe     Down   Down       1                 02/03/2010 21:59:01
2            Epipe     Down   Down       1                 02/03/2010 21:59:02
5            Apipe     Down   Down       1                 02/03/2010 21:59:02
6            VPRN      Down   Down       1                 03/01/2010 18:55:58
23           IES       Down   Down       1                 02/03/2010 21:59:01
100          Ipipe     Down   Down       1                 02/03/2010 21:59:02
-------------------------------------------------------------------------------
Matching Services : 6
-------------------------------------------------------------------------------
===============================================================================
*A:ALU-12>show>service#
```

Table 123: Show Service Service-Using Output Fields

Label	Description
Service Id	The service identifier
Type	The service type configured for the service ID
Name	The service name
Description	A description of the service
Adm	The desired state of the service
Opr	The operating state of the service

Table 123: Show Service Service-Using Output Fields (Continued)

Label	Description
CustomerID	The ID of the customer who owns this service
Last Mgmt Change	The date and time of the most recent management-initiated change to this service
Matching Services	The number of services of the same type

Show Router Commands

Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

router

Syntax **router** [*service-id*]

Context show

Description This command provides access to the show commands for the service specified by the *service-id*.

Using *service-id* with this command applies only to VPRN service.

Parameters *service-id* — specifies the *service-id*

Values 1 to 2147483647 or *service-name*

aggregate

Syntax **aggregate** [**active**]

Context show>router

Description This command displays aggregated routes.

Parameters **active —** filters out inactive aggregates

Output The following output is an example of aggregate route information, and Table 124 describes the fields.

Sample Output

```
*A:ALU-12# show router 104 aggregate
===============================================================================
Aggregates (Service: 104)
===============================================================================
Prefix                                          Aggr IP-Address   Aggr AS
   Summary                                         AS Set            State
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of Aggregates: 0
===============================================================================
*A:ALU-12#
```

Table 124: Show Aggregate Table Output Fields

Label	Description
Prefix	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	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

arp

Syntax **arp** [*ip-address* | *ip-int-name* | **mac** *ieee-mac-addr*] [**sdp** *sdp-id:vc-id*] [**summary**]

Context show>router

Description 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* — the IP address for which ARP entries will be displayed

ip-int-name — the interface name for which ARP entries will be displayed

ieee-mac-addr — the MAC address for which ARP entries will be displayed

sdp-id — the SDP ID for which ARP entries will be displayed

Values 1 to 17407

Default all SDPs

vc-id — the virtual circuit ID on the SDP ID

Values 1 to 4294967295

summary — displays summary APR table information

Output The following output is an example of ARP table information, and Table 125 describes the fields.

Sample Output

```
*A:ALU-12# show router 6 arp
===============================================================================
ARP Table (Service: 6)
===============================================================================
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   ALU-1-2
-------------------------------------------------------------------------------
No. of ARP Entries: 2
===============================================================================
*A:ALU-12#
```

Table 125: Show 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

bgp

Syntax **bgp**

Context show>router

Description This command enables the context to display BGP-related information.

damping

Syntax **damping** [*ip-prefix*[/*ip-prefix-length*]] [*damp-type*] [**detail**] [**ipv4**]
damping [*ip-prefix*[/*ip-prefix-length*]] [*damp-type*] [**detail**] **ipv6**
damping [*ip-prefix*[/*ip-prefix-length*]] [*damp-type*] [**detail**] **vpn-ipv4**
damping [*ip-prefix*[/*ip-prefix-length*]] [*damp-type*] [**detail**] **vpn-ipv6**

Context show>router>bgp

Description This command displays BGP routes that have been dampened due to route flapping. This command can be entered with or without a route parameter. If no parameters are included, all dampened routes are listed.

If the keyword **detail** is included, more detailed information is displayed.

If a *damp-type* is specified, only those types of dampened routes (decayed, history, or suppressed) are displayed. Routes that have a state of **decayed** have gained penalties for flapping but have not yet reached the suppression limit. Routes that have a state of **history** have had a route flap and have been withdrawn. Routes that have a state of **suppressed** have reached the suppression limit and are not considered in BGP path selection.

Parameters *ip-prefix/ip-prefix-length* — displays damping information for the specified IP address

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 30
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 to FFFF]H d: [0 to 255]D
ipv6-prefix-length	0 to 128

damp-type — displays damping information for routes with the specified damp type

Values decayed, history, suppressed

detail — displays detailed information

ipv4 — displays dampened routes for the IPv4 address family

ipv6 — displays dampened routes for the IPv6 address family

vpn-ipv4 — displays dampened routes for the VPN-IPv4 address family

vpn-ipv6 — displays dampened routes for the VPN-IPv6 address family

Output The following output is an example of BGP damping information, and Table 126 describes the fields.

Sample Output

```
*A:ALU-12# show router 6 bgp damping
===============================================================================
 BGP Router ID:0.0.0.14         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, b - backup
===============================================================================
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:ALU-12#

*A:ALU-12# show router 6 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:ALU-12#
```

```
*A:ALU-12# show router 6 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:ALU-12#

*A:ALU-12# show router 6 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:ALU-12#
```

Table 126: Show 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, it is the same value as the AS
Network	The IP prefix and mask length for the route
Flag(s)	Legend: Status codes: u-used, s-suppressed, h-history, d-decayed, *-valid (if an * is not present, the status is invalid) Origin codes: i-IGP, e-EGP, ?-incomplete, >-best
From	The originator ID path attribute value
Reuse/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	The BGP next hop for the route
Peer AS	The autonomous system number of the advertising router

Table 126: Show BGP Damping Output Fields (Continued)

Label	Description
Peer Router-Id	The router ID of the advertising router
Local Pref	The BGP local preference path attribute for the route
Age	The time elapsed since the service was enabled
Last update	The time that BGP was last updated
FOM Present	The current Figure of Merit (FOM) value
FOM Last upd.	The last updated 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

group

Syntax **group** [*name*] [**detail**]

Context show>router>bgp

Description 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. Valid states are:

- Up — BGP global process is configured and running
- Down — BGP global process is administratively shut down and not running
- Disabled — BGP global process is operationally disabled. The process must be restarted by the operator.

Parameters *name* — displays information for the specified BGP group

detail — displays detailed information

Output The following output is an example of BGP group information, and Table 127 describes the fields.

Sample Output

```
*A:ALU-12# show router 6 bgp group
===============================================================================
BGP Group
===============================================================================
-------------------------------------------------------------------------------
Group            : bgp_group
-------------------------------------------------------------------------------
Group Type       : No Type              State            : Up
Peer AS          : n/a                  Local AS         : 1
Local Address    : n/a                  Loop Detect      : Ignore
Import Policy    : None Specified / Inherited
Export Policy    : None Specified / Inherited
Hold Time        : 90                   Keep Alive       : 30
Cluster Id       : None                 Client Reflect   : Enabled
NLRI             : Unicast              Preference       : 170
TTL Security     : Enabled              Min TTL Value    : 10
Graceful Restart : Enabled              Stale Routes Time: 360
Auth key chain   : n/a
Bfd Enabled      : Disabled

List of Peers
- 44.44.44.44 :

Total Peers      : 1                    Established      : 0
-------------------------------------------------------------------------------
Peer Groups : 1

*A:ALU-12# show router 6 bgp group detail
===============================================================================
BGP Group  (detail)
===============================================================================
-------------------------------------------------------------------------------
Group            : bgp_group
-------------------------------------------------------------------------------
Group Type       : No Type              State            : Up
Peer AS          : n/a                  Local AS         : 1
Local Address    : n/a                  Loop Detect      : Ignore
Connect Retry    : 120                  Authentication   : None
Local Pref       : 100                  MED Out          : 0
Multihop         : 0 (Default)          AS Override      : Disabled
Min Route Advt.  : 30                   Min AS Originate : 15
Prefix Limit     : No Limit             Passive          : Disabled
Next Hop Self    : Disabled             Aggregator ID 0  : Disabled
Remove Private   : Disabled             Damping          : Enabled
Import Policy    : None Specified / Inherited
Export Policy    : None Specified / Inherited
Hold Time        : 90                   Keep Alive       : 30
Cluster Id       : None                 Client Reflect   : Enabled
NLRI             : Unicast              Preference       : 170
TTL Security     : Enabled              Min TTL Value    : 10
Graceful Restart : Enabled              Stale Routes Time: 360
Auth key chain   : n/a
Bfd Enabled      : Disabled

List of Peers
- 44.44.44.44 :
```

```
Total Peers      : 1                    Established      : 0
-------------------------------------------------------------------------------
Peer Groups : 1
===============================================================================
*A:ALU-12#
```

Table 127: Show BGP Group Output Fields

Label	Description
Group	The 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
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 that is assigned to advertised routes
Multihop	The maximum number of router hops a BGP connection can traverse
AS Override	The setting of the AS override
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

Table 127: Show BGP Group Output Fields (Continued)

Label	Description
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 a BGP connection with a neighbor in the specified peer group Enabled: BGP will not actively attempt to establish a BGP connection with a neighbor in the specified peer group
Next Hop Self	Disabled: BGP is not configured to send only its own IP address as the BGP next hop in route updates to neighbors in the peer group Enabled: BGP sends only its own IP address as the BGP next hop 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
Damping	Disabled: the peer group is configured not to dampen route flaps Enabled: the peer group is configured to dampen route flaps
Import Policy	The configured import policies for the peer group
Export Policy	The configured export policies for the peer group
Hold Time	The configured hold time setting
Keep Alive	The configured keepalive setting
Cluster Id	The configured route reflector 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 network layer reachability information that the specified peer group can accept Unicast: IPv4 unicast routing information can be carried

Table 127: Show BGP Group Output Fields (Continued)

Label	Description
Preference	The configured route preference value for the peer group
TTL Security	Enabled: TTL security is enabled Disabled: TTL security is disabled
Min TTL Value	The minimum TTL value configured for the peer
Graceful Restart	The state of graceful restart
Stale Routes Time	The length of time that stale routes are kept in the route table
Auth key chain	The value for the authentication key chain
Bfd Enabled	Enabled: BFD is enabled Disabled: BFD is disabled
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
Peer Groups	The number of peer groups

neighbor

Syntax **neighbor** [*ip-address* [[*family*] *filter1* [**brief**]]]
neighbor [*as-number* [[*family*] *filter2*]]
neighbor [*ip-address*] **orf** [*filter3*]
neighbor [*ip-address*] **graceful-restart**

Context show>router>bgp

Description 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 displays.

Note: This information is not available when using SNMP.

Parameters *ip-address* — the specified IP address for which to display information

Values *ipv4-address*: a.b.c.d (host bits must be 0)

as-number — the specified AS number for which to display information

Values 1 to 4294967295

family — the type of routing information to be distributed by this peer group

Values **ipv4** — displays only those BGP peers that have the IPv4 family enabled and not those capable of exchanging IP-VPN routes

vpn-ipv4 —displays the contents of the multicast routing table

filter1 — displays information for the specified IP address

Values **received-routes** — displays the number of routes received from this peer

advertised-routes — displays the number of routes advertised by this peer

history — displays statistics for dampened routes

suppressed — displays the number of paths from this peer that have been suppressed by damping

detail — displays detailed information pertaining to *filter1*



Note: When either **received-routes** or **advertised-routes** is specified, the routes that are received from or sent to the specified peer are listed. When either **history** or **suppressed** is specified, the routes that are learned from those peers that either have a history or are suppressed are listed.

filter2 — displays information for the specified AS number

Values **history** — displays statistics for dampened routes

suppressed — displays the number of paths from this peer that have been suppressed by damping

detail — displays detailed information pertaining to *filter2*

brief — displays information in a brief format. This parameter is only supported with **received-routes** and **advertised-routes**.

orf — displays outbound route filtering for the BGP instance. ORF (Outbound Route Filtering) is used to inform a neighbor of targets (using target-list) that it is willing to receive. This mechanism helps lessen the update exchanges between neighbors and saves CPU cycles to process routes that could have been received from the neighbor only to be dropped/ignored.

filter3 — displays path information for the specified IP address

Values **send** — displays the number of paths sent to this peer

receive — displays the number of paths received from this peer

graceful-restart — displays neighbors configured for graceful restart

Output The following outputs are examples of BGP neighbor information:

- BGP neighbor (standard and detailed) (Sample Output - BGP Neighbor (standard and detailed), Table 128)
- BGP neighbor (advertised and received) (Sample Output - BGP Neighbor (advertised-routes and received-routes), Table 129)
- BGP neighbor (graceful restart) (Sample Output - BGP Neighbor (graceful restart), Table 130)

Sample Output - BGP Neighbor (standard and detailed)

```
*A:ALU-12# show router 6 bgp neighbor
===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer  : 10.10.10.12
Group : ibgp_group
-------------------------------------------------------------------------------
Peer AS               : 65000            Peer Port            : 49550
Peer Address          : 10.10.10.12
Local AS              : 65000            Local Port           : 179
Local Address         : 10.10.10.1
Peer Type             : Internal
State                 : Established      Last State           : Established
Last Event            : recvKeepAlive
Last Error            : Cease
Local Family          : IPv4 VPN-IPv4
Remote Family         : IPv4 VPN-IPv4
Hold Time             : 90               Keep Alive           : 30
Active Hold Time      : 90               Active Keep Alive    : 30
Cluster Id            : None
Preference            : 170              Num of Flaps         : 0
Recd. Paths           : 19
IPv4 Recd. Prefixes   : 600              IPv4 Active Prefixes : 563
IPv4 Suppressed Pfxs  : 0                VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs   : 8656             VPN-IPv4 Active Pfxs : 8656
Mc IPv4 Recd. Pfxs.   : 0                Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs   : 0
Input Queue           : 0                Output Queue         : 0
i/p Messages          : 1141             o/p Messages         : 1041
i/p Octets            : 449029           o/p Octets           : 163814
i/p Updates           : 151              o/p Updates          : 50
TTL Security          : Enabled          Min TTL Value        : 10
Graceful Restart      : Disabled         Stale Routes Time    : n/a
Advertise Inactive    : Disabled         Peer Tracking        : Disabled
Auth key chain        : n/a
Bfd Enabled           : Enabled
Local Capability      : RouteRefresh MP-BGP
Remote Capability     : RouteRefresh MP-BGP
Import Policy         : None Specified / Inherited
Export Policy         : stmt1
-------------------------------------------------------------------------------
Neighbors : 1
===============================================================================
*A:ALU-12>#
```

```
*A:ALU-12# show router 6 bgp neighbor 10.10.10.12 detail
===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer  : 10.10.10.12
Group : iBGP
-------------------------------------------------------------------------------
Peer AS               : 65000            Peer Port            : 49550
Peer Address          : 10.10.10.12
Local AS              : 65000            Local Port           : 179
Local Address         : 10.10.10.1
Peer Type             : Internal
State                 : Established      Last State           : Established
Last Event            : recvKeepAlive
Last Error            : Cease
Local Family          : IPv4 VPN-IPv4
Remote Family         : IPv4 VPN-IPv4
Connect Retry         : 120              Local Pref.          : 70
Min Route Advt.       : 30               Min AS Orig.         : 15
Multihop              : 0 (Default)      AS Override          : Disabled
Damping               : Disabled         Loop Detect          : Ignore
MED Out               : No MED Out       Authentication       : None
Next Hop Self         : Disabled         AggregatorID Zero    : Disabled
Remove Private        : Disabled         Passive              : Disabled
Peer Identifier       : 10.10.10.12      Fsm Est. Trans       : 1
Fsm Est. Time         : 22h42m46s        InUpd Elap. Time     : 22h54m31s
Prefix Limit          : No Limit
Hold Time             : 90               Keep Alive           : 30
Active Hold Time      : 90               Active Keep Alive    : 30
Cluster Id            : None             Client Reflect       : Disabled
Preference            : 170              Num of Flaps         : 0
Recd. Paths           : 19
IPv4 Recd. Prefixes   : 600              IPv4 Active Prefixes : 563
IPv4 Suppressed Pfxs  : 0                VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs   : 8656             VPN-IPv4 Active Pfxs : 8656
Mc IPv4 Recd. Pfxs.   : 0                Mc IPv4 Active Pfxs. : 0
Mc IPv4 Suppr. Pfxs   : 0
Input Queue           : 0                Output Queue         : 0
i/p Messages          : 2881             o/p Messages         : 2777
i/p Octets            : 482089           o/p Octets           : 196798
i/p Updates           : 151              o/p Updates          : 50
TTL Security          : Enabled          Min TTL Value        : 10
Graceful Restart      : Disabled         Stale Routes Time    : n/a
Advertise Inactive    : Disabled         Peer Tracking        : Disabled
Auth key chain        : n/a
Bfd Enabled           : Enabled
Local Capability      : RouteRefresh MP-BGP
Remote Capability     : RouteRefresh MP-BGP
Import Policy         : None Specified / Inherited
Export Policy         : stmt1
-------------------------------------------------------------------------------
Neighbors : 1
===============================================================================
*A:ALU-12#
```

```
*A:ALU-12# show router 6 bgp neighbor 10.10.10.11 orf
===============================================================================
BGP Neighbor 10.0.0.11 ORF
===============================================================================
Send List (Automatic)
-------------------------------------------------------------------------------
target:65535:10
target:65535:20
===============================================================================
*A:ALU-12#

*A:ALU-12# show router 6 bgp neighbor 10.10.10.1 orf
===============================================================================
BGP Neighbor 10.0.0.1 ORF
===============================================================================
Receive List
-------------------------------------------------------------------------------
target:65535:10
target:65535:20
===============================================================================
*A:ALU-12#
```

Table 128: Show BGP Neighbor (Standard and Detailed) Output Fields

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
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

Table 128: Show BGP Neighbor (Standard and Detailed) Output Fields (Continued)

Label	Description
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
	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 session 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 connections with 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

Table 128: Show BGP Neighbor (Standard and Detailed) Output Fields (Continued)

Label	Description
Last Event	start: BGP has initialized the BGP neighbor
	stop: BGP has disabled the BGP neighbor
	open: BGP transport connection is opened
	close: BGP transport connection is closed
	openFail: BGP transport connection failed to open
	error: BGP transport connection error
	connectRetry: the connect retry timer expired
	holdTime: the hold time timer expired
	keepAlive: the keepalive timer expired
	recvOpen: BGP has received an OPEN message
	revKeepalive: BGP has received a KEEPALIVE message
	recvUpdate: BGP has received an UPDATE message
	recvNotify: BGP has received a NOTIFICATION message
	None: no events have occurred
Last Error	The last BGP error and subcode to occur on the BGP neighbor
Local Family	The configured local family value
Remote Family	The configured remote family value
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
Damping	Disabled: the BGP neighbor is configured not to dampen route flaps
	Enabled: the BGP neighbor is configured to dampen route flaps

Table 128: Show BGP Neighbor (Standard and Detailed) Output Fields (Continued)

Label	Description
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
MED Out	The configured or inherited MED value that is assigned to advertised routes
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 next hop in route updates to the specified neighbor
	Enabled: BGP will send only its own IP address as the BGP next hop 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
Peer Identifier	The IP identifier for the peer router
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

Table 128: Show BGP Neighbor (Standard and Detailed) Output Fields (Continued)

Label	Description
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
Preference	The configured route preference value for the peer group
Num of Flaps	The number of route flaps in the neighbor connection
Recd. Prefixes	The number of routes received from the BGP neighbor
Recd. Paths	The number of unique sets of path attributes received from the BGP neighbor
IPv4 Recd. Prefixes	The number of unique sets of IPv4 path attributes received from the BGP neighbor
IPv4 Active Prefixes	The number of IPv4 routes received from the BGP neighbor and active in the forwarding table
IPv4 Suppressed Pfxs	The number of unique sets of IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
VPN-IPv4 Suppr. Pfxs	The number of unique sets of VPN-IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
VPN-IPv4 Recd. Pfxs	The number of unique sets of VPN-IPv4 path attributes received from the BGP neighbor
VPN-IPv4 Active Pfxs	The number of VPN-IPv4 routes received from the BGP neighbor and active in the forwarding table
Mc IPv4 Recd. Pfxs	The number of unique sets of multiclass IPv4 path attributes received from the BGP neighbor
Mc IPv4 Active Pfxs	The number of multiclass IPv4 routes received from the BGP neighbor and active in the forwarding table
Mc IPv4 Suppr. Pfxs	The number of unique sets of multiclass IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
Input Queue	The number of BGP messages to be processed

Table 128: Show BGP Neighbor (Standard and Detailed) Output Fields (Continued)

Label	Description
Output Queue	The number of BGP messages to be transmitted
i/p Messages	The total number of packets received from the BGP neighbor
o/p Messages	The total number of packets sent to the BGP neighbor
i/p Octets	The total number of octets received from the BGP neighbor
o/p Octets	The total number of octets sent to the BGP neighbor
i/p Updates	The total number of updates received from the BGP neighbor
o/p Updates	The total number of updates sent to the BGP neighbor
TTL Security	Enabled: TTL security is enabled Disabled: TTL security is disabled
Min TTL Value	The minimum TTL value configured for the peer
Graceful Restart	The state of graceful restart
Stale Routes Time	The length of time that stale routes are kept in the route table
Auth key chain	The value for the authentication key chain
Bfd Enabled	Enabled: BFD is enabled Disabled: BFD is disabled
Local Capability	The capability of the local BGP speaker; for example, route refresh, MP-BGP, ORF
Remote Capability	The capability of the remote BGP peer; for example, route refresh, MP-BGP, ORF
Export Policy	The configured export policies for the peer group
Import Policy	The configured import policies for the peer group

Sample Output - BGP Neighbor (advertised-routes and received-routes)

```
*A:ALU-12# show router 6 bgp neighbor 44.44.44.44 advertised-routes
===============================================================================
 BGP Router ID : 55.55.55.55       AS : 1       Local AS : 1
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                              LocalPref   MED
      Nexthop                                                         VPN Label
      As-Path
-------------------------------------------------------------------------------
? 10.0.0.02/32 100 none
10.0.0.16  -
No As-Path
? 10.0.6.04/24 100 none
10.0.0.16-
No As-Path
-------------------------------------------------------------------------------
Routes : 2
===============================================================================
*A:ALU-12#

*A:ALU-12# show router 6 bgp neighbor 10.10.10.12 advertised-routes brief
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network
-------------------------------------------------------------------------------
?     10.10.10.1/32
?     12.12.1.0/24
?     20.0.0.0/24
?     21.0.0.0/24
?     88.88.1.0/24
-------------------------------------------------------------------------------
Routes : 5
===============================================================================
*A:ALU-12#

*A:ALU-12# show router 6 bgp neighbor 44.44.44.44 received-routes
===============================================================================
 BGP Router ID : 55.55.55.55       AS : 1       Local AS : 1
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
```

```
BGP IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop                                                      VPN Label
      As-Path
-------------------------------------------------------------------------------
? 10.0.0.16/32 100 none
10.0.0.16  -
No As-Path
? 10.0.6.0/24  100 none
10.0.0.16-
No As-Path
? 10.0.8.0/24  100 none
10.0.0.16-
No As-Path
? 10.0.12.0/24 100 none
10.0.0.16-
No As-Path
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
*A:ALU-12#
```

Table 129: Show BGP Neighbor (Advertised-Routes and Received-Routes) Output Fields

Label	Description
BGP Router ID	The local BGP router ID
AS	The configured autonomous system number
Local AS	The configured local AS setting. If not configured, then it is the same value as the AS.
Flag(s)	Legend: Status codes: u - used s - suppressed h - history d - decayed * - valid If an * is not present, then the status is invalid Origin codes: i - IGP e - EGP ? - incomplete > - best
Network	The route IP prefix and mask length for the route

Table 129: Show BGP Neighbor (Advertised-Routes and Received-Routes) Output Fields (Continued)

Label	Description
Next Hop	The BGP next hop for the route
LocalPref	The BGP local preference path attribute for the route
MED	The BGP Multi-Exit Discriminator (MED) path attribute for the route
AS-Path	The BGP AS path for the route

Sample Output - BGP Neighbor (graceful restart)

```
*A:ALU-12# show router 6 bgp neighbor 20.10.120.44 graceful-restart
===============================================================================
BGP Neighbor 20.10.120.44 Graceful Restart
===============================================================================
Graceful Restart locally configured for peer  : Enabled
Peer's Graceful Restart feature               : Enabled
NLRI(s) that peer supports restart for        : IPv4-Unicast IPv4-MPLS IPv4-VPN
NLRI(s) that peer saved forwarding for        : IPv4-Unicast IPv4-MPLS IPv4-VPN
NLRI(s) that restart is negotiated for        : None
NLRI(s) of received end-of-rib markers        : IPv4-Unicast
NLRI(s) of all end-of-rib markers sent        : IPv4-Unicast
Restart time locally configured for peer      : 120 seconds
Restart time requested by the peer            : 390 seconds
Time stale routes from peer are kept for      : 360 seconds
Graceful restart status on the peer           : Not currently being helped
Number of Restarts                            : 328
Last Restart at                               : 08/20/2006 12:22:06
===============================================================================
*A:ALU-12#
```

Table 130: BGP Neighbor (Graceful Restart) Output Fields

Label	Description
BGP Neighbor	The IP address of the BGP neighbor
Graceful Restart locally configured for peer	The configured state of graceful restart for the local router
Peer's Graceful Restart feature	The configured state of graceful restart for the peer router
NLRI(s) that peer supports restart for	The families supported by the peer router for graceful restart
NLRI(s) that peer saved forwarding for	The families for which the peer router continued to forward packets after graceful restart

Table 130: BGP Neighbor (Graceful Restart) Output Fields (Continued)

Label	Description
NLRI(s) that restart is negotiated for	The families that negotiate restart during graceful restart
NLRI(s) of received end-of-rib markers	The families for which end-of-RIB markers have been received
NLRI(s) of all end-of-rib markers sent	The families for which end-of-RIB markers have been sent
Restart time locally configured for peer	The length of time configured on the local router for the peer router's graceful restart
Restart time requested by the peer	The length of time requested by the peer router for graceful restart
Time stale routes from peer are kept for	The length of time that the local router continues to support stale routes
Graceful restart status on the peer	The status of graceful restart on the peer router
Number of Restarts	The number of restarts since graceful restart is enabled between peers
Last Restart at	The local time of the last graceful restart

next-hop

Syntax **next-hop** [*family*] [*ip-address*] [**detail**]

Context show>router>bgp

Description This command displays BGP next-hop information.

Parameters *family* — the type of routing information to be distributed by the BGP instance

Values **ipv4** — displays only those BGP peers that have the IPv4 family enabled and not those capable of exchanging IP-VPN routes

ip-address — displays the next hop information for the specified IP address

Values *ipv4-address*: a.b.c.d (host bits must be 0)

detail — displays the more detailed version of the output

Output The following output is an example of BGP next-hop information, and Table 131 describes the fields.

Sample Output

```
*A:ALU-12# show router 6 bgp next-hop
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================
===============================================================================
BGP Next Hop
===============================================================================
Next Hop                                                        Pref Owner
   Resolving Prefix                                                  Metric
   Resolved Next Hop                                                 Ref. Count
-------------------------------------------------------------------------------
10.10.10.12                                                     15   ISIS
   0.0.0.0/0                                                         10
   88.88.1.2                                                         592
10.10.10.12                                                     15   ISIS
   0.0.0.0/0                                                         10
   88.88.2.2                                                         592
27.0.0.1                                                        15   ISIS
   27.0.0.0/24                                                       20
   88.88.1.2                                                         8
27.0.0.1                                                        15   ISIS
   27.0.0.0/24                                                       20
   88.88.2.2                                                         8
-------------------------------------------------------------------------------
Next Hops : 2
===============================================================================
*A:ALU-12#

*A:ALU-12# show router 6 bgp next-hop 27.0.0.1
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================
===============================================================================
BGP Next Hop
===============================================================================
Next Hop                                                        Pref Owner
   Resolving Prefix                                                  Metric
   Resolved Next Hop                                                 Ref. Count
-------------------------------------------------------------------------------
27.0.0.1                                                        15   ISIS
   27.0.0.0/24                                                       20
   88.88.1.2                                                         8
27.0.0.1                                                        15   ISIS
   27.0.0.0/24                                                       20
   88.88.2.2                                                         8
-------------------------------------------------------------------------------
Next Hops : 1
===============================================================================
*A:ALU-12#
```

```
*A:ALU-12# show router 6 bgp next-hop 27.0.0.1 detail
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================
===============================================================================
BGP Next Hop
===============================================================================
Next Hop: 27.0.0.1
-------------------------------------------------------------------------------
Resolving Prefix : 27.0.0.0/24
Preference       : 15                   Metric           : 20
Reference Count  : 8                    Owner            : ISIS
Resolved Next Hop: 88.88.1.2
Egress Label     : N/A
Resolved Next Hop: 88.88.2.2
Egress Label     : N/A
Resolved Next Hop: 88.88.3.2
Egress Label     : N/A
-------------------------------------------------------------------------------
Next Hops : 1
===============================================================================
*A:ALU-12>#
```

Table 131: Show BGP Next-Hop Output Fields

Label	Description
BGP Router ID	The local BGP router ID
AS	The configured autonomous system number
Local AS	The configured local AS setting. If not configured, then the value is the same as the AS.
Next Hop	The next-hop address
Resolving Prefix	The prefix of the best next hop
Pref: Preference	The BGP preference attribute for the routes
Metric	The metric derived from the IGP for a particular next hop
Reference Count	The number of routes using the resolving prefix
Owner	The routing protocol used to derive the best next hop
Resolved Next Hop	The IP address of the next hop
Egress Label	The VPN label used for VPN-IPv4 data
Next Hops	The number of next hops

paths

Syntax **paths**

Context show>router>bgp

Description This command displays a summary of BGP path attributes.

Output The following output is an example of BGP path information, and Table 132 describes the fields.

Sample Output

```
*A:ALU-12# show router 6 bgp paths
===============================================================================
 BGP Router ID : 55.55.55.55 AS : 65000   Local AS : 65000
===============================================================================
===============================================================================
BGP Paths
===============================================================================
Path: No As-Path
-------------------------------------------------------------------------------
Next Hop         : 44.44.10.12
Origin           : Incomplete           Segments         : 0
MED              : None                 Local Preference : 4294967295
Refs             : 1080                 ASes             : 0
Flags            : IBGP-learned
-------------------------------------------------------------------------------
Path: No As-Path
-------------------------------------------------------------------------------
Next Hop         : 88.88.1.2
Origin           : IGP                  Segments         : 0
MED              : 10                   Local Preference : None
Refs             : 4                    ASes             : 0
Flags            : Imported
-------------------------------------------------------------------------------
Path: No As-Path
-------------------------------------------------------------------------------
Next Hop         : 44.44.10.21
Origin           : IGP                  Segments         : 0
MED              : None                 Local Preference : 100
Refs             : 1082                 ASes             : 0
Flags            : IBGP-learned
Cluster          : 10.10.10.12
Originator Id    : 10.10.10.21
-------------------------------------------------------------------------------
Paths : 3
===============================================================================
*A:ALU-12#
```

Table 132: Show BGP Path Output Fields

Label	Description
BGP Router ID	The local BGP router ID
AS	The configured autonomous system number
Local AS	The configured local AS setting. If not configured, then the value is the same as the AS.
Path	The AS path attribute
Next Hop	The advertised BGP next hop
Origin	EGP: the NLRI is learned by an EGP protocol
	IGP: the NLRI is interior to the originating AS
	Incomplete: NLRI was learned another way
Segments	The number of segments in the AS path attribute
MED	The Multi-Exit Discriminator value
Local Preference	The local preference value. This value is used if the BGP route arrives from a BGP peer without the Local Pref attribute set. It is overridden by any value set via a route policy.
Refs	The number of routes using a specified set of path attributes
ASes	The number of autonomous system numbers in the AS path attribute
Flags	IBGP-learned: path attributes learned by an IBGP peering
Community	The BGP community attribute list
Cluster List	The route reflector cluster list
Originator ID	The originator ID path attribute value

routes

Syntax **routes** [*family*] [**received**] [**url** *file-url*]
routes [*family*] [**brief**]
routes [*family*] *prefix* [**detail** | **longer** | **hunt** [**brief**]]
routes [*family*] **community** *comm-id*
routes [*family*] **aspath-regex** *reg-ex*

Context show>router>bgp

Description This command displays BGP route information.

When this command is issued without any parameters, the entire BGP routing table displays.

When this command is issued with an IP prefix/mask or IP address, the best match for the parameter displays.

Parameters *family* — the type of routing information to be distributed by the BGP instance

Values **ipv4** — displays only those BGP peers that have the IPv4 family enabled and not those capable of exchanging IP-VPN routes

vpn-ipv4 — displays the BGP peers that are IP-VPN capable

prefix — the type of routing information to display

Values *rd* | [*rd*]:[*ip-prefix*[/*ip-prefix-length*]]

where

rd: ip-addr:*comm-val* | *2byte-asnumber*:*ext-comm-val* | *4-byte-asnumber*:*comm-val*

where

ip-addr	a.b.c.d
comm-val:	0 to 65535
2-byte-asnumber	0 to 65535
ext-comm-val	0 to 4294967295
4-byte-asnumber	*asn1*.*asn2* (two 2-byte pieces)
	asn1: 1 to 65535
	asn2: 0 to 65535
ip-prefix	a.b.c.d
ip-prefix-length:	0 to 32

received — generates a bulk report (or file) for all received BGP routes

url *file-url* — provides information on the BGP routes stored in the URL file

Values		
	local-url:	[*cflash-id/*][*file-path*] 255 characters maximum
	cflash-id	cf3: | cf3-A: | cf3-B:
	remote-url:	[{*ftp://* | *tftp://*}*login*:*pswd*@*remote-locn*/][*file-path*] 255 characters maximum
	remote-locn:	*hostname* | *ipv4-address*
	ipv4-address	a.b.c.d

brief — provides a summarized display of the set of peers to which a BGP route is advertised

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 — displays the more detailed version of the output

community *comm-id* — displays all routes with the specified BGP community

Values		
	comm-id:	[*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 to 65535
	comm-val1:	0 to 65535
	type:	**target**, **origin** (keywords)
	ip-address:	a.b.c.d
	comm-va 2	0 to 4294967295
	as-number2	0 to 4294967295
	well-known-comm:	**no-export**, **no-export-subconfed**, **no-advertise** (keywords)

aspath-regex *reg-exp* — displays all routes with an AS path matching the specified regular expression *reg-exp*

Output The following output is an example of BGP route information, and Table 133 describes the fields.

Sample Output

```
*A:ALU-12# show router 6 bgp routes
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP IPv4 Routes
```

```
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop                                                      VPN Label
      As-Path
-------------------------------------------------------------------------------
u*>?  10.10.10.12/32                                   100         None
      10.10.10.12                                                  -
      No As-Path

u*>?  11.11.1.0/24                                     100         None
      10.10.10.12                                                  -
      No As-Path

*?    23.0.0.0/24                                      100         None
      10.10.10.12                                                  -
      No As-Path

*?    24.0.0.0/24                                      100         None
      10.10.10.12                                                  -
      No As-Path
-------------------------------------------------------------------------------
Routes : 4
===============================================================================


*A:ALU-12# show router 6 bgp routes brief
===============================================================================
 BGP Router ID : 10.10.10.1        AS : 65000   Local AS : 65000
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network
-------------------------------------------------------------------------------
u*>?  10.10.10.12/32
u*>?  11.11.1.0/24
*?    23.0.0.0/24
*?    24.0.0.0/24
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
*A:ALU-12#


*A:ALU-12# show router 6 bgp routes 13.1.0.0/24 detail
===============================================================================
BGP Router ID : 10.128.0.161 AS : 65535 Local AS : 65535
===============================================================================
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
===============================================================================
BGP Routes
===============================================================================
Original Attributes
Network        : 13.1.0.0/24         Nexthop : 10.20.1.20
Route Dist.    : 10070:100           VPN Label : 152784
```

```
From           : 10.20.1.20          Res. Nexthop: 10.130.0.2
Local Pref.    : 100
Aggregator AS  : none                Aggregator: none
Atomic Aggr.   : Not Atomic          MED : none
Community      : target:10070:1
Cluster        : No Cluster Members
Originator Id  : None                Peer Router Id: 10.20.1.20
Flags          : Used Valid Best IGP
AS-Path        : 10070 {14730}

Modified Attributes

Network        : 13.1.0.0/24         Nexthop : 10.20.1.20
Route Dist.    : 10001:100           VPN Label : 152560
From           : 10.20.1.20          Res. Nexthop : 10.130.0.2
Local Pref.    : 100
Aggregator AS  : none                Aggregator: none
Atomic Aggr.   : Not Atomic          MED : none
Community      : target:10001:1
Cluster        : No Cluster Members
Originator Id  : None                Peer Router Id: 10.20.1.20
Flags          : Used Valid Best IGP
AS-Path        : No As-Path
-------------------------------------------------------------------------------
===============================================================================


*A:ALU-12# show router 6 bgp routes 100.0.0.0/30 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        : 100.0.0.0/30
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:ALU-12*
```

Table 133: Show BGP Route Output Fields

Label	Description
BGP Router ID	The local BGP router ID
AS	The configured autonomous system number
Local AS	The configured local AS setting. If not configured, then the value is the same as the AS.
Flag(s)	Legend: Status codes: u - used s - suppressed h - history d - decayed * - valid If an * is not present, then the status is invalid Origin codes: i - IGP e - EGP ? - incomplete > - best
Network	The IP prefix and mask length
Nexthop	The BGP next hop
AS-Path	The BGP AS path attribute
Local Pref.	The local preference value. This value is used if the BGP route arrives from a BGP peer without the Local Pref attribute set. It is overridden by any value set via a route policy.
MED	The MED metric value
	none: MED metrics are present
VPN Label	The label generated by the PE's label manager
Original Attributes	The received BGP attributes of a route from a peer without any modification from any policy
Modified Attributes	The final BGP attributes of a route after the policies evaluation
Route Dist.	The route distinguisher identifier attached to routes that distinguishes the VPN it belongs to
From	The advertising BGP neighbor's IP address

Table 133: Show BGP Route Output Fields (Continued)

Label	Description
Res. Nexthop	The resolved next hop
Aggregator AS	The aggregator AS value
	none: aggregator AS attributes are not present
Aggregator	The aggregator attribute value
	none: aggregator attributes are not present
Atomic Aggr.	Atomic: the atomic aggregator flag is set
	Not Atomic: the atomic aggregator flag is not set
Community	The BGP community attribute list
Cluster	The route reflector cluster list
Originator Id	The originator ID path attribute value
	none: the originator ID attribute is not present
Peer Router Id	The router ID of the advertising router
VPRN Imported	The VPRNs where a particular BGP-VPN received route has been imported and installed

summary

Syntax **summary** [**all**]
summary [**family** *family*] [**neighbor** *ip-address*]

Context show>router>bgp

Description 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 shut down 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 *family* — the type of routing information to be distributed by the BGP instance

Values **ipv4** — displays only those BGP peers that have the IPv4 family enabled

vpn-ipv4 — displays the BGP peers that are IP-VPN capable

ip-address — clears damping information for entries received from the BGP neighbor

Values *ipv4-address*: a.b.c.d

Output The following output is an example of BGP summary information, and Table 134 describes the fields.

Sample Output

```
*A:ALU-12# show router 6 bgp summary
===============================================================================
 BGP Router ID : 55.55.55.1        AS : 65000   Local AS : 65000
===============================================================================
BGP Admin State         : Up          BGP Oper State              : Up
Total Peer Groups       : 1           Total Peers                 : 1
Total BGP Paths         : 74          Total Path Memory           : 9128
Total IPv4 Remote Rts   : 600         Total IPv4 Rem. Active Rts  : 563
Total Supressed Rts     : 0           Total Hist. Rts             : 0
Total Decay Rts         : 0
Total VPN Peer Groups   : 0           Total VPN Peers             : 0
Total VPN Local Rts     : 8672
Total VPN-IPv4 Rem. Rts : 8656        Total VPN-IPv4 Rem. Act. Rts: 8656
Total VPN Supp. Rts     : 0           Total VPN Hist. Rts         : 0
Total VPN Decay Rts     : 0
===============================================================================
BGP Summary
===============================================================================
Neighbor
                  AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                     PktSent OutQ
-------------------------------------------------------------------------------
44.44.10.12
                65000     654    0 04h11m01s 600/563/569 (IPv4)
                          557    0           8656/8656/8672 (VpnIPv4)
===============================================================================
*A:ALU-12#

*A:ALU-12# show router 6 bgp summary all
===============================================================================
BGP Summary
===============================================================================
Neighbor
ServiceId          AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                     PktSent OutQ
-------------------------------------------------------------------------------
44.44.10.12
Def. Instance   65000     662    0 04h14m52s 600/563/569 (IPv4)
                          564    0           8656/8656/8672 (VpnIPv4)
===============================================================================
*A:ALU-12#
```

```
*A:ALU-12# show router 6 bgp summary neighbor 44.44.10.12
===============================================================================
 BGP Router ID : 44.44.10.1        AS : 65000   Local AS : 65000
===============================================================================
BGP Admin State         : Up          BGP Oper State              : Up
Total Peer Groups       : 1           Total Peers                 : 1
Total BGP Paths         : 74          Total Path Memory           : 9128
Total IPv4 Remote Rts   : 600         Total IPv4 Rem. Active Rts  : 563
Total Supressed Rts     : 0           Total Hist. Rts             : 0
Total Decay Rts         : 0

Total VPN Peer Groups   : 0           Total VPN Peers             : 0
Total VPN Local Rts     : 8672
Total VPN-IPv4 Rem. Rts : 8656        Total VPN-IPv4 Rem. Act. Rts: 8656
Total VPN Supp. Rts     : 0           Total VPN Hist. Rts         : 0
Total VPN Decay Rts     : 0

===============================================================================
BGP Summary
===============================================================================
Neighbor
                   AS PktRcvd InQ  Up/Down   State|Rcv/Act/Sent (Addr Family)
                      PktSent OutQ
-------------------------------------------------------------------------------
44.44.10.12
                65000     673    0 04h20m24s 600/563/569 (IPv4)
                          575    0           8656/8656/8672 (VpnIPv4)
===============================================================================
*A:ALU-12#

*A:ALU-12# show router 6 bgp summary family ipv4
===============================================================================
 BGP Router ID : 44.44.10.1        AS : 65000   Local AS : 65000
===============================================================================
BGP Admin State         : Up          BGP Oper State              : Up
Total Peer Groups       : 1           Total Peers                 : 1
Total BGP Paths         : 74          Total Path Memory           : 9128
Total IPv4 Remote Rts   : 600         Total IPv4 Rem. Active Rts  : 563
Total Supressed Rts     : 0           Total Hist. Rts             : 0
Total Decay Rts         : 0

Total VPN Peer Groups   : 0           Total VPN Peers             : 0
Total VPN Local Rts     : 8672
Total VPN-IPv4 Rem. Rts : 8656        Total VPN-IPv4 Rem. Act. Rts: 8656
Total VPN Supp. Rts     : 0           Total VPN Hist. Rts         : 0
Total VPN Decay Rts     : 0
===============================================================================
BGP IPv4 Summary
===============================================================================
Neighbor
                   AS PktRcvd PktSent  InQ OutQ Up/Down   State|Recv/Actv/Sent
-------------------------------------------------------------------------------
44.44.10.12
                65000     679     581    0    0 04h23m36s 600/563/569
===============================================================================
*A:ALU-12#
```

Table 134: Show BGP Summary Output Fields

Label	Description
BGP Router ID	The local BGP router ID
AS	The configured autonomous system number
Local AS	The configured local AS setting. If not configured, then the value is the same as the AS.
BGP Admin State	Down: BGP is administratively disabled
	Up: BGP is administratively enabled
BGP Oper State	Down: BGP is operationally disabled
	Up: BGP is operationally enabled
Total Peer Groups	The total number of configured BGP peer groups
Total Peers	The total number of configured BGP peers
Total BGP Paths	The total number of unique sets of BGP path attributes learned from BGP peers
Total Path Memory	The total amount of memory used to store the path attributes
Total IPv4 Remote Rts	The total number of IPv4 routes learned from BGP peers
Total IPv4 Remote Act. Rts	The total number of IPv4 routes used in the forwarding table
Total Supressed Rts	The total number of suppressed routes due to route damping
Total Hist. Rts	The total number of routes with history due to route damping
Total Decay Rts	The total number of decayed routes due to route damping
Total VPN Peer Groups	The total number of configured VPN peer groups
Total VPN Peers	The total number of configured VPN peers
Total VPN Local Rts	The total number of configured local VPN routes
Total VPN-IPv4 Rem. Rts	The total number of configured remote VPN-IPv4 routes
Total VPN-IPv4 Rem. Act. Rts	The total number of active remote VPN-IPv4 routes used in the forwarding table
Total VPN Supp. Rts	The total number of suppressed VPN routes due to route damping
Total VPN Hist. Rts	The total number of VPN routes with history due to route damping
Total VPN Decay Rts	The total number of decayed routes due to route damping

Table 134: Show BGP Summary Output Fields (Continued)

Label	Description
Neighbor	The BGP neighbor address
AS (Neighbor)	The BGP neighbor autonomous system number
PktRcvd	The total number of packets received from the BGP neighbor
PktSent	The total number of packets sent to the BGP neighbor
InQ	The number of BGP messages to be processed
OutQ	The number of BGP messages to be transmitted
Up/Down	The amount of time that the BGP neighbor has either been established or not established depending on its current state
State|Recv/Actv/Sent (Addr Family)	The BGP neighbor's current state (if not established) or the number of received routes, active routes and sent routes (if established), along with the address family

statistics

Syntax **statistics** [*ip-int-name* | *ip-address*]

Context show>router>dhcp

Description This command displays statistics for DHCP Relay.

If no IP address or interface name is specified, then all configured interfaces are displayed.

If an IP address or interface name is specified, then only data regarding the specified interface is displayed.

Parameters *ip-int-name* | *ip-address* — displays statistics for the specified IP interface

Output The following output is an example of DHCP statistics information, and Table 135 describes the fields.

Sample Output

```
*A:ALU-1# show router dhcp statistics
===============================================================
DHCP Global Statistics (Service: x)
===============================================================
Rx Packets                          : 0
Tx Packets                          : 0
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 0
Client Packets Discarded            : 0
Client Packets Relayed              : 0
Server Packets Discarded            : 0
```

```
Server Packets Relayed               : 0
===============================================================
*A:ALU-1#
```

Table 135: Show DHCP Statistics Output Fields

Label	Description
DHCP Global Statistics (Service: x)	
Rx Packets	The number of packets received
Tx Packets	The number of packets transmitted
Rx Malformed Packets	The number of malformed packets received
Rx Untrusted Packets	The number of untrusted packets received
Client Packets Discarded	The number of packets from the DHCP client that were discarded
Client Packets Relayed	The number of packets from the DHCP client that were forwarded
Server Packets Discarded	The number of packets from the DHCP server that were discarded
Server Packets Relayed	The number of packets from the DHCP server that were forwarded

summary

Syntax **summary**

Context show>router>dhcp

Description This command displays a summary of DHCP configuration.

Output The following output is an example of DHCP summary information, and Table 136 describes the fields.

Sample Output

```
A:ALU-48# show router 6 dhcp summary
===============================================================================
DHCP Summary (Service: 6)
===============================================================================
Interface Name                   Arp      Used/                Info    Admin
  SapId/Sdp                      Populate Provided             Option  State
-------------------------------------------------------------------------------
vprn_interface                   No       0/0                  Keep    Down
  sap:1/5/2                               0/0
-------------------------------------------------------------------------------
Interfaces: 1
===============================================================================
A:ALU-48#
```

Table 136: Show DHCP Summary Output Fields

Label	Description
DHCP Summary (Service: x)	
Interface Name SapID/ Sdp	The name of the interface
Arp Populate	Specifies whether ARP populate is enabled or disabled
Used/Provided:	Used—the number of lease-states that are currently in use on the specified interface; that is, the number of clients on the interface that got an IP address by DHCP. This number is always less than or equal to the "Provided" field.
	Provided—lease-populate value configured for the specified interface
Info Option	Keep—the existing, information is kept on the packet and the router does not add any additional information
	Replace—on ingress, the existing information-option is replaced with the information-option from the router
	Drop—the packet is dropped and an error is logged
Admin State	The administrative state

interface

Syntax **interface** [{[*ip-address* | *ip-int-name*] [**detail**]}] | [**summary**] | [**exclude-services**]

Context show>router

Description This command displays the router IP interface table sorted by interface index.

Parameters *ip-address* — the IP address of the interface for which to display information

ip-int-name — the IP interface name for which to display information

detail — displays detailed IP interface information for the router

summary — displays summary IP interface information for the router

exclude-services — displays IP interface information, excluding IP interfaces configured for customer services. Only core network IP interfaces are displayed.

Output The following output is an example of standard IP interface information, and Table 137 describes the fields.

Sample Output

```
*A:ALU-12# show router 6 interface
===============================================================================
Interface Table (Service: 6)
===============================================================================
Interface-Name                   Adm         Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                    PfxState
-------------------------------------------------------------------------------
vprn_interface                   Up          Down/Down   VPRN    1/5/2
   -                                                             -
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
*A:ALU-12#
```

Table 137: Show IP Interface Output Fields

Label	Description
Interface Table (Service: x)	
Interface-Name	The name of the interface
IP-Address	The IP address of the interface
Adm	The administrative state of the interface
Opr (v4/v6)	The operational state of the interface (only ipv4 addresses apply)
Type	The service type
Port/SapId PfxState	The port or SAP associated with the interface

route-table

Syntax **route-table** [*ip-prefix /*[*mask*] **[longer]** | [**protocol** *protocol*] | [**summary**]]

Context show>router

Description 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]* — the IP prefix and optional mask for which to display routes

longer — displays routes matching the ip-prefix/mask and routes with longer masks

protocol **—** displays routes learned from the specified protocol

Values bgp, isis, local, ospf, rip, static, aggregate

summary — displays a route table summary information

Output The following output is an example of route table information, and Table 138 describes the fields.

Sample Output

```
*A:ALU-12# show router 6 route-table
===============================================================================
Route Table (Service: 6)
===============================================================================
Dest Prefix                                    Type    Proto    Age         Pref
       Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.10.0.1/32                                   Remote  OSPF    65844        10
       10.10.13.1                                                    1001
10.10.0.2/32                                   Remote  OSPF    65844        10
       10.10.13.1                                                    2001
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of Routes: 2
===============================================================================
*A:ALU-12#
```

Table 138: Show Route Table Output Fields

Label	Description
Dest Prefix	The route destination address and mask
Next Hop	The next hop IP address for the route destination
Type	Local—the route is a local route
	Remote—the route is a remote route
Proto	The protocol through which the route was learned
Age	The route age in seconds for the route
Metric	The route metric value for the route
Pref	The route preference value for the route
No. of Routes:	The number of routes displayed in the list

sgt-qos

Syntax **sgt-qos**

Context show>router

Description This command displays QoS information about self-generated traffic.

Parameters *service-id* — specifies the service identifier of the service

Values 1 to 2147483647 or *service-name*

application

Syntax **application** [*app-name*] [**dscp** | **dot1p**]

Context show>router>sgt-qos

Description This command displays application QoS settings.

Parameters *app-name* — the specified application

Values arp, bgp, dhcp, dns, ftp, icmp, isis, ldp, ntp, ospf, ptp, radius, rsvp, snmp, snmp-notification, ssh, syslog, tacplus, telnet, tftp, traceroute

Note:

- PTP in the context of SGT QoS is defined as Precision Timing Protocol and is an application in the 7705 SAR. The PTP application name is also used in areas such as event-control and logging. Precision Timing Protocol is defined in IEEE 1588-2008.
- PTP in the context of IP filters is defined as Performance Transparency Protocol. IP protocols can be used as IP filter match criteria; the match is made on the 8-bit protocol field in the IP header.

dscp — specifies to show all DSCP applications

dot1p — specifies to show all dot1p applications

Output The following output is an example of application QoS information, and Table 139 describes the fields.

Sample Output

```
*A:ALU-1>show# router 6 sgt-qos application
===============================================================================
DSCP Application Values
===============================================================================
Application         DSCP Value                    Default DSCP Value
-------------------------------------------------------------------------------
bgp                 none                          none
dhcp                none                          none
dns                 none                          none
ftp                 none                          none
icmp                none                          none
ldp                 none                          none
ntp                 none                          none
ospf                none                          none
radius              none                          none
rsvp                none                          none
snmp                none                          none
snmp-notification   none                          none
ssh                 none                          none
syslog              none                          none
tacplus             none                          none
telnet              none                          none
```

```
tftp                none                            none
traceroute          none                            none
ptp                 none                            none
===============================================================================

===============================================================================
Dot1p Application Values
===============================================================================
Application         Dot1p Value                     Default Dot1p Value
-------------------------------------------------------------------------------
arp                 none                            none
isis                none                            none
===============================================================================
*A:ALU-1>show#
```

Table 139: Application QoS Output Fields

Label	Description
Application	The DSCP or dot1p application
DSCP Value	The DSCP name or value assigned to the application; if you assign a value to the application (0 to 63), the DSCP name that maps to the value is displayed
Default DSCP Value	The default DSCP value
Dot1p Value	The dot1p priority assigned to the application (applies only to ARP and IS-IS)
Default Dot1p Value	The default dot1p value

dscp-map

Syntax **dscp-map** [*dscp-name*]

Context show>router>sgt-qos

Description This command displays the DSCP-to-FC mappings.

Parameters *dscp-name* — the specified DSCP name.

Values be | ef | | cp1 | cp2 | cp3 | cp4 | cp5 | cp6 | cp7 | cp9 | cs1 | cs2 | cs3 | cs4 | cs5 | nc1 | nc2 | af11 | af12 | af13 | af21 | af22 | af23 | af31 | af32 | af33 | af41 | af42 | af43 | cp11 | cp13 | cp15 | cp17 | cp19 | cp21 | cp23 | cp25 | cp27 | cp29 | cp31 | cp33 | cp35 | cp37 | cp39 | cp41 | cp42 | cp43 | cp44 | cp45 | cp47 | cp49 | cp50 | cp51 | cp52 | cp53 | cp54 | cp55 | cp57 | cp58 | cp59 | cp60 | cp61 | cp62 | cp63

Output The following output is an example of DSCP-to-FC mapping information, and Table 140 describes the fields.

Sample Output

```
A:ALU-1# show router 6 sgt-qos dscp-map
===============================================================================
DSCP to FC Mappings
===============================================================================
DSCP Value          FC Value            Default FC Value
-------------------------------------------------------------------------------
be                  nc                  nc
cp1                 be                  be
cp2                 be                  be
cp3                 be                  be
cp4                 be                  be
cp5                 be                  be
cp6                 be                  be
cp7                 be                  be
cs1                 be                  be
cp9                 be                  be
af11                af                  af
cp11                be                  be
af12                af                  af
cp13                be                  be
af13                af                  af
cp15                be                  be
cs2                 be                  be
cp17                be                  be
af21                l1                  l1
cp19                be                  be
af22                l1                  l1
cp21                be                  be
af23                l1                  l1
cp23                be                  be
cs3                 be                  be
cp25                be                  be
af31                l1                  l1
cp27                be                  be
af32                l1                  l1
cp29                be                  be
af33                l1                  l1
cp31                be                  be
cs4                 be                  be
cp33                be                  be
af41                nc                  nc
cp35                be                  be
af42                af                  h2
cp37                be                  be
af43                h2                  h2
cp39                be                  be
cs5                 be                  be
cp41                be                  be
cp42                be                  be
cp43                be                  be
cp44                be                  be
cp45                be                  be
ef                  ef                  ef
cp47                be                  be
nc1                 nc                  nc
cp49                be                  be
cp50                h2                  h2
```

```
cp51                be                  be
cp52                be                  be
cp53                be                  be
cp54                be                  be
cp55                be                  be
nc2                 nc                  nc
cp57                be                  be
cp58                be                  be
cp59                be                  be
cp60                be                  be
cp61                be                  be
cp62                be                  be
cp63                be                  be
===============================================================================
A:ALU-1#
```

Table 140: DSCP-to-FC Mapping Output Fields

Label	Description
DSCP Value	The DSCP values (displayed as names) of the self-generated traffic
FC Value	The FC value mapped to each DSCP value
Default FC Value	The default FC value

static-arp

Syntax **static-arp** [*ip-address* | *ip-int-name* | **mac** *ieee-mac-addr*]

Context show>router

Description 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* — the IP address for which static ARP entries are displayed

ip-int-name — the interface name for which static ARP entries are displayed

ieee-mac-addr — the MAC address for which static ARP entries are displayed

Output The following output is an example of static ARP table information, and Table 141 describes the fields.

Sample Output

```
*A:ALU-12# show router 6 static-arp
===============================================================================
ARP Table (Service: 6)
===============================================================================
IP Address      MAC Address       Expiry    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:ALU-12#
```

Table 141: Show Static ARP Table Output Fields

Label	Description
IP Address	The IP address of the static ARP entry
MAC Address	The MAC address of the static ARP entry
Expiry	The age of the ARP entry. Static ARPs always have 00:00:00 for the age.
Type	Inv—the ARP entry is an inactive static ARP entry (invalid)
	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

static-route

Syntax **static-route** [*ip-prefix /mask*] | [**preference** *preference*] | [**next-hop** *ip-address* [**detail**]

Context show>router

Description 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 only the static routes matching the specified IP prefix and mask

preference — displays static routes with the specified route preference

Values 0 to 65535

ip-address — displays only the static routes with the specified next-hop IP address

detail — displays detailed information about the static route

Output The following output is an example of static route information, and Table 142 describes the fields.

Sample Output

```
*A:ALU-12# show router 6 static-route
===============================================================================
Static Route Table (Service: 6)  Family: IPv4
===============================================================================
Prefix                                    Tag         Met    Pref Type Act
   Next Hop                                Interface
-------------------------------------------------------------------------------
192.168.250.0/24                                       1     5    NH    Y
    10.200.10.1                           to-ser1
192.168.252.0/24                                       1     5    NH    N
    10.10.0.254                           na
===============================================================================
*A:ALU-12#
```

Table 142: Show Static Route Output Fields

Label	Description
Prefix	The static route destination address and mask
Tag	The 32-bit integer tag added to the static route
Met	The route metric value for the static route
Pref	The route preference value for the static route
Type	BH—the static route is a blackhole route, where the next hop for this type of route is black-hole
	ID—the static route is an indirect route, where the next hop for this type of route is the non-directly connected next hop
	NH—the route is a static route with a directly connected next hop
Act	N—the static route is inactive; for example, the static route is disabled or the next hop IP interface is down
	Y—the static route is active
Next Hop	The next hop for the static route destination
No. of Routes:	The number of routes displayed in the list

tunnel-table

Syntax **tunnel-table** [*ip-address*[*/mask*]] [**protocol** *protocol* | **sdp** *sdp-id*]
tunnel-table [**summary**]

Context show>router

Description This command displays tunnel table information.

Note that **auto-bind** GRE tunnels are not displayed in the **show** command output. GRE tunnels are not the same as SDP tunnels that use the GRE encapsulation type. When the **auto-bind** command is used when configuring a VPRN service, it means the MP-BGP NH resolution is referring to the 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]* — the IP address and mask for the tunnel table destination

protocol **—** displays LDP protocol information

sdp-id **—** displays information about the specified SDP

summary — displays summary tunnel table information

Output The following output is an example of tunnel table information, and Table 143 describes the fields.

Sample Output

```
*A:ALU-12>config>service# show router 6 tunnel-table
===============================================================================
Tunnel Table (Service: 6)
===============================================================================
Destination        Owner Encap TunnelId  Pref     Nexthop        Metric
-------------------------------------------------------------------------------
10.0.0.1/32        sdp   GRE   10        5        10.0.0.1       0
10.0.0.1/32        sdp   GRE   21        5        10.0.0.1       0
10.0.0.1/32        sdp   GRE   31        5        10.0.0.1       0
10.0.0.1/32        sdp   GRE   41        5        10.0.0.1       0
===============================================================================
*A:ALU-12>
```

Table 143: Show Tunnel Table Output Fields

Label	Description
Destination	The route's destination address and mask
Owner	The tunnel owner
Encap	The tunnel encapsulation type
TunnelID	The tunnel (SDP) identifier
Pref	The route preference for routes learned from the configured peer(s)
Nexthop	The next hop for the route's destination
Metric	The route metric value for the route

Clear Service Commands

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears data for a specific service.

Parameters *service-id* — the ID that uniquely identifies a service

Values 1 to 2147483647 or *service-name* (64 characters maximum)

arp

Syntax **arp**

Context clear>service>id

Description This command clears all ARP entries.

dhcp

Syntax **dhcp**

Context clear>service>id

Description This command enables the context to clear and reset DHCP entities.

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id*:*vc-id* | **interface** {*ip-int-name* | *ip-address*}]

Context clear>service>id>dhcp

Description This command clears DHCP statistics for a specified IP interface.

Parameters *ip-int-name* — the name of the IP interface on which to clear DHCP statistics, up to 32 characters (must start with a letter)

sap-id — the SAP ID on which to clear DHCP statistics. See Table 39 for a full list of SAP IDs.

sdp-id — the spoke SDP ID on which to clear DHCP statistics

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID for which to clear DHCP information

Values 1 to 4294967295

ip-address — the IP address of the IP interface

dhcp6

Syntax **dhcp6**

Context clear>service>id

Description This command enables the context to clear and reset DHCPv6 entities.

statistics

Syntax **statistics** [**interface** {*ip-int-name* | *ipv6-address*}]

Context clear>service>id>dhcp6

Description This command clears DHCPv6 statistics for this IP interface.

Parameters **interface** — clears only DHCPv6 statistics for a specified interface name

ip-int-name — the name of the IPv6 interface

ipv6-address — the address of the IPv6 interface

fdb

Syntax **fdb** {**all** | **mac** *ieee-address* | **sap** *sap-id* | **mesh-sdp** *sdp-id*[*:vc-id*] | **spoke-sdp** *sdp-id:vc-id*}

Context clear>service>id

Description This command clears FDB entries for the service.

Parameters **all** — clears all FDB entries

ieee-address — clears only FDB entries in the FDB table for the specified 48-bit MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

sap-id — clears the physical port identifier portion of the SAP definition for the specified SAP ID. See Table 39 for a full list of SAP IDs.

mesh-sdp — clears only service FDB entries associated with the specified mesh SDP ID. For a mesh SDP, the VC ID is optional.

spoke-sdp — clears only service FDB entries associated with the specified spoke SDP ID. For a spoke SDP, the VC ID must be specified.

sdp-id — the SDP ID for which to clear associated FDB entries

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID for which to clear associated FDB entries

Values 1 to 4294967295

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[*:vc-id*] **ingress-vc-label**

Context clear>service>id

Description This command clears and resets the mesh SDP binding for the service.

Parameters *sdp-id* — the mesh SDP ID to be reset

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

mfib

Syntax **mfib**

Context clear>service>id

Description This command enables the context to clear Multicast Forwarding Information Base (MFIB)-related data.

statistics

Syntax **statistics** {**all** | **ip** | **mac** | **group** *grp-address* }

Context clear>service>id>mfib

Description This command clears MFIB statistics.

Parameters **all —** clears all MFIB statistics for the service ID

ip — clears MFIB statistics for the specified IP address

Values a.b.c.d

mac — clears MFIB statistics for the specified MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

grp-ip-address — clears MFIB statistics for the specified group address

Values a.b.c.d

spoke-sdp

Syntax **spoke-sdp** *sdp-id:vc-id* **ingress-vc-label**

Context clear>service>id

Description This command clears and resets the spoke SDP binding for the service.

Parameters *sdp-id* **—** the spoke SDP ID to be reset

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

ingress-vc-label — the locally assigned ingress service label

statistics

Syntax **statistics**

Context clear>service>statistics

Description This command enables the context to clear statistics data.

cem

Syntax **cem**

Context clear>service>statistics>id

Description This command clears Circuit Emulation (CEM) statistics for the service.

counters

Syntax **counters**

Context clear>service>statistics>id

Description This command clears all traffic queue counters statistics associated with the service.

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[*:vc-id*] {**all** | **counters** | **stp**}

Context clear>service>statistics>id

Description This command clears and resets the mesh SDP binding statistics for the service.

Parameters *sdp-id* — the mesh SDP ID to be reset

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

spoke-sdp

Syntax **spoke-sdp** *sdp-id:vc-id* {**all** | **counters**}

Context clear>service>statistics>id

Description This command clears statistics for the spoke SDP bound to the service.

Parameters *sdp-id* — the spoke SDP ID for which to clear statistics

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID to be reset

Values 1 to 4294967295

all — clears all queue statistics associated with the SDP

counters — clears all queue counters associated with the SDP

sap

Syntax **sap** *sap-id* {**all** | **cem** | **counters**}

Context clear>service>statistics

Description This command clears statistics for the SAP bound to the service.

Parameters *sap-id* — specifies the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

all — clears all queue statistics associated with the SAP

cem — clears all CEM statistics associated with the SAP

counters — clears all queue counters associated with the SAP

sap-aggregation-group

Syntax **sap-aggregation-group** *svcId-id:groupName* {**all** | **counters**}

Context clear>service>statistics

Description This command clears statistics for a specified SAP aggregation group.

Parameters *svcId-id:groupName* — the service ID or group name, up to 96 characters

all — clears all queue statistics associated with the SAP aggregation group

counters — clears all queue counters associated with the SAP aggregation group

sdp

Syntax **sdp** *sdp-id* **keep-alive**

Context clear>service>statistics

Description This command clears keepalive statistics associated with the SDP ID.

Parameters *sdp-id* — the SDP ID for which to clear statistics

Values 1 to 17407

keep-alive — clears the keepalive history associated with this SDP ID

Debug Service Commands

id

Syntax **id** *service-id*
no id *service-id*

Context debug>service

Description This command enables the debugging context for a specific service.

The **no** form of the command disables debugging for the service.

Parameters *service-id* — the ID that uniquely identifies a service

Values 1 to 2147483647 or *service-name*

dhcp

Syntax **dhcp**
no dhcp

Context debug>service>id

Description This command enables the context for DHCP debugging.

The **no** form of the command disables DHCP debugging.

detail-level

Syntax **detail-level** {**low** | **medium** | **high**}
no detail-level

Context debug>service>id>dhcp
debug>service>id>dhcp6

Description This command enables DHCP and DHCPv6 detail level tracing.

The **no** form of the command disables the detail level tracing.

mac

Syntax **mac** *ieee-address*
no mac *ieee-address*

Context debug>service>id>dhcp
debug>service>id>dhcp6

Description This command enables debugging for a specified MAC address.

The **no** form of the command disables debugging for the MAC address.

Parameters *ieee-address* — the MAC address

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

mode

Syntax **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
no mode

Context debug>service>id>dhcp

Description This command enables the DHCP tracing mode.

The **no** form of the command disables the tracing mode.

sap

Syntax **sap** *sap-id*
no sap *sap-id*

Context debug>service>id>dhcp
debug>service>id>dhcp6

Description This command enables debugging for a specific SAP.

The **no** form of the command disables the debugging for the SAP.

Parameters *sap-id* — the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP ID

sdp

Syntax **sdp** *sdp-id:vc-id*
no sdp *sdp-id:vc-id*

Context debug>service>id>dhcp

Description This command enables debugging for a specific SDP.

The **no** form of the command disables the debugging for the SDP.

Parameters *sdp-id* — the mesh SDP or spoke SDP ID; for a mesh SDP, the VC ID is optional

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID

Values 1 to 4294967295

dhcp6

Syntax **dhcp6**
no dhcp6

Context debug>service>id

Description This command enables the context for DHCPv6 debugging.

The **no** form of the command disables DHCPv6 debugging.

mode

Syntax **mode** {**all** | **dropped-only**}
no mode

Context debug>service>id>dhcp6

Description This command enables the DHCPv6 tracing mode.

The **no** form of the command disables the tracing mode.

event-type

Syntax **event-type** {**config-change** | **svc-oper-status-change** | **sap-oper-status-change** | **sdpbind-oper-status-change**}
no event-type

Context debug>service>id

Description This command enables debugging for an event type.

The **no** form of the command disables debugging on the event type.

sap

Syntax **sap** *sap-id*
no sap *sap-id*

Context debug>service>id

Description This command enables debugging for a specific SAP.

The **no** form of the command disables debugging for the SAP.

Parameters *sap-id* — the physical port identifier portion of the SAP definition. See Table 39 for a full list of SAP IDs.

sdp

Syntax **sdp** *sdp-id:vc-id*
no sdp *sdp-id:vc-id*

Context debug>service>id

Description This command enables debugging for a specific SDP.

The **no** form of the command disables the debugging for the SDP.

Parameters *sdp-id* — the mesh SDP or spoke SDP ID

Values 1 to 17407

vc-id — the virtual circuit ID on the SDP ID

Values 1 to 4294967295

IPSec

In This Chapter

This chapter provides information to configure security parameters. Topics in this chapter include:

- IPSec Overview
- Public Key Infrastructure (PKI)
- Best Practices Recommendations
- Configuration Notes
- Configuring IPSec with CLI
- IPSec Command Reference

IPSec Overview

This section contains the following topics:

- IPSec Implementation
- X.509v3 Certificate Overview
- Using Certificates For IPSec Tunnel Authentication
- Trust Anchor Profile
- Certificate Profile
- Certificate Management Protocol Version 2 (CMPv2)
- OCSP
- Applications
- NAT-Traversal for IKEv2 and IPSec
- BFD over IPSec Tunnel
- QoS for IPSec
- Fragmentation and IP MTU
- Support for Private VPRN Service Features
- Routing in Private Services
- IPSec on the 10-port 1GigE/1-port 10GigE X-Adapter Card
- IPSec Sequence Number
- PBR and MFC
- Statistics
- Security Support

IPSec Implementation

This section contains the following topics:

- IPSec Overview
- Hardware Support
- IPSec Encryption Features
- SHA2 Support
- IPSec Security Policy, IKE Policy, and IPSec Transform
- Tunnel Group
- Tunnel Interfaces and SAPs
- IPSec Tunnel Configuration

IPSec Overview

IPSec is a structure of open standards to ensure private, secure communications over Internet Protocol (IP) networks by using cryptographic security services.

For IPSec, the 7705 SAR supports VPRN for the private side of the tunnel and IES for the public side of the tunnel. A public service instance (IES) connects to the public network and a private service instance (VPRN) connects to the private network, which originates the traffic that is to be encrypted (see Figure 107).

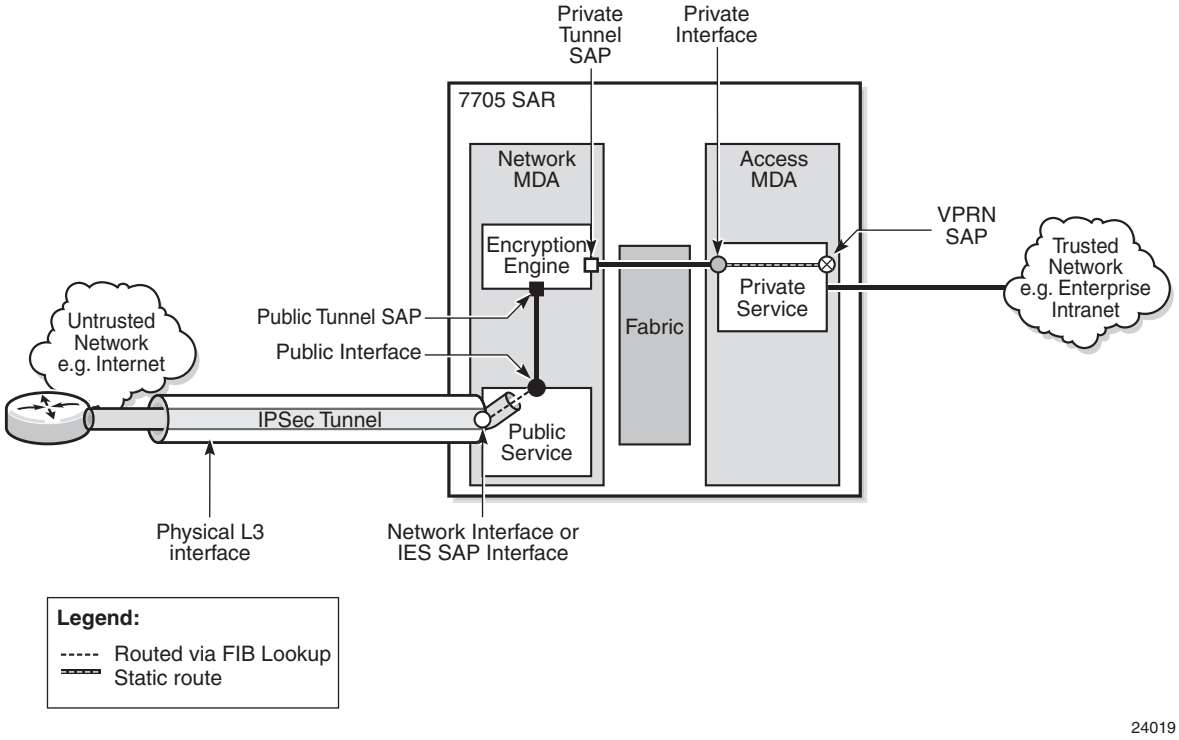
In Figure 107, all ingress customer traffic from the trusted network is aggregated into the private VPRN service, where a VPRN static route directs the traffic into the encryption engine. The encryption engine encrypts the customer traffic using configurable encryption and authentication protocols, and adds the IPSec tunnel outer IP header. The source IP address of the outer IP header is the local security gateway address, and the destination IP address is the peer security gateway address.

The encrypted IPSec packet exits the node via an IES or router interface that is configured on an encryption-capable adapter card; it gets routed to its destination via a standard FIB lookup.

On ingress from an untrusted network, all arriving IPSec traffic is routed to the appropriate security gateway as configured on the VPRN service. The incoming IPSec traffic is processed and checked against the tunnel security configuration.

If the traffic passes all security checks, it is decrypted and the customer traffic is routed through the associated VPRN. Any traffic that does not match the tunnel security configuration is dropped.

Figure 107: IPSec Implementation Architecture



Hardware Support

The 7705 SAR supports IPSec on the following nodes and adapter cards:

- 7705 SAR-8 with 8-port Gigabit Ethernet Adapter card, version 3
- 7705 SAR-18 with either:
  - → 8-port Gigabit Ethernet Adapter card, version 3, or
  - → 10-port 1GigE/1-port 10GigE X-Adapter card, version 2
- 7705 SAR-H
- 7705 SAR-Hc
- 7705 SAR-W
- 7705 SAR-Wx
- 7705 SAR-X

Note: On the 7705 SAR-X, each Ethernet port has its own encryption engine.

IPSec Encryption Features

IPSec provides a variety of encryption features required to establish bidirectional IPSec tunnels, including:

Control plane:

- manual keying
- dynamic keying: IKEv2
- authentication: pre-shared-key (PSK)
- perfect forward secrecy (PFS)
- dead peer detection (DPD)
- NAT-traversal (NAT-T)
- security policy

Data plane:

- ESP (with authentication) tunnel mode
- IPSec transform (NULL cannot be used for authentication and encryption at the same time):
  - → authentication algorithm: NULL/MD5/SHA1/SHA256/SHA384/SHA512
  - → encryption algorithm: NULL/DES/3DES/AES128/AES192/AES256
- IPSec IKE policy (NULL is not supported):
  - → authentication algorithm: MD5/SHA1/SHA256/SHA384/SHA512
  - → encryption algorithm: DES/3DES/AES128/AES192/AES256
- DH-Group: 1/2/5/14/15

Note: The 7705 SAR uses a configured authentication algorithm in an IKE policy for the pseudorandom function (PRF).

SHA2 Support

The 7705 SAR supports RFC 4868. For data origin authentication and integrity verification functions in the IKEv2 and Encapsulating Security Payload (ESP) protocols, the 7705 SAR supports the following HMAC-SHA-256+ algorithms:

- AUTH_HMAC_SHA2_256_128
- AUTH_HMAC_SHA2_384_192
- AUTH_HMAC_SHA2_512_256

For pseudorandom functions (PRF) with IKEv2, the 7705 SAR supports the following HMAC-SHA-256+ algorithms:

- PRF_HMAC_SHA2_256_128
- PRF_HMAC_SHA2_384_192
- PRF_HMAC_SHA2_512_256

IPSec Security Policy, IKE Policy, and IPSec Transform

An IPSec security policy defines the type of traffic allowed to pass in or out of an IPSec tunnel. The policy does this through the configuration of local and remote IP address pairs. The behavior of an IPSec security policy is similar to IP filtering. IPSec security policies are created for a VPRN service context and applied to an IPSec tunnel in that service.

An IKE policy defines how the 7705 SAR encrypts and authenticates an IPSec tunnel that uses that policy. Its configuration includes specifics on Diffie-Hellman key derivation algorithms, encryption and authentication protocols to be used for establishing phase 1 and phase 2 security associations, and so on.

An IPSec transform defines the algorithms used for IPSec SA. The transform configuration dictates the algorithms that customer traffic uses for encryption and authentication.

Tunnel Group

A tunnel group is a collection of IPSec tunnels. The 7705 SAR supports one tunnel group that always uses tunnel ID 1.

An IPSec tunnel belongs to only one tunnel group. The `show>isa>tunnel-group` command allows the operator to view information about the configured tunnel group.

Tunnel Interfaces and SAPs

There are two types of tunnel interfaces and associated SAPs:

- public tunnel interface: configured in the public IES service; outgoing tunnel packets have a source IP address (local gateway address) in this subnet

  → public tunnel SAP: associated with the public tunnel interface
- private tunnel interface: configured in the private VPRN service

  → private tunnel SAP: associated with the private tunnel interface, logically linked to the public tunnel SAP

Public Tunnel SAPs

An IES service (the delivery service) must have at least one IP interface associated with a public tunnel SAP to receive and process the following types of packets associated with IPSec tunnels:

- IPSec ESP (IP protocol 50)
- IKEv2 (UDP)

The public tunnel SAP type has the format `tunnel-`*`tunnel-group-id`*`.public:`*`tag`*, where *`tunnel-group-id`* is always 1. See Configuring IPSec and IPSec Tunnels in Services for a CLI configuration example.

Private Tunnel SAPs

The private (VPRN) service must have an IP interface to an IPSec tunnel in order to forward IP packets into the tunnel, causing them to be encapsulated according to the tunnel configuration, and to receive IP packets from the tunnel after the encapsulation has been removed (and decrypted). That IP interface is associated with a private tunnel SAP.

The private tunnel SAP has the format `tunnel-`*`tunnel-group-id`*`.private:`*`tag`*, where *tunnel-group-id* is always 1. The `tunnel` keyword must be used when creating the private tunnel interface. See Configuring IPSec and IPSec Tunnels in Services for a CLI configuration example.

IP Interface Configuration

The IP MTU of a private tunnel SAP interface can be configured. This sets the maximum payload IP packet size (including IP header) that can be sent into the tunnel and applies to the packet size before the tunnel encapsulation is added. When an IPv4 payload packet that needs to be forwarded to the tunnel is larger than M bytes, one of the following behaviors occurs.

- If the DF bit is clear (not set), the payload packet is fragmented to the MTU size prior to tunnel encapsulation.
- If the DF bit is set, the payload packet is discarded and (if allowed by the ICMP setting of the sending interface) an ICMP type 3/code 4 is returned to the sender (with the MTU of the private tunnel SAP interface in the payload).

IPSec Tunnel Configuration

To bind an IPSec tunnel to a private tunnel SAP, the `ipsec-tunnel` command is configured under the SAP context, where the `ipsec-tunnel` context provides access to the following parameters:

```
security-policy 1
local-gateway-address 80.22.22.2 peer 80.11.11.2 delivery-service 21
dynamic-keying
ike-policy 1
pre-shared-key "alcatel"
transform 1
```

The local gateway address must belong to the same subnet as the delivery-service (IES) public tunnel interface address. The local gateway address and peer gateway address are the source and destination addresses for the outgoing IPSec traffic.

A private tunnel SAP can have only one IPSec tunnel.

X.509v3 Certificate Overview

X.509v3 is an ITU-T standard that consists of a hierarchical system of certificate authorities (CAs) that issue certificates that bind a public key to particular entity's identification. The entity's identification could be a distinguished name or an alternative name, such as a fully qualified domain name (FQDN) or an IP address.

An end entity (EE) is an entity that is not a CA. For example, an end entity can be a web server, a VPN client, or a VPN gateway.

A CA issues a certificate by signing an entity's public key with its own private key. A CA can issue certificates for an end entity as well as for another CA. When a CA certificate is issued for itself (signed by its own private key), this CA is called the root CA. Therefore, an end entity's certificate can be issued by the root CA or by a subordinate CA (that is, issued by another subordinate CA or root CA). When there are multiple CAs involved, this is called a chain of CAs.

In addition to issuing certificates, the public key infrastructure (PKI) also includes a mechanism for revoking certificates due to reasons such as a compromised private key.

A certificate can be used for authentication. Typically, the certificate authentication process functions as follows.

- The system trusts a CA as the trust anchor CA (which typically is a root CA). This means that all certificates issued by a trust anchor CA, or the certificates issued by a subordinate CA that have been issued by the trust anchor CA, are consider trusted.
- A peer that is to be authenticated presents its certificate along with a signature over some shared data between the peer and system, and the certificate is signed using a private key.
- The signature is verified by using the public key in the certificate. In addition, the certificate itself is verified as being issued by the trust anchor CA or a subordinate CA that is part of the chain leading up to the trust anchor CA. The system can also check if the peer's certificate has been revoked. Only when all these verifications succeed does the certificate authentication succeed.

X.509v3 Certificate Support on the 7705 SAR

The 7705 SAR PKI implementation supports the following features:

- certificate enrollment:
  - → locally generated RSA/DSA key
  - → offline enrollment via PKCS#10 (public key cryptography standards)
  - → online enrollment via CMPv2
- support for CA chain
- certificate revocation check:
  - → certificate revocation list (CRL) for both EE and CA certificates
  - → online certificate status protocol (OCSP) for EE certificate only

Local Storage

The 7705 SAR requires the following objects to be stored locally as a file:

- CA certificate
- CRL
- the system's own certificate
- the system's own key

All these objects must be imported with the `admin certificate import` command before they can be used by the 7705 SAR. The import process converts the format of the input file to distinguished encoding rules (DER), encrypts the key file, and saves it in the `cf3:/system-pki` directory.

The imported file can also be exported using a specified format by means of the `admin certificate export` command.

The `admin certificate import` and `admin certificate export` commands support the following formats:

- certificates can be imported and exported using the following formats:
  - → PKCS #12
  - → PKCS #7 (DER and PEM) (privacy enhanced mail)
  - → PEM
  - → DER

  If there are multiple certificates in the file, only the first one is used.

- key pairs can be imported and exported using the following formats:
  - → PKCS #12
  - → PEM
  - → DER
- the CRL can be imported and exported using the following formats:
  - → PKCS #7 (DER and PEM)
  - → PEM
  - → DER
- The PKCS #12 file can be encrypted with a password.

CA Profile

On the 7705 SAR, the CA-related configuration is stored in a CA profile that contains the following configurable items:

- name and description
- CA's certificate — an imported certificate
- CA's CRL— an imported CRL
- revocation check method — specifies the way the CA checks the revocation status of the certificate it issued
- CMPv2 — a CMPv2 server-related configuration
- OCSP— an OCSP responder-related configuration

When a user enables a `ca-profile` (`no shutdown`), the system loads the specified CA certificate and CRL into memory. The following checks are performed:

- for the CA certificate:
  - → all mandatory fields defined in section 4.1 of RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*, exist and conform to the RFC 5280 defined format
  - → the version field value is 0x2
  - → the validity field indicates that the certificate is still in its validity period
  - → the X.509 Basic Constraints extension exists and the CA Boolean value is true
  - → if the Key Usage extension exists, at the least), the **keyCertSign** and **cRLSign** are asserted
- for the CRL:
  - → all mandatory fields defined in section 5.1 of RFC 5280 exist and conform to the RFC 5280 defined format

→ if the version field exists, the value is 0x1

→ the delta CRL indicator does not exist (delta CRL is not supported)

→ the CRL is signed by the configured CA certificate

The CRL is required in order to enable `ca-profile`.

CA Chain Computation

When verifying a certificate with a CA or a chain of CAs, the system must identify the issuer CA of the certificate. The 7705 SAR looks through all configured
CA profiles to find the issuer CA. The following is the method that the system uses to find the issuer CA:

- the issuer CA's certificate subject must match the issuer field of the certificate in question
- if present, the authority key identifier of the certificate in question must match the subject key identifier of the issuer CA's certificate
- if present, the key usage extension of the issuer CA's certificate must permit certificate signing

Certificate Enrollment

The 7705 SAR supports two certificate enrollment methods:

- the offline method using PKCS #10
- the online method using CMPv2

To use the offline method, perform the following steps:

Step 1. Generate a key pair using the **`admin certificate gen-keypair`** command.

For example: **`admin certificate gen-keypair`** `cf3:/segw.key` **`size`** `2048` **`type`** `rsa`

Step 2. Generate a PKCS#10 certificate signing request with the key generated in Step 1 using the `admin certificate gen-local-cert-req` command.

For example: **`admin certificate gen-local-cert-req keypair`** `cf3:/segw.key` **`subject-dn`** `C=US,ST=CA,O=ALU,CN=SeGW` **`domain-name`** `segw-1.alu.com` **`file`** `cf3:/segw.pkcs10`

As well as specifying the subject of the certificate request, you can optionally specify an FQDN and an IP address as SubjectAltName.

Step 3. Import the key file using the **admin certificate import** command.

Example: **admin certificate import type** key **input** cf3:/segw.key **output** segw.key **format** der

Step 4. Because the key is imported, remove the key file generated in Step 1 for security reasons.

Step 5. Send the PKCS #10 file to the CA via an offline method such as email.

The CA signs the request and returns the certificate.

Step 6. Import the returned certificate using the **admin certificate import** command.

Example: **admin certificate import type** cert **input** cf3:/segw.cert **output** segw.cert **format** pem

For the online method using CMPv2-based enrollment, see the Certificate Management Protocol Version 2 (CMPv2).

Certificate Revocation Check

A revocation check is a process that checks whether a certificate has been revoked by the issuer CA.

The 7705 SAR supports two methods for the certificate revocation check:

- CRL
- OCSP

The CRL can be used for both EE and CA certificate checks, while OCSP can only be used for an EE certificate check.

For an IPSec application, users can configure multiple check methods with a priority order for an EE certificate. Using the `status-verify` command in the `ipsec-tunnel` configuration context, users can configure a primary method, a secondary method, and a default result. The primary and secondary methods can be either OCSP or CRL. The default result is either `good` or `revoked`. If the system does not get an answer from the primary method, it falls back to the secondary method. If the secondary method does not return an answer, the system uses the default result.

By default, the system uses the CRL to check the revocation status of a certificate, whether it is an end entity certificate or a CA certificate. This makes the CRL a mandatory configuration in the `ca-profile`.

For details about OCSP, see OCSP.

Certificate, CRL, and Key Cache

Configured certificates, CRLs, and keys are cached in memory before they are used by the system.

Every certificate, CRL, and key has one system-wide cache copy.

For a CA certificate and a CRL, the cache is created when there is a CA profile and when a `no shutdown` is performed and removed.

For an IPSec tunnel using legacy `cert` and `key` configurations, the cache is created only when the first tunnel using the cache is in a `no shutdown` state, and it is cleared when the last tunnel that used it is shut down.

For an IPSec tunnel using a `cert-profile`, the cache is created when the first `cert-profile` using the cache is in a `no shutdown` state, and it is removed when the last `cert-profile` that used it is shut down.

If a certificate or key is configured with both a `cert-profile` and legacy `cert` or `key` command, the cache is created when the first object (an `ipsec-tunnel` or a `cert-profile`) using it is in a `no shutdown` state, and it is removed when the last object using it is shut down.

In order to update a certificate or key without a shutdown `ca-profile` or `ipsec-tunnel`, the CLI command `admin>certificate>reload` manually reloads the certificate and key cache.

Using Certificates For IPSec Tunnel Authentication

The 7705 SAR supports X.509v3 certificate authentication for an IKEv2 tunnel (LAN-to-LAN tunnel and remote-access tunnel). The 7705 SAR also supports asymmetric authentication. This means that the 7705 SAR and the IKEv2 peer can use different methods to authenticate. For example, one side of the tunnel could use a pre-shared key and the other side could use a certificate.

The 7705 SAR supports certificate chain verification. For a static LAN-to-LAN tunnel, the command `trust-anchor-profile` specifies which CAs are expected to be present in the certificate chain before reaching the root CA (self-signed CA) configured in the system.

The key and certificate for the 7705 SAR are also configurable on a per-tunnel basis.

When using certificate authentication, the 7705 SAR uses the subject of the configured certificate as its ID by default.

Trust Anchor Profile

The 7705 SAR supports multiple trust anchors for each IPSec tunnel. A trust anchor profile can be configured with up to eight CAs. The system builds a certificate chain by using the certificate in the first certificate payload in the received IKEv2 message. If any of the configured trust anchor CAs in the trust anchor profile appear in the chain, then authentication is successful; otherwise, authentication fails.



Note: The 7705 SAR only supports processing of up to 16 hashes for the trust anchor list from other products. If the remote end sends more than 16 hashes and a certificate match is in the 17th or later hash, the tunnel remains down due to authentication failure.

Certificate Profile

The 7705 SAR supports sending different certificates and chains according to the received IKEv2 certificate-request payload. This is done by configuring a `cert-profile` that allows up to eight entries. Each entry includes a certificate and a key and, optionally, a chain of CA certificates.

The system loads the certificate and/or key in the `cert-profile` into memory and builds a compare-chain for the certificate configured in each entry of the `cert-profile` upon a `no shutdown` of the `cert-profile`. These chains are used for IKEv2 certificate authentication. If a chain computation cannot be completed for a configured certificate, the corresponding compare-chain will be empty or only partially computed.

Because there can be multiple entries configured in the `cert-profile`, the system must pick the certificate and key in the entry that the other side expects to receive. This is done by looking up the CAs within the received certificate request payload in the compare-chain and picking the first entry that has a certificate request CA appearing in its chain. If there is no such `cert`, the system picks the first entry in the `cert-profile`. The first entry is the first configured entry in the `cert-profile`. The *`entry-id`* of the first entry does not have to be “1”.

For example, assume there are three CAs listed in the certificate-request payload: CA-1, CA-2 and CA-3, and there are two entries configured in the `cert-profile`, as shown in the following configuration:

```
cert-profile "cert-profile-1"
  entry 1
    cert "cert-1"
    key "key-1"
  entry 2
    cert "cert-2"
    key "key-2"
    send-chain
      ca-profile "CA-1"
      ca-profile "CA-2"
```

The system builds two compare-chains: chain-1 for cert-1 and chain-2 for cert-2. Assume CA-2 appears in chain-2, but CA-1 and CA-3 do not appear in either chain-1 or chain-2. In that case, the system will pick entry 2.

After a certificate profile entry is selected, the system generates the AUTH payload by using the configured key in the selected entry. The system also sends the certificate in the selected entry as "certificate" payload to the peer.

If a chain is configured in the selected entry, then one certificate payload is needed for each certificate in the configured chain. The first certificate payload in the IKEv2 message will be the signing certificate, which is configured by the `cert` command in the chosen `cert-profile` entry. In the preceding example, the system will send three certificate payloads: cert-2, CA-1, and CA-2.

The following CA chain-related enhancements are supported.

- The `no shutdown` of a `ca-profile` triggers a recomputation of the compute-chain in related certificate profiles. The system also generates a new log-1 to indicate that a new compute-chain has been generated; the log includes the CA profile names on the new chain. Another log, log-2, is generated if the `send-chain` in a `cert-profile` entry is not in a compute-chain due to this CA profile change. Another log is generated if the hash calculation for a certificate under a ca-profile has changed.
- When performing a `no shutdown` command on a `cert-profile`, the system allows the CAs in the send-chain, not in the compute-chain. The system also generates log-2, as above.
- The system allows changes to the configuration of the send-chain without shutting down `cert-profile`.

Certificate Management Protocol Version 2 (CMPv2)

CMPv2 is a protocol between a Certificate Authority (CA) and an end entity (EE) based on RFC 4210, *Internet X.509 Public Key Infrastructure Certificate Management Protocol* (*CMP*). It provides multiple certificate management functions, such as certificate enrollment, certificate update, and so on.

The 7705 SAR supports the following CMPv2 operations:

- initial registration — the process by which the 7705 SAR enrolls a certificate with a certain CA for the first time, where:
  - → a public/private key pair must be pre-provisioned before enrollment by means of local generation or other methods
  - → optionally, users can include a certificate or certificate chain in the extraCerts field of the initial registration request
- key pair update — a process by which the 7705 SAR updates an existing certificate for any reason (for example, a refresh of a key or certificate before it expires)
- certificate update — a process by which an initialized 7705 SAR obtains additional certificates
- polling — in some cases, the CA may not immediately return the certificate for reasons such as "request processing needs manual intervention". In such cases, the 7705 SAR supports polling requests and responds as described in Section 5.3.22, Polling Request and Response, in RFC 4210.

The following list provides implementation details.

- HTTP is the only supported transport protocol for CMPv2. HTTP 1.1 and 1.0 are supported and configurable.
- All CMPv2 messages sent by the 7705 SAR consist of only one PKI message. In all cases, the size of the sequence for PKI messages is 1.
- Both password-based MAC and public key-based signature CMPv2 message protection are supported.
- The 7705 SAR only allows one outstanding `ir/cr/kur` request for each CMPv2 server. The means that no new requests are allowed if a pending request is present.

OCSP

The Online Certificate Status Protocol (OCSP) is used by 7705 SAR applications to determine the revocation state of an identified certificate, based on RFC 2560, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*. Unlike the CRL, which relies on checking against an offline file, OCSP provides timely, online information regarding the revocation status of a certificate.

IPSec is the only supported application that uses OCSP. The 7705 SAR supports both the CRL and OCSP as the certificate revocation status checking method. For a given IPSec tunnel, the user can configure a primary method, a secondary method, and a default result to achieve a hierarchical fallback mechanism. If the primary method fails to return a result, the system falls back to the secondary method. If the secondary method fails, the system uses the default result.

The following list provides implementation details.

- Only an OCSP client function is supported.
- HTTP is the only supported transport protocol.
- OCSP server access via a management routing instance is not supported.
- The 7705 SAR does not sign an OCSP request.
- The OCSP response must be signed. The system will verify the response by using the signer's certificate included in the response. If there is no such certificate, the CA certificate in the `ca-profile` will be used.
- If a nextUpdate exists in the OCSP response, the 7705 SAR checks the current time to determine if it is earlier than the nextUpdate. If yes, then the response is valid; otherwise, the response is considered unreliable and the 7705 SAR moves to the next revocation checking method.
- The revocation status result from a valid OCSP response is cached in the system.
- OCSP can only be used to verify the revocation status of the EE certificate. The CRL is still needed to verify the status of a CA certificate.

Applications

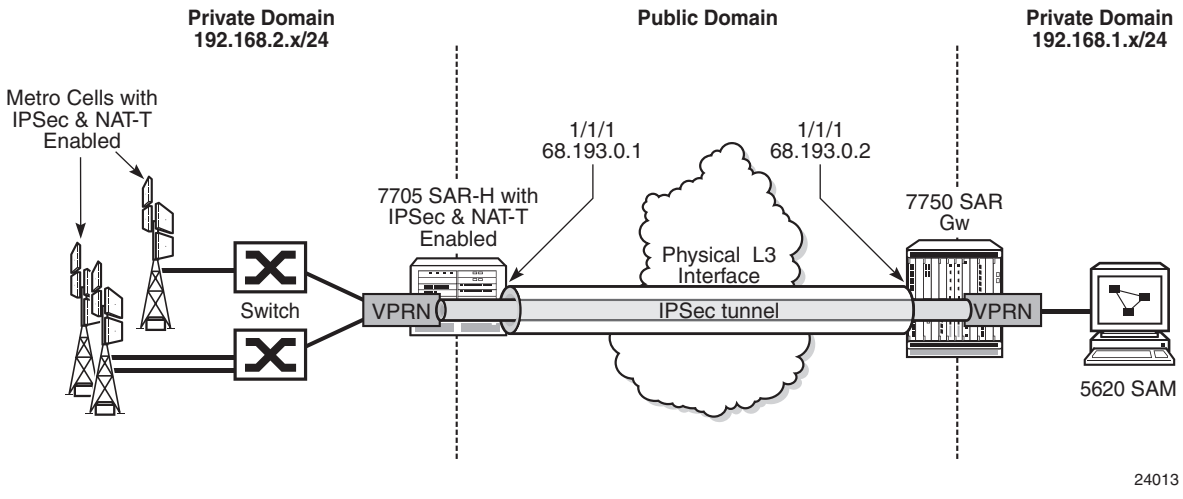
Two mobile backhaul applications are described in this section:

- Metrocell Deployment: a solution for providers who are looking for security in the transmission medium to manage remote private networks in a metrocell deployment. IPSec is used as an encrypted uplink for OAM and mobile traffic to connect the remote network to the MTSO. The 7705 SAR-initiated IPSec tunnels can provide a secure means for managing the 7705 SAR and any private network behind the 7705 SAR, while NAT can provide scalability of IPSec tunnels over a single public IP address.
- Small Business Deployment: the use of LTE and IP NodeBs, as an alternative to PWs, to provide a better match for an operator's choice of transport network (that is, IPSec over public network compared to MPLS/PWs over a private network)

Metrocell Deployment

As shown in Figure 108, in a typical metrocell deployment, the cell site network is divided into two separate segments: the private domain and the public domain. An IPSec tunnel generated from the 7705 SAR-H is used to backhaul the management and OAM traffic of the private network, including the management traffic of the switches and the 7705 SAR-H itself. All OAM traffic is aggregated within a VPRN service and uses the IPSec tunnel as the uplink tunnel to the 7750 SR gateway.

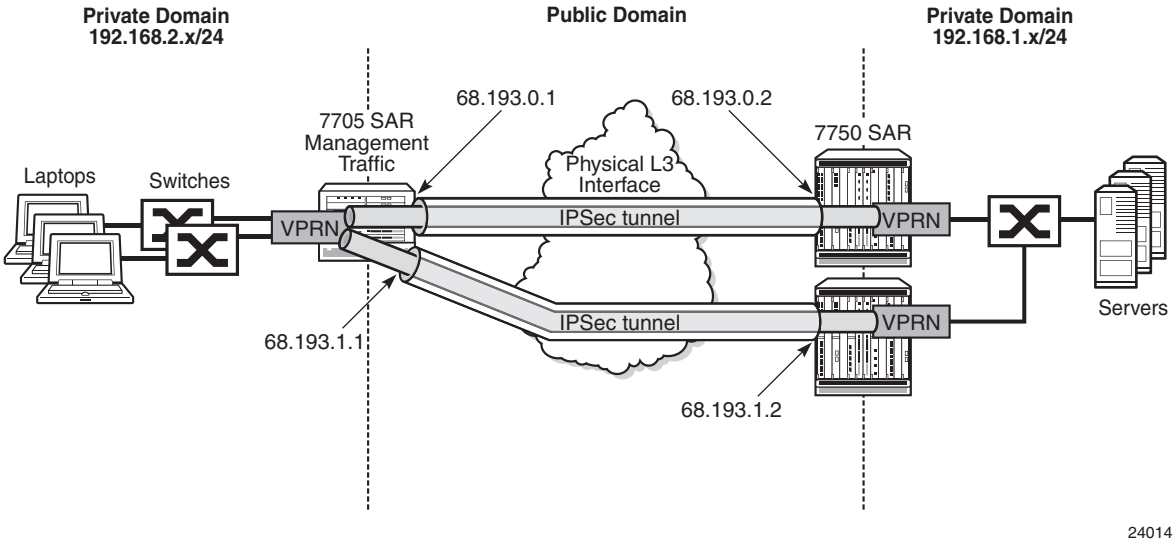
Figure 108: Typical Metrocell Deployment



Small Business Deployment

In a small business deployment, the network is usually designed with a hub and spoke topology. The spoke sites connect to the hub through a leased line or a public non-secure domain. IPSec provides the security and encryption needed to connect the spoke sites to the centralized office (hub). The hub and spoke topology in a small business deployment is favorable because of the security that the hub side can provide to the entire network. IPS/IDS and anti-virus appliances can be deployed to the hub site, which examines arriving traffic from the spoke sites. SPAM and viruses can be filtered out on the hub site by these appliances. If additional spoke-to-spoke connectivity is required, then additional IPSec tunnels can be established. See Figure 109.

Figure 109: Typical Small Business Deployment



NAT-Traversal for IKEv2 and IPSec

The 7705 SAR supports NAT-T (Network Address Translation-Traversal) for IKEv2, as described in section 2.23 of RFC 5996, *Internet Key Exchange Protocol Version 2 (IKEv2)*. NAT-T is functionality belonging to IPSec and IKEv2. It is not functionality belonging to the NAT device.

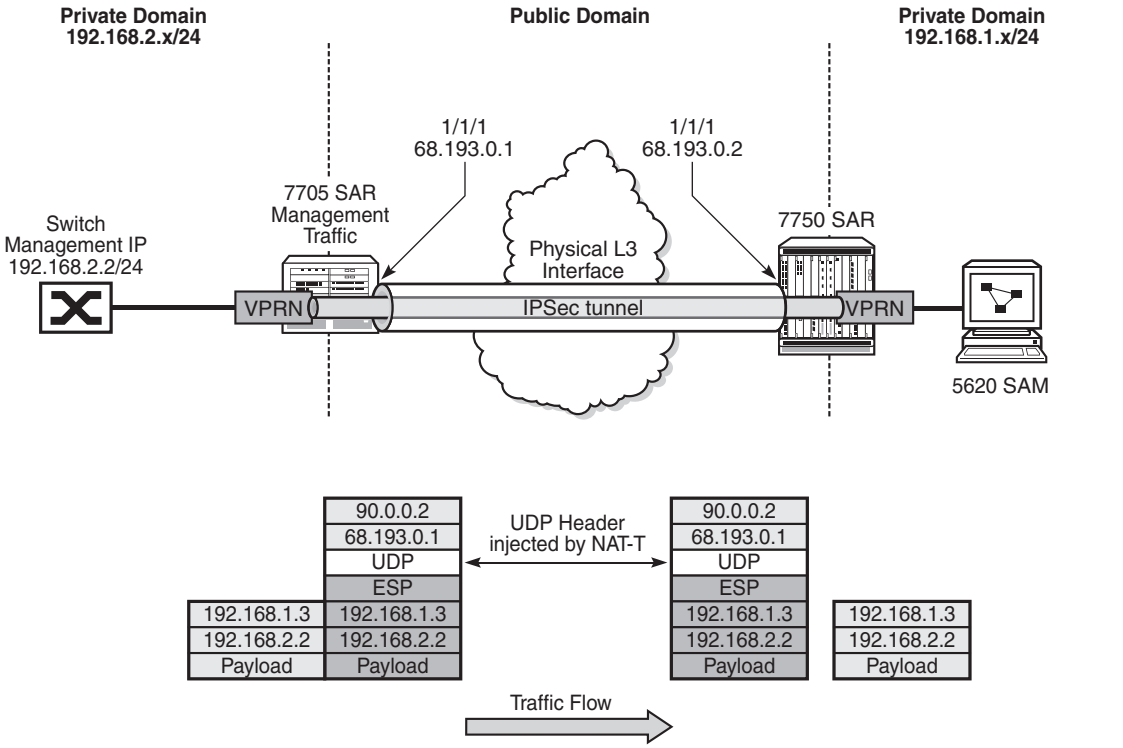
In a private network where the entire network is hidden behind a single public IP address, NAT-T for IPSec is used to support the fan-out of multiple IPSec tunnels in the private network.

IPSec is an IP protocol and as such does not use ports. Figure 110 illustrates how the UDP header is injected into the packet as well as the many-to-one to one-to-many mappings. NAT relies on port mapping, so in order to allow traversal of a NAT device, NAT-T adds a UDP header with port 4500 to the IPSec traffic when the NAT device is detected. The UDP header is added to the IPSec packet above the ESP header and IKEv2 already uses UDP port 500. This UDP header can be used by the NAT device to uniquely map each IPSec tunnel and assign a different source port to each individual tunnel. That is, many IP addresses using UDP 4500 lead to a NAT mapping where a single public IP address uses many UDP ports.

In Figure 110, the 7750 SR performs the following functions:

- tracks the different metrocell IKEv2 port-to-session mappings
- tracks the different metrocell IPSec port-to-tunnel mappings
- transmits traffic to each metrocell on the appropriate UDP port

Figure 110: UDP Header Injected by a NAT-T-enabled IPSec Tunnel



500/udp - Internet Key Exchange (IKE)
4500/udp NAT Traversal

IKE Establishment

srcIP 192.168.2.2, UDPsrcPort 500 ⇒ srcIP 68.193.0.1, UDPsrcPort 5000

dstIP 192.168.2.2, UDPdstPort 500 ⇐ dstIP 68.193.0.1, UDPdstPort 5000

IPSec with NAT-T Establishment

srcIP 192.168.2.2, UDPsrcPort 4500 ⇒ srcIP 68.193.0.1, UDPsrcPort 6000

dstIP 192.168.2.2, UDPdstPort 4500 ⇐ dstIP 68.193.0.1, UDPdstPort 6000

24016

BFD over IPSec Tunnel

To configure BFD for an IPSec tunnel, do the following:

- configure BFD on a loopback interface in the private VPRN
- configure at least two IPSec tunnels:
  - → one tunnel is a BFD-designate tunnel over which BFD packets are exchanged; this BFD-designate tunnel does not go down when BFD goes down
  - → the other tunnels are tunnels that use the BFD-designate tunnel's BFD session; these tunnels go down when BFD goes down
- configure a static route in the private VPRN, where the static route points to the destination node's private-side loopback interface, using the BFD-designate tunnel as the next hop
- configure BFD under the BFD-designate tunnel using the loopback interface and point to the far-end loopback address
- configure BFD under the protected tunnels also using the loopback interface (same configuration as under the BFD-designate tunnel)

QoS for IPSec

This section contains information on the following topics:

- Network and Access Ingress QoS (Decryption QoS)
- Network and Access Egress QoS (Encryption QoS)

Network and Access Ingress QoS (Decryption QoS)

IPSec traffic arriving on network ingress is classified based on network policy and network queue policy when the uplink interface is a network interface (refer to the 7705 SAR OS Quality of Service Guide). This classification is done based on the DSCP marking of the IPSec outer IP header.

The IPSec (encrypted) traffic destined for the secure gateway (SeGW) of the 7705 SAR is mapped to two queues (expedited and best effort) of the decryption engine on the ingress adapter card. This means that encrypted traffic is mapped to the decryption queue.

The encrypted traffic is mapped to one of these two queues based on the queue-type of its mapped network ingress queue, as determined by the result of the network ingress classification. For example, at network ingress, the QoS network policy determines a forwarding class (FC) for the packet. Then, the network-ingress queue policy maps the FC to a queue. The configured queue-type of network-ingress queue (`expedited`, `best-effort`, or `auto-expedite`) is used to choose the queue for the decryption engine on the ingress adapter card (where the uplink interface resides).

The uplink interface for the SeGW can be configured as a network interface or as an access IES interface. For an access IES interface, the decryption-queue mapping is based on the queue-type of the access ingress queue of the IES interface SAP. For example, at IES ingress, the SAP ingress policy determines a forwarding class (FC) for the packet. The FC is mapped to a SAP ingress queue. The queue-type of this SAP ingress queue (`expedited`, `best-effort`, or `auto-expedite`) is used to choose the queue for the decryption engine on the ingress adapter card where the IES interface SAP resides.

The decrypted customer traffic with removed IPSec tunnel header is queued on network and access ingress queues (the uplink interface can be a network interface or an IES interface) based on the network and access ingress policy, and the DSCP bits of the IPSec outer IP header are used for the classification.

Network Ingress QoS Tunnel Override

When a network QoS policy is enabled on an ingress network interface, IPSec packets arriving at that interface are assigned to encryption queues based on the IPSec outer IP header. However, if the QoS network policy of the arriving network interface is configured with `ler-use-dscp`, then after decryption, all the datapath or firewall queuing is based on the DSCP marking of the IPSec inner IP header rather than on the DSCP marking of the IPSec outer IP header. For more information about the `ler-use-dscp` command, refer to the 7705 SAR OS Quality of Service Guide, “QoS Policy Network Commands”.

Network and Access Egress QoS (Encryption QoS)

Customer packets arriving on access ingress in the VPRN are classified based on the SAP ingress policy (refer to the 7705 SAR OS Quality of Service Guide).

Customer packets arriving in the VPRN that are destined for the IPSec tunnel are enqueued before the encryption engine on the egress adapter card. There are three queues servicing the encryption engine on the egress adapter card (expedited, best effort, and CTL).

All CSM traffic over IPSec (BFD, ping, and so on) is queued in the CTL queue, while data (customer) traffic is mapped to the expedited or best-effort queue.

The customer traffic to the two data queues is mapped based on the queue-type of the ingress SAP queue. For example, at access VPRN SAP ingress, the ingress SAP policy determines a forwarding class (FC) for the packet. The FC is mapped to a SAP ingress queue. The queue-type of the SAP ingress queue (`expedited`, `best-effort`, or `auto-expedite`) is used to choose the queue for the decryption engine on the egress adapter card (where the uplink interface resides).

Note: If DSCP egress re-marking is configured on the network interface or access interface (uplink interface), DSCP bits are re-marked on the IPSec outer IP header.

Fragmentation and IP MTU

On the 7705 SAR, unencrypted IP packets arriving on a VPRN access interface and destined for an IPSec uplink will be fragmented if the incoming packet is larger than:

- the VPRN private interface MTU
- the IPSec tunnel MTU
- the difference between the uplink MTU and the IPSec overhead (uplink interface MTU minus IPSec overhead), where the IPSec overhead values are calculated as follows:

  → IPSec overhead if NAT-T is enabled

  IPSec overhead = outer IPSec (20) + UDP (8) + ESP (24) + trailer (17) + ICV (32)
  = 101 bytes

  → IPSec overhead if NAT-T is disabled (`no nat-t`)

  IPSec overhead = outer IP (20) + ESP (24) + trailer (17) + ICV (32)
  = 93 bytes

For example, if a private tunnel interface has its IP MTU set to 1000 bytes, then a packet larger than 1000 bytes will be fragmented. As another example, if an uplink interface has its IP MTU set to 1000 bytes, then a packet that is larger than 1000 – IPSec overhead will be fragmented. Both these examples assume that the DF bit is not set or the `clear-df-bit` command is enabled.

Fragmentation Configuration

By default, the 7705 SAR sets the DF bit on the IPSec tunnel IP header.

There are some configurations where the customer IP header DF bit needs to be copied into the IPSec tunnel IP header. The `copy-df-bit` command under the `config>service>vprn>if>sap>ipsec-tunnel` context enables copying the customer clear text IP header DF bit into IPSec tunnel IP header.

The `clear-df-bit` command , also under the `ipsec-tunnel` context, clears the customer clear text IP header DF bit (if it is set). This allows the customer packet to be fragmented into the IPSec tunnel even if the customer packet has the DF bit set. However, the fragmented customer clear text packet is not be reassembled at the far end of IPSec tunnel.

Reassembly

The 7705 SAR does not support reassembly of fragmented IPSec packets.

Support for Private VPRN Service Features

Private VPRN access features are only supported on non-IPSec interfaces. That is, they are only supported for Layer 3 interfaces that are not configured with a private IPSec tunnel.

Some of these features include r-VPLS, VRRP, ECMP, and LAG. See VPRN Services for information on these features.

Routing in Private Services

For static LAN-to-LAN tunnels, the static route with the IPSec tunnel as the next-hop could be exported to a routing protocol by a route policy. The protocol type remains static.

IPSec on the 10-port 1GigE/1-port 10GigE X-Adapter Card

The 10-port 1GigE/1-port 10GigE X-Adapter card has two encryption engines that share the encryption/decryption load. Therefore, the 10-port 1GigE/1-port 10GigE X-Adapter card has the potential for double the encryption/decryption throughput when compared with other adapter cards and nodes with a single encryption engine (the 8-port Gigabit Ethernet Adapter card, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-W, and 7705 SAR-Wx).

To utilize the potential of the 10-port 1GigE/1-port 10GigE X-Adapter card, the IPSec security associations (SAs) are load-balanced between the two engines based on a round-robin algorithm. When there is an SA download to the 10-port 1GigE/1-port 10GigE X-Adapter card, the upper-layer software load-balances the SA on the two engines.

IPSec Sequence Number

The IPSec sequence number is used to prevent replay attacks. A replay attack is a network attack in which valid data transmission is repeated or delayed for fraudulent purposes. The 7705 SAR supports a 32-bit sequence number, where the transmission of each packet increments the sequence number. If there is a sequence number rollover, the 7705 SAR performs the rollover by resignaling the phase-2 IKE negotiation.

PBR and MFC

Both PBR (policy-based routing) and MFC (multi-field classification) are part of the ingress ACL (access control list) configuration on the 7705 SAR. Hence, both PBR and MFC are supported by IPSec on the 7705 SAR, as discussed below:

- PBR
- MFC

PBR

PBR configuration can be applied in two places for an IPSec service.

The first place is for VPRN and applies to all incoming access traffic into a private VPRN. In this case, PBR can be used to direct the customer traffic into uplink IPSec tunnels by means of ACL matching criteria. The filtering action of forwarding to an indirect next hop can be used to direct customer traffic into the appropriate IPSec tunnel. Note that the security policy works only on the original (customer packet) IP header; that is, the PBR next hop is not used in making the security policy decision.

The second place is for IPSec traffic entering the 7705 SAR from the public domain. A PBR filter can be placed on the network interface or the IES interface to direct the IPSec packet based on the matching/forwarding criteria. In this case, IPSec packets are processed by the PBR filter in the same way as any other IP packet.

For more information on PBR, refer to the "Policy-Based Routing" section in the 7705 SAR OS Router Configuration Guide.

Note:

- All routing decisions are made based on the PBR configuration; therefore, it is possible that even if the packet is destined for the local node security gateway (SeGW), the PBR filter might redirect the packet to another interface.
- Alternatively, for IPSec packets that are not destined for the local node SeGW, PBR can force the packets into the local node SeGW. In this case, the encapsulating security payload (ESP) index of the IPSec packet will not match the SeGW ESP configuration and the packet will be dropped. Thus, it is the responsibility of the network administrator to ensure that the PBR configuration is correct and meets the network needs.

MFC

Multi-field classification (MFC) is supported on the private IPSec service (VPRN). MFC functions in the same manner as the VPRN configuration of traditional services.

For more information on MFC, refer to the "Multi-field Classification" section in the 7705 SAR OS Router Configuration Guide.

Statistics

All statistics for security association and tunnel statistics on the 7705 SAR are retrieved from the datapath on demand. When the user requests the statistics for a tunnel, the statistics are retrieved directly from the datapath; the retrieved statistics are not cached on the 7705 SAR. This means that on redundant platforms (that is, on the 7705 SAR-8 or 7705 SAR-18), statistics will not synchronize over to the inactive CSM, and at the time of a CSM switchover, all statistics will be lost. Also, in the case of statistics rollover in the datapath, the newly retrieved statistics will start from 0 (zero) again.

Security Support

IPSec on the 7705 SAR requires a public service (IES or VPRN) and a private service (VPLS). All IPSec traffic on the public service is encrypted. By the time the traffic is routed to the private service, it has been decrypted and can be examined by the 7705 SAR firewall engine.

The public interface of IPSec service is a secure gateway (SeGW). Because packets arriving on and leaving this interface are encrypted with an IPSec header, they are not firewalled. However, NAT can be applied to traffic traversing the public interface of IPSec security zones.

After being decrypted, the customer traffic may traverse a second security zone configured within the VPRN, to sanitize any of these packets according to the firewall rules. This security zone can be extended to have NAT performed on the customer clear text packets.

For more information about configuring security parameters, refer to the 7705 SAR OS Router Configuration Guide, “Configuring Security Parameters”.

Public Key Infrastructure (PKI)

Public Key Infrastructure (PKI) is a cryptographic technique that enables users to securely communicate on an insecure public network, and reliably verify and authenticate the identity of a user through the use of digital signatures.

PKI is a system for creating, storing, and distributing digital certificates, which are used to verify that a particular public key belongs to a particular entity. PKI creates digital certificates that map public keys to entities and securely stores these certificates in a central repository, revoking them as needed.

PKI includes the following components:

- X.509v3 identity certificates
- a certificate authority (CA), which has the following properties:
  - → has a secure server
  - → can sign and publish X.509v3 certificates
  - → is trusted by all users of the system
- public/private key pairs
- the ability to be deployed as flat architecture or hierarchical architecture (chained certificates)

CA Role in PKI

The role of a CA in PKI includes the following items.

- The CA is a trusted third-party organization or company that issues digital certificates.
- The CA may or may not be a third party from the end entity's (EE's) point of view.
- The CA often belongs to the same organization as the EEs it supports.
- The CA can be a root CA or a subordinate CA:
  - → root CA – a CA that is directly trusted by an end entity.
  - → subordinate CA – CAs that are not root CAs. The first subordinate CA in a hierarchy obtains its CA certificate from the root CA. This first subordinate CA can, in turn, use this key to issue certificates that verify the integrity of another subordinate CA.
- The CA verifies digital certificates using a chain of trust, the root CA being the trust anchor for the digital certificate.

- The root CA issues a root certificate, which is the top-most certificate of the certificate tree. The root certificate's private key is used to sign other certificates. All certificates immediately below the root certificate inherit the trustworthiness of the root certificate.

Digital Signature and Certificates

Digital signatures and digital certificates are not the same objects.

A digital signature is an electronic signature that can be used to demonstrate the authenticity of a message. Digital signatures use hashing and asymmetric encryption. There are two aspects to a digital signature:

- signature construction
- signature verification

A digital certificate is an electronic document that uses a digital signature to bind a public key with an identity that has been digitally signed by a CA.

Certificates

PKI uses two types of certificates:

- root certificate
  - → created by a well-known root CA (see CA Role in PKI)
  - → used to validate the authenticity of provider certificates
  - → must be installed on the 7705 SAR manually, where the installation method can be FTP, SFTP, or SC, into the cf3: directory of the 7705 SAR
- provider certificate
  - → created by each vendor for its own authentication needs
  - → contains the public RSA or DSA keys that are used in PKI for the encryption of the phase 1 IKE message
  - → must be signed by the root CA to prove the authenticity of the vendor certificate to the far-end node

Vendor Certificate Signature by the Root CA

For each vendor, the 7705 SAR must have a vendor certificate signed by the CA and stored internally.

The sequence of events is as follows:

1. The 7705 SAR creates an X.509v3 certificate and the key pair (public/private key).
2. The 7705 SAR sends the certificate to the CA to be signed by the CA private key.
   i. The CA runs the hash over the X.509v3 certificate.
   ii. The result of the hash is encrypted via the CA public key (digital signature).
   iii. The digital signature is appended to the certificate and, consequently, the signed vendor certificate.

Vendor Certificate Authentication by a Peer

While the IKE is being initiated, phase 1 of the IKE can be authenticated via PKI. As such, the certificate from Vendor Certificate Signature by the Root CA is sent to the peer as part of the IKE authorization stage. The peer must ensure that this certificate is generated by the correct 7705 SAR and not an intermediary node.

Upon receiving the certificate, the peer does the following.

1. The peer runs a hash over the X.509v3 part of the certificate.
2. The peer decrypts the digital signature of the certificate.
3. If the hash calculated in step 1 and the hash decrypted in step 2 match, then the certificate is authenticated.
4. The peer can use the public key in the certificate to encrypt the IKE phase 2 channel.

Example of PKI Operation

In the following scenario, which provides an example of PKI operation, the 7705 SAR and 7750 SR trust the same CAs and they have already obtained the CA's certificate (which includes the CA's public key) using an out-band method. The CA issues a certificate for the 7705 SAR, as follows.

- The CA calculates a hash of the contents of the 7705 SAR certificate, which includes the 7705 SAR public key, name, validity date, allowed uses for the certificate, and so on.
- The CA encrypts this hash using its private key, and attaches the resulting signature to the 7705 SAR certificate.

The 7705 SAR can now present this signed certificate to the 7750 SR during the IKE phase.

The 7750 SR verifies the received certificate as follows.

- The 7750 SR calculates its own hash of the 7705 SAR certificate contents.
- The 7750 SR decrypts the signed hash in the certificate using the CA's public key, and if the hashes match the 7750 SR knows that this certificate was signed by the CA.
- Optionally, the 7750 SR consults the CA's certificate revocation list (CRL) and checks that the certificate from the 7705 SAR has not been revoked.

Certificate Chain

A peer's certificate may be issued by an intermediate CA. In this case, a user must have and trust all the intermediate CA certificates up to and including the root CA installed in the peer for authentication of an end entity.

The 7705 SAR supports the following implementation.

- When receiving a certificate, all subordinate CAs must be installed locally (that is, only an EE certificate can be received from the peer and processed by the 7705 SAR). Even if a certificate chain is received from the peer, the 7705 SAR processes the EE certificate only.
- The `send-chain` command under `cert-profile` is for third-party peers that support receiving a chain certificate. In this case, the 7705 SAR can send a chain certificate to be used in the entire chain if the third-party peer supports receiving a chained certificate. The 7705 SAR does not support receiving chained certificates.

Certificate Storage

The 7705 SAR IPSec configuration expects the keys and certificates to be stored in a particular directory on the 7705 SAR compact flash. This directory is called cf3:\system-pki and is created automatically when the first file is imported into this folder.

The following files can be imported and exported to and from the cf3:\system-pki directory. A sample directory is shown after the list.

- key pair – this file is encrypted during the import process
- certificates
- certificate revocation list (CRL)

```
A:ALU-A>file cf3:\system-pki\ # dir
Directory of cf3:\system-pki\
09/09/2015  09:17a     <DIR>        ./
09/09/2015  09:17a     <DIR>        ../
09/22/2015  11:38a                906 CMS1-ROOTCA-CERT
09/22/2015  11:41a                458 CMS1-ROOTCA-CRL
09/24/2015  08:18a                864 cert-1
09/25/2015  08:18a               1192 SAR-key-1
09/25/2015  09:32a                905 cal_cert_CMS1-SUBCA
09/25/2015  09:32a                457 cal_crl_CMS1-SUBCA
          6 Files(s)              4732 bytes.
          2 Dir(s)                65605632 bytes free.
```

Note: Always use the `import` and `export` commands to move files in and out of this directory. Do not copy any files directly to or from the system PKI directory.

CMPv2 Certificate Management

The Certificate Management Protocol (CMP) is an Internet protocol used for obtaining X.509v3 digital certificates in a PKI. It is described in RFC 4210. CMP messages are encoded in ASN.1 using the DER method and are usually transported over HTTP.

A CA issues the certificates and acts as the secure server in PKI using CMP. One of the clients obtains its digital certificates by means of this protocol and is called the end entity (EE).

The 7705 SAR supports the following CMPv2 operations:

- initial registration (`ir`)

  This is the process that the 7705 SAR uses to enrol a certificate with a certain CA for the first time.

  A public/private key pair must be pre-provisioned before enrolment by means of local generation or another method.
- certificate update (`cr`)

  This is the process whereby an initialized 7705 SAR obtains additional certificates.
- key pair update (`kur`)

  This is the process where the 7705 SAR updates an existing certificate due to any reason, such as a key or certificate refresh before the key or certificate expires.
- polling

  In some cases, the CA may not return the certificate immediately, for reasons such as “request processing needs manual intervention”. In such cases, the 7705 SAR supports polling requests and responses, as described in Section 5.3.22, *Polling Request and Response*, in RFC 4210.

CMPv2 Initial Registration

Initial registration is a process that the end entity uses to enroll a certificate with a certain CA for the first time. The result of this process is that a CA issues a certificate for an end entity's public key, returning that certificate to the end entity or posting that certificate in a public repository (or both).

The 7705 SAR must be pre-provisioned with the operator CA certificate.

The 7705 SAR public/private key pair is always pre-provisioned before enrollment by means of local generation or another method.

The 7705 SAR uses the CMPv2 initial registration process to enroll its pre-provisioned key with an operator's CA. The result of this process is a certificate issued by the operator's CA.

There are two authentication methods (PKI message protection) in this process, which are chosen using the CLI:

- MSG_MAC_ALG: uses a pre-shared key and a reference number that is pre-issued by the CA
- MSG_SIG_ALG: uses a CLI-provided protection key to sign the message; if a protection key is not provided, the key to be certified is used

Key Update

When a key pair is due to expire, the relevant end entity (EE) may request a key update. That is, the EE may request that the CA issue a new certificate for a new key pair or, under certain circumstances, a new certificate for the same key pair. The request is made using a key update request (`kur`) message, also known as a certificate update operation.

This command requests a new certificate from the CA in order to update an existing certificate due to reasons such as the need to refresh a key or to replace a compromised key.

If the EE already has a signing key pair with a corresponding verification certificate, then communication between the EE and the CA will be protected by the EE's digital signature.

If the request is successful, the CA returns the new certificate in a key update response (`kup`) message, which is syntactically identical to a CertRepMessage.

CRL

In the operation of some cryptosystems, such as PKIs, a certificate revocation list (CRL) is used. A CRL is list of certificates (or more specifically, a list of serial numbers for certificates) that have been revoked. Entities presenting those revoked certificates should not be trusted.

The CRL must be obtained and imported via CMPv2.

OCSP

The online certificate status protocol (OCSP) enables applications to determine the revocation status of an X.509v3 digital certificate. OCSP was created as an alternative to using CRLs and can be used to obtain additional status information. OCSP is described in RFC 6960.

Messages communicated via OCSP are encoded in ASN.1 and are usually communicated over HTTP. OCSP consists of a request message and a response message.

An OCSP responder may return a signed response indicating that the certificate specified in the request is good, revoked, or unknown. The Enterprise Java Bean Certificate Authority (EJBCA) server contains, by default, an internal OCSP responder and, therefore, can be used in conjunction with the 7705 SAR.

Ensure that both the 7705 SAR and the OCSP server are running NTP so that both devices are synchronized with respect to their timing.

Best Practices Recommendations

This section provides best practices recommendations.

IPSec Best Practices

To avoid high CPU loads and some complex cases, the following are suggestions to configure the IKEv2 lifetime.

- Both the IKE_SA and CHILD_SA lifetime on one side should be approximately 2 or 3 times larger than the other side.
- With the previous rule, the lifetime of the side with the smaller lifetime should NOT be too small:
  - → IKE_SA: greater than or equal to 86 400 s
  - → CHILD_SA: greater than or equal to 3600 s
- With the first rule, on the side with the smaller lifetime, the IKE_SA lifetime should be at least 3 times larger than the CHILD_SA lifetime.

The IKE protocol is the control plane of IPSec; thus, the IKE packet should be treated as high QoS priority in the end-to-end path of public service.

- On a public interface, a SAP-ingress QoS policy should be configured to ensure that the IKE packet gets high QoS priority.

Configuration Notes

This section describes operational conditions and IPSec configuration caveats.

- A tunnel group that is in use cannot be deleted. Changes are allowed only when the tunnel group is in a shutdown state.
- A change to the IPSec transform policy is allowed at any time. The change will not impact tunnels that have been established until they are renegotiated. If the change is required immediately, the tunnel must be cleared (reset) for force renegotiation.
- A change to the IKE policy is allowed at any time. The change will not impact tunnels that have been established until they are renegotiated. If the change is required immediately, the tunnel must be cleared (reset) for force renegotiation.
- An IPSec tunnel must be shut down before the transform policy can be modified.
- The public interface address can be changed at any time (current behavior). If changed, tunnels that were configured to use it will require a configuration change. If the subnet has been changed, the tunnels will be in an operationally down state until their configuration is corrected. The public service cannot be deleted while tunnels are configured to use it. A public service is the IES service that holds an interface with a public tunnel SAP that connects the node to the public network. A private service connects to the private protected service.
- The 7705 SAR supports only one tunnel group (tunnel-group 1).
- A change to the security policy is not allowed while a tunnel is active and using the policy.
- The tunnel local gateway address, peer address, or delivery router parameters cannot be changed while the tunnel is operationally up (`shutdown` makes the tunnel both administratively down and operationally down).
- A tunnel security policy cannot be changed while the tunnel is operationally up. An IPSec transform policy or IKE policy assignment to a tunnel requires the tunnel to be shut down.

Reference Sources

For information on supported IEEE standards, IETF drafts and standards as well as standard and proprietary MIBs, refer to Standards and Protocol Support.

Configuring IPSec with CLI

This section provides information to configure IPSec using the command line interface.

Topics in this section include:

- Basic Configuration Overview
- Common Configuration Tasks
- Service Management Tasks

Basic Configuration Overview

The following list provides a high-level outline for setting up IPSec on the 7705 SAR.

- Create an IPSec tunnel group.
- Configure an IPSec IKE policy.
- Configure an IPSec transform policy.
- Create a private side tunnel interface on the VPRN service.
  - → Create an interface using the tunnel keyword and private tunnel SAP.
  - → Create the IPSec tunnel and configure its parameters, which include local and peer gateway IP addresses, IP MTU, keying (manual or dynamic), and so on.
- Create a public side tunnel interface on the IES.
  - → Create an interface and public tunnel SAP.
- Configure a VPRN static route for the IPSec tunnel.

Common Configuration Tasks

This section provides a brief overview of the following common configuration tasks that must be performed to configure IPSec:

- Configuring an IPSec Tunnel Group
- Configuring Router Interfaces for IPSec
- Configuring IPSec Parameters
- Configuring IPSec and IPSec Tunnels in Services
- Configuring IPSec IPv6 Parameters for a VPRN Service
- Configuring X.509v3 Certificate Parameters
- Configuring CMPv2
- Configuring OCSP

Configuring an IPSec Tunnel Group

The following output displays an IPSec group configuration in the ISA context. The 7705 SAR supports only one tunnel-group. The *tunnel-group-id* is always 1.

```
*A:7705custDoc:Sar18>config>isa# info detail
----------------------------------------------
        tunnel-group 1 create
            shutdown
            no description
        exit
----------------------------------------------
*A:7705custDoc:Sar18>config>isa#
```

Configuring Router Interfaces for IPSec

An IPSec tunnel requires the following three interfaces:

- public tunnel interface (under IES)
- private tunnel interface (under VPRN)
- physical untrusted network/Internet-facing interface: either IES or router

The physical interface is the one that must reside on an encryption-capable adapter card.

The following example displays an interface (“internet”) configured using a network port (1/1/1) and an IES interface (“public”) configuration using SAP 1/1/8.

```
*A:ALU-49>config# info
----------------------------------------------
...
    router
        interface "internet"
            address 10.10.7.118/24
            port 1/1/1
        exit
        interface "system"
            address 10.20.1.118/32
        exit
         autonomous-system 123
    exit
...
----------------------------------------------
*A:ALU-49>config#


*A:7705:Dut-A>config>service>ies# info
----------------------------------------------
            description "ies interface towards internet"
            interface "public" create
                address 200.1.1.1/24
                sap 1/1/8 create
                    description "sap-100-200.1.1.1"
                exit
            exit
            no shutdown
----------------------------------------------
```

Configuring IPSec Parameters

Under the IPSec context, configure the IKE policy and IPSec transform parameters.

The following example displays the IPSec parameter configuration output.

```
*A:7705custDoc:Sar18>config>ipsec# info
#------------------------------------------------
    ipsec
        ike-policy 2 create
            own-auth-method psk
            dh-group 14
            ipsec-lifetime 48000
            isakmp-lifetime 60000
            pfs dh-group 5
            auth-algorithm sha384
            encryption-algorithm aes192
            nat-traversal keep-alive-interval 240
            dpd interval 25
        exit
        ipsec-transform 2 create
            esp-auth-algorithm md5
            esp-encryption-algorithm 3des
        exit
    exit
#------------------------------------------------
```

Configuring IPSec and IPSec Tunnels in Services

IPSec is configured under IES and VPRN services.

Under VPRN service, configure IPSec security policies, and create tunnel interfaces, private tunnel SAPs, and IPSec tunnels along with setting the IPSec tunnel parameters. Use the `tunnel` keyword when creating an interface for a private tunnel SAP.

Under IES service, create an interface and public tunnel SAP. The `tunnel` keyword is not used when creating an interface for a public tunnel SAP.

Private and public tunnels function in pairs, where a pair is defined by the *`tag`* in the *`sap-id`*; for example, SAP `tunnel-1.private:22` and SAP `tunnel-1.public:22` are a pair.

The local gateway address and delivery service configured under the VPRN `ipsec-tunnel>local-gateway-address` command correspond to the IES interface address and service ID where the public tunnel interface is defined. In the example below, the `local-gateway-address` is 10.10.10.11 and the `delivery-service` is 10.

The following example displays the configuration output when configuring IPSec for VPRN services and for IES.

```
*A:7705custDoc:Sar18>config>service>vprn# info detail
----------------------------------------------
...
            ipsec
                security-policy 1 create
                    entry 1 create
                        local-ip any
                        remote-ip any
                    exit
                    entry 2 create
                        local-ip 192.100.100.10/32
                        remote-ip 192.100.200.10/32
                    exit
                exit
                security-policy 15 create
                    entry 15 create
                        no local-ip
                        no remote-ip
                    exit
                exit
            exit
...
            interface "vprn_tunnel" tunnel create
                no ip-mtu
                sap tunnel-1.private:22 create
                    no description
                    ingress
                        qos 1
                    exit
                    egress
```

```
                        qos 1
                        no filter
                        no agg-rate-limit
                    exit
                    ipsec-tunnel "ipsec_tunnel_tag1" create
                        shutdown
                        no description
                        security-policy 1 2
                        local-gateway address 10.10.10.11 peer 11.11.11.11
                           delivery-service 10
                        no bfd-designate
                        no clear-df-bit
                        no ip-mtu
                    exit
                    no shutdown
                exit
                no shutdown
            exit
            no service-name
            static-route 192.100.200.10/32 ipsec-tunnel "ipsec_tunnel_tag1"
-----------------------------------------------
*A:7705custDoc:Sar18>config>service>vprn#


*A:7705custDoc:Sar18>config>service>ies# info detail
-----------------------------------------------
  ...
        ies 10 customer 1 create
            interface "ies_tunnelPublicSide_1" create
                address 10.10.10.1/24
                sap tunnel-1.public:22 create
                    no description
                    ingress
                        qos 1
                    exit
                    egress
                        qos 1
                        no filter
                        no agg-rate-limit
                    exit
                    no collect-stats
                    no accounting-policy
                    no shutdown
                exit
            exit
            no service-name
-----------------------------------------------
*A:7705custDoc:Sar18>config>service>ies#
```

Configuring IPSec IPv6 Parameters for a VPRN Service

Use the following CLI syntax to configure IPSec IPv6 parameters for a VPRN service:

CLI Syntax:

```
config>service# vprn service-id [customer customer-id]
  [create]
      ipsec
           security-policy security-policy-id [create]
                entry entry-id [create]
                     local-v6-ip {ipv6-prefix/prefix-
                          length | any}
                     remote-v6-ip {ipv6-prefix/prefix-
                          length | any}
```

Example:

```
A:ALU-41>config>service# vprn 1011
A:ALU-41>config>service>vprn$ ipsec
A:ALU-41>config>service>vprn>ipsec>security-policy$ 1
 create
A:ALU-41>config>service>vprn>ipsec>sec-plcy>entry$ 1
 create
A:ALU-41>config>service>vprn>ipsec>sec-
 plcy>entry>local-v6-ip$ fd00:1:1:1::/64
A:ALU-41>config>service>vprn>ipsec>sec-
 plcy>entry>local-v6-ip$ exit
A:ALU-41>config>service>vprn>ipsec>sec-
 plcy>entry>remote-v6-ip$ fd00:1:129:1::/64
A:ALU-41>config>service>vprn>ipsec>sec-
 plcy>entry>remote-v6-ip$ exit
A:ALU-41>config>service>vprn>ipsec>sec-plcy>entry$ exit
A:ALU-41>config>service>vprn>ipsec>security-policy$
 exit
A:ALU-41>config>service>vprn>ipsec$ exit
```

The following example displays IPSec IPv6 parameters configuration output.

```
*A:7705:Dut-A>config>service>vprn# info
----------------------------------------------
            ipsec
                security-policy 1 create
                    entry 1 create
                        local-v6-ip fd00:1:1:1::/64
                        remote-v6-ip fd00:1:129:1::/64
                    exit
                exit
            exit
```

Configuring X.509v3 Certificate Parameters

Perform the following steps to configure certificate enrollment.

Step 1. Generate a key:

```
admin certificate gen-keypair cf3:/key_plain_rsa2048
size 2048 type rsa
```

Step 2. Generate a certificate request:

```
admin certificate gen-local-cert-req keypair cf3:/
key_plain_rsa2048 subject-dn "C=US,ST=CA,CN=7705" file
7705_req.csr
```

Step 3. Send the certificate request to CA-1 to sign and get the signed certificate.

Step 4. Import the key:

```
admin certificate import type key input cf3:/
key_plain_rsa2048 output key1_rsa2048 format der
```

Step 5. Import the signed certificate:

```
admin certificate import type cert input cf3:/
7705_cert.pem output 7705cert format pem
```

Perform the following steps to import the CA certificate and CRL.

Step 1. Import the CA certificate:

```
admin certificate import type cert input cf3:/
CA_1_cert.pem output ca_cert format pem
```

Step 2. Import the CA's CRL:

```
admin certificate import type crl input cf3:/
CA_1_crl.pem output ca_crl format pem
```

The following example displays a certificate authentication for IKEv2 static LAN-to-LAN tunnel configuration.

```
config>system>security>pki# info
----------------------------------------------
                ca-profile "alu-root" create
                    cert-file "alu_root.cert"
                    crl-file "alu_root.crl"
                    no shutdown
                exit
----------------------------------------------
config>ipsec# info
----------------------------------------------
        ike-policy 1 create
            auth-method cert-auth
        exit
        ipsec-transform 1 create
        exit
```

```
        cert-profile "segw" create
             entry 1 create
                 cert segw.cert
                 key segw.key
             exit
             no shutdown
         exit
         trust-anchor-profile "alu" create
             trust-anchor "alu-root"
         exit

config>service>vprn>if>sap
-----------------------------------------------
                ipsec-tunnel "t50" create
                    security-policy 1
                    local-gateway-address 192.168.55.30 peer 192.168.33.100 delivery-
                                    service 300
                    dynamic-keying
                        ike-policy 1
                        transform 1
                        cert
                            trust-anchor-profile "alu"
                            cert-profile "segw"
                        exit
                    exit
                    no shutdown
                exit
```

The following example displays the syntax to import a certificate from the PEM format.

```
*A:ALU-A# admin certificate import type cert input cf3:/pre-import/R1-
0cert.pem output R1-0cert.der format pem
```

The following example displays the syntax to export a certificate to the PEM format.

```
*A:ALU-A#  admin certificate export type cert input R1-0cert.der output cf3:/
R1-0cert.pem format pem
```

Configuring CMPv2

CMPv2 server information is configured under a corresponding `ca-profile` by using the following CLI commands:

CLI Syntax:
```
config>system>security>pki>ca-profile
     cmpv2
          url url-string [service-id service-id]
          response-signing-cert filename
          key-list
               key password [hash | hash2] reference
                  reference-number
```

The `url` command specifies the HTTP URL of the CMPv2 server and the `service-id` specifies the routing instance that the system used to access the CMPv2 server (if the service ID is omitted, the system uses the base routing instance).

The service ID is only needed for in-band connections to the server via VPRN services. IES services are not referenced by the service ID, since an IES service routing instance is considered to be a base routing instance.

The `response-signing-cert` command specifies an imported certificate that is used to verify CMP response messages if they are protected by a signature. If this command is not configured, the CA's certificate is used.

The `key-list` command specifies a list of pre-shared-keys used for CMPv2 initial registration message protection.

Example:
```
config>system>security>pki>ca-profile>
     cmpv2
          url "http://cmp.example.com/request" service-
            id 100
          key-list
               key passwordToBeUsed [hash | hash2]
                  reference "1"
```

All CMPv2 operations are invoked by using the `admin certificate cmpv2` command.

If there is no `key-list` defined under the `cmpv2` configuration, the system defaults to the `cmpv2` transaction that was input for the command line related to authenticating a message without a sender ID. If there is no sender ID in the response message and there is a `key-list` defined, the system chooses the lexicographical first entry only, and if that fails, there is a fail result for the transaction.

The system supports optional commands (such as `always-set-sender-ir`) to support inter-operation with CMPv2 servers. Refer to the IPSec Command Reference for details about syntax and usage.

Configuring OCSP

OCSP server information is configured under the corresponding `ca-profile`:

CLI Syntax:
```
config>system>security>pki>ca-profile>
     ocsp
          responder-url url-string
          service service-id
```

The `responder-url` command specifies the HTTP URL of the OCSP responder. The `service` command specifies the routing instance that the system used to access the OCSP responder.

Example:
```
config>system>security>pki>ca-profile>
     ocsp
          responder-url "http://ocsp.example.com/
            request"
          service 100
```

For a given IPSec tunnel, the user can configure a primary method, a secondary method, and a default result.

```
config>service>vprn>if>sap>ipsec-tun>
     cert
          status-verify
          primary {ocsp | crl}
          secondary {ocsp | crl}
          default-result {revoked | good}
```

Example:
```
config>service>vprn>if>sap>ipsec-tun>
     cert
          status-verify
          primary ocsp
          secondary crl
```

Service Management Tasks

This section provides a brief overview of the following service management tasks:

- Deleting an IPSec IKE Policy or an IPSec Transform
- Deleting a Public IPSec Tunnel SAP and Interface
- Deleting a Private IPSec Tunnel SAP and Interface
- Deleting an IPSec Security Policy
- Deleting an IPSec Tunnel

Deleting an IPSec IKE Policy or an IPSec Transform

An IPSec IKE policy or transform cannot be deleted if it is being used by an IPSec tunnel. To delete an IKE policy or IPSec transform:

CLI Syntax:

```
config>service>vprn>if>sap>ipsec-tunnel# dynamic-keying
config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying#
 no ike-policy
config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying#
 no transform
config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying#
 exit all
config>ipsec# no ike-policy ike-policy-id
config>ipsec# no ipsec-transform transform-id
```

Example:

```
config>service>vprn>if>sap>ipsec-tunnel# dynamic-keying
config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying#
 no ike-policy
config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying#
 no transform
config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying#
 exit all
config>ipsec# no ike-policy 2
config>ipsec# no ipsec-transform 2
```

Deleting a Public IPSec Tunnel SAP and Interface

A public IPSec tunnel interface and SAP are created under IES. To delete a public IPSec tunnel interface and SAP:

CLI Syntax:
```
config>service>ies>interface# no sap tunnel-id.public:tag
config>service>ies# no interface ip-int-name
```

Example:
```
config>service>ies>interface# no sap tunnel-1.public:22
config>service>ies# no interface ies_tunnelPublicSide_1
```

Deleting a Private IPSec Tunnel SAP and Interface

A private IPSec tunnel interface and SAP are created under a VPRN service. To delete a private IPSec tunnel interface and SAP:

CLI Syntax:
```
config>service>vprn>interface# no sap tunnel-id.private:tag
config>service>vprn# no interface ip-int-name
```

Example:
```
config>service>vprn>interface# no sap tunnel-1.private:22
config>service>vprn# no interface vprn-tunnel
```

Deleting an IPSec Security Policy

Security policies are created under the VPRN service. To delete an IPSec security policy:

CLI Syntax:
```
config>service>vprn>ipsec# no security-policy security-policy-id
```

Example:
```
config>service>vprn# no security-policy 35
```

Deleting an IPSec Tunnel

IPSec tunnels are created under the VPRN service. To delete an IPSec tunnel:

CLI Syntax: `config>service>vprn>if>sap# no ipsec-tunnel` *`ipsec-tunnel-name`*

Example:
```
config>service>vprn>if>sap# no ipsec-tunnel
 ies_tunnelPublicSide_1
```

IPSec Command Reference

Command Hierarchies

- IPSec Configuration Commands
  - → ISA Tunnel Commands
  - → IPSec Commands
  - → Service Configuration Commands
  - → Service Interface Tunnel Commands
  - → Service Static Route Commands
- PKI Configuration Commands
  - → X.509 and Certificate Commands
  - → PKI Infrastructure Commands
  - → IPSec PKI Commands
  - → IKE PKI Commands
- Show Commands
- Clear Commands
- Debug Commands

IPSec Configuration Commands

ISA Tunnel Commands

```
config
    — [no] isa
        — tunnel-group tunnel-group-id [create]
        — no tunnel-group tunnel-group-id
            — description description-string
            — no description
            — [no] shutdown
```

IPSec Commands

```
— config
    — ipsec
        — ike-policy ike-policy-id [create]
        — no ike-policy ike-policy-id
            — auth-algorithm {md5 | sha1 | sha256 | sha384 | sha512}
            — no auth-algorithm
            — auth-method psk
            — no auth-method
            — description description-string
            — no description
            — dh-group {1 | 2 | 5 | 14 | 15}
            — no dh-group
            — dpd [interval interval] [max-retries max-retries] [reply-only]
            — no dpd
            — encryption-algorithm {des | 3des | aes128 | aes192 | aes256}
            — no encryption-algorithm
            — ipsec-lifetime ipsec-lifetime
            — no ipsec-lifetime
            — isakmp-lifetime isakmp-lifetime
            — no isakmp-lifetime
            — nat-traversal [force] [keep-alive-interval keep-alive-interval]
                [force-keep-alive]
            — no nat-traversal
            — own-auth-method psk
            — no own-auth-method
            — pfs [dh-group {1 | 2 | 5}]
            — no pfs
        — ipsec-transform transform-id [create]
        — no ipsec-transform transform-id
            — esp-auth-algorithm {null | md5 | sha1| sha256 | sha384 | sha512}
            — no esp-auth-algorithm
            — esp-encryption-algorithm {null | des | 3des | aes128 | aes192 | aes256}
            — no esp-encryption-algorithm
```

Service Configuration Commands

```
config
    — service
        — vprn service-id
            — ipsec
                — security-policy security-policy-id [create]
                — no security-policy security-policy-id
                    — entry entry-id [create]
                    — no entry entry-id
                        — local-ip {ip-prefix/prefix-length | ip-prefix netmask | any}
                        — no local-ip
                        — local-v6-ip {ipv6-prefix/prefix-length | any}
                        — no local-v6-ip
                        — remote-ip {ip-prefix/prefix-length | ip-prefix netmask | any}
                        — no remote-ip
                        — remote-v6-ip {ipv6-prefix/prefix-length | any}
                        — no remote-v6-ip
```

Service Interface Tunnel Commands

```
config
    — service
        — ies
            — interface ip-int-name [tunnel] [create]
            — no interface ip-int-name
                — sap sap-id [create]
                — no sap sap-id
```

```
config
    — service
        — vprn
            — interface ip-int-name [tunnel] [create]
            — no interface ip-int-name
                — sap sap-id [create]
                — no sap sap-id
                    — ipsec-tunnel ipsec-tunnel-name [create]
                    — no ipsec-tunnel ipsec-tunnel-name
                        — [no] bfd-designate
                        — bfd-enable service service-id interface interface-name dst-ip
                          bip-address
                        — no bfd-enable
                        — [no] clear-df-bit
                        — [no] copy-df-bit
                        — description description-string
                        — no description
                        — [no] dynamic-keying
                            — [no] auto-establish
                            — ike-policy ike-policy-id
                            — no ike-policy
                            — local-id type {ipv4 | fqdn} value value
```

```
            — no local-id
            — pre-shared-key key
            — no pre-shared-key
            — transform transform-id [transform-id...(up to 4 max) ]
            — no transform
      — ip-mtu octets
      — no ip-mtu
      — local-gateway-address ip-address peer ip-address delivery-
          service service-id
      — no local-gateway-address
      — [no] manual-keying
            — security-association security-entry-id
                  authentication-key authentication-key
                  encryption-key encryption-key spi spi transform
                  transform-id direction {inbound | outbound}
            — no security-association security-entry-id direction
                  {inbound | outbound}
      — security-policy security-policy-id
      — no security-policy
```

Service Static Route Commands

```
config
    — service
        — vprn service-id
            — [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
                [preference preference] [metric metric] [tag tag] [enable | disable]
                {next-hop ip-address | ipsec-tunnel ipsec-tunnel-name}
                [bfd-enable | {cpe-check cpe-ip-address [interval seconds] [drop-count count]
                [log]}
```

See VPRN Services Command Reference for the command descriptions.

PKI Configuration Commands

X.509 and Certificate Commands

```
admin
    — certificate
        — clear-ocsp-cache [entry-id]
        — cmpv2
            — cert-request ca ca-profile-name current-key key-filename current-cert cert-filename [hash-alg hash-algorithm] newkey key-filename subject-dn subject-dn save-as save-path-of-result-cert
            — clear-request ca ca-profile-name
            — initial-registration ca ca-profile-name key-to-certify key-filename protection-alg {password password reference ref-number | signature [cert cert-file-name [send-chain [with-ca ca-profile-name]]] [protection-key key-filename] [hash-alg {md5 | sha1 | sha224 | sha256 | sha384 | sha512}]} subject-dn dn save-as save-path-of-result-cert
            — key-update ca ca-profile-name newkey key-filename oldkey key-filename oldcert cert-filename [hash-alg hash-algorithm] save-as save-path-of-result-cert
            — poll ca ca-profile-name
            — show-request [ca ca-profile-name]
        — display type {cert | key | crl | cert-request} url-string format {pkcs10 | pkcs12 | pkcs7-der | pkcs7-pem | pem | der} [password password]
        — export type {cert | key |crl} input input-filename output url-string format output-format [password password] [pkey pkey-filename]
        — gen-keypair url-string [size {512 | 1024 | 2048}] [type {rsa | dsa}]
        — gen-local-cert-req keypair url-string subject-dn subject-dn [domain-name domain-name] [ip-addr {ip-address | ipv6-address}] file url-string [hash-alg hash-algorithm] [use-printable]
        — import type {cert | key | crl} input url-string output filename format input-format [password password]
        — reload type {cert | key} filename [key-file filename]
```

PKI Infrastructure Commands

```
configure
    — system
        — security
            — pki
                — ca-profile name [create]
                — no ca-profile name
                    — cert-file filename
                    — no cert-file
                    — cmpv2
                        — [no] accept-unprotected-errormsg
                        — [no] accept-unprotected-pkiconf
                        — [no] always-set-sender-for-ir
                        — http-response-timeout timeout
```

```
                        — no http-response-timeout
                        — http-version {1.0 | 1.1}
                        — key-list
                            — key password [hash | hash2] reference reference-
                                  number
                            — no key reference reference-number
                        — response-signing-cert filename
                        — no response-signing-cert
                        — [no] same-recipnonce-for-pollreq
                        — url url-string [service-id service-id]
                        — no url
                    — crl-file filename
                    — no crl-file
                    — description description-string
                    — no description
                    — ocsp
                        — responder-url url-string
                        — no responder-url
                        — service service-id
                        — no service
                    — [no] shutdown
                — certificate-display-format {ascii | utf8}
                — maximum-cert-chain-depth level
                — no maximum-cert-chain-depth
```

IPSec PKI Commands

```
config
    — ipsec
        — cert-profile profile-name [create]
        — no cert-profile profile-name
            — entry entry-id [create]
            — no entry entry-id
                — cert cert-filename
                — no cert
                — key key-filename
                — no key
                — [no] send-chain
                    — [no] ca-profile name
        — ike-policy ike-policy-id [create]
        — no ike-policy ike-policy-id
            — auth-method {psk | cert-auth}
            — no auth-method
            — own-auth-method {psk |cert}
            — no auth-method
        — [no] shutdown
        — trust-anchor-profile name [create]
        — no trust-anchor-profile name
```

IKE PKI Commands

```
config
    — service
        — vprn
            — interface ip-int-name [tunnel] [create]
            — no interface ip-int-name
                — sap sap-id [create]
                — no sap sap-id
                    — ipsec-tunnel ipsec-tunnel-name [create]
                    — no ipsec-tunnel ipsec-tunnel-name
                        — cert
                            — cert-profile profile
                            — no cert-profile
                            — status-verify
                                — default-result {revoked | good}
                                — no default-result
                                — primary {crl | ocsp}
                                — no primary
                                — secondary {crl | ocsp}
                                — no secondary
                            — trust-anchor-profile profile-name
                            — no trust-anchor-profile
```

Show Commands

```
show
    — certificate
        — ca-profile name [association]
        — ocsp-cache entry-id
        — statistics
    — ipsec
        — cert-profile name association
        — cert-profile [name]
        — cert-profile name entry [1..8]
        — ike-policy ike-policy-id
        — ike-policy
        — security-policy service service-id [security-policy-id security-policy-id]
        — security-policy
        — transform [transform-id]
        — trust-anchor-profile trust-anchor-profile association
        — trust-anchor-profile [trust-anchor-profile]
        — tunnel
        — tunnel ipsec-tunnel-name
        — tunnel count
```

```
show
    — mda slot/mda
        — statistics {source-mda | dest-mda | security [encryption]}   (for 7705 SAR-18 and
            7705 SAR-8 (with CSMv2) )
    — mda aggregate-statistics   (for 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-W, and 7705 SAR-Wx)
```

Refer to the section "Card, Adapter Card, and Port Command Reference" in the 7705 SAR OS Interface Configuration Guide for information on the **show>mda** commands.

```
show
    — router
        — interface ip-int-name statistics
```

Refer to the section "IP Router Command Reference" in the 7705 SAR OS Router Configuration Guide for information on the **show>router >interface statistics** command.

Clear Commands

```
clear
    — mda {slot/mda | all}
    — mda all statistics
    — mda slot/mda statistics security [encryption]
```

Debug Commands

debug
- — [**no**] **cmpv2**
  - — [**no**] **ca-profile** *profile-name*
- — **ipsec**
  - — [**no**] **certificate** *filename*
  - — **tunnel** [*ipsec-tunnel-name*] [**detail**]
  - — **no tunnel** [*ipsec-tunnel-name*]

Command Descriptions

- IPSec Configuration Commands
- PKI Configuration Commands
- Show Commands
- Clear Commands
- Debug Commands

IPSec Configuration Commands

- Generic Commands
- ISA Tunnel Commands
- Internet Key Exchange (IKE) and Transform Commands
- Service Configuration Commands
- Service Interface Tunnel Commands

Generic Commands

description

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

Context config>ipsec>ike-policy
config>isa>tunnel-group
config>service>ies>interface
config>service>ies>if>sap
config>service>vprn>interface
config>service>vprn>if>sap
config>service>vprn>if>sap>ipsec-tunnel

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

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

Default No description is associated with the configuration context.

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

shutdown

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

Context config>isa>tunnel-group
config>service>ies>interface
config>service>ies>if>sap
config>service>vprn>interface
config>service>vprn>if>sap

Description The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many objects must be shut down before they may be deleted. Many entities must be explicitly enabled using the **no shutdown** command.

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

Services are created in the administratively down state (**shutdown**). When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state.

ISA Tunnel Commands

isa

Syntax **[no] isa**

Context config

Description This command creates an ISA tunnel configuration context.

The **no** form of this command removes the context.

Default n/a

tunnel-group

Syntax **tunnel-group** *tunnel-group-id* [**create**]
no tunnel-group *tunnel-group-id*

Context config>isa

Description This command enables a tunnel group to be created or edited. The 7705 SAR can have only one tunnel group (**tunnel-group** 1).

The **no** form of the command deletes the specified tunnel group from the configuration.

Default n/a

Parameters *tunnel-group-id* — specifies an integer value that uniquely identifies the tunnel group

Values 1 to 16 (1 is the only valid value)

create — mandatory keyword required when creating a tunnel group. The **create** keyword requirement can be enabled/disabled in the **environment**>**create** context.

Internet Key Exchange (IKE) and Transform Commands

ipsec

Syntax **ipsec**

Context config

Description This command enables the context to configure Internet Protocol security (IPSec) parameters. IPSec is a structure of open standards to ensure private, secure communications over Internet Protocol (IP) networks by using cryptographic security services.

ike-policy

Syntax **ike-policy** *ike-policy-id* [**create**]
no ike-policy *ike-policy-id*

Context config>ipsec

Description This command enables provisioning of IKE policy parameters.

The **no** form of the command removes the IKE policy.

Parameters *ike-policy-id* — specifies a policy ID value to identify the IKE policy

Values 1 to 2048

create — mandatory keyword required when creating an IKE policy. The **create** keyword requirement can be enabled/disabled in the **environment**>**create** context.

auth-algorithm

Syntax **auth-algorithm** {**md5** | **sha1** | **sha256** | **sha384** | **sha512**}
no auth-algorithm

Context config>ipsec>ike-policy

Description This command specifies which hashing algorithm to use for the IKE authentication function. The **no** form of the command returns the parameter to its default value.

Default sha1

Parameters **md5 —** specifies the hmac-md5 algorithm for authentication

sha1 — specifies the hmac-sha1 algorithm for authentication

sha256 — specifies the sha256 algorithm for authentication

sha384 — specifies the sha384 algorithm for authentication

sha512 — specifies the sha512 algorithm for authentication

auth-method

Syntax **auth-method psk**
no auth-method

Context config>ipsec>ike-policy

Description This command specifies the authentication method used with this IKE policy. Configuring the policy for pre-shared key (PSK) or **no auth-method** produces the same result since PSK is both the default value and the only option.

The **no** form of the command returns the parameter to its default value (psk).

Default no auth-method

Parameters **psk —** both the client and the gateway authenticate each other by a hash derived from a secret PSK. Both client and gateway must have the PSK. This works with both IKEv1 and IKEv2.

dh-group

Syntax **dh-group** {**1** | **2** | **5** | **14** | **15**}
no dh-group

Context config>ipsec>ike-policy

Description This command specifies which Diffie-Hellman group is used to calculate session keys:

- Group1: 768 bits
- Group2: 1024 bits
- Group5: 1536 bits
- Group14: 2048 bits
- Group15: 3072 bits

More bits provide a higher level of security but require more processing.

The **no** form of the command returns the parameter to its default value (Group2).

Default no dh-group (Group2)

dpd

Syntax **dpd** [**interval** *interval*] [**max-retries** *max-retries*] [**reply-only**]
no dpd

Context config>ipsec>ike-policy

Description This command controls the dead peer detection (DPD) mechanism to detect a dead IKE peer.

The **no** form of the command disables DPD and returns the parameters to their default values.

Default no dpd

Parameters *interval* — specifies the interval that will be used to test connectivity to the tunnel peer. If the peer initiates the connectivity check before the interval timer, it will be reset.

Values 10 to 300 s

Default 30

max-retries — specifies the maximum number of retries before the tunnel is removed

Values 2 to 5

Default 3

reply-only — specifies to only reply to DPD keepalives. Issuing the command without the **reply-only** keyword disables the reply-only behavior.

encryption-algorithm

Syntax **encryption-algorithm** {**des** | **3des** | **aes128** | **aes192** | **aes256**}
no encryption-algorithm

Context config>ipsec>ike-policy

Description This command specifies the encryption algorithm to use for the IKE session.

The **no** form of the command returns the algorithm to its default value (aes128).

Default aes128

Parameters **des** — configures the 56-bit **des** algorithm for encryption. This is an older algorithm, with relatively weak security. It should only be used when a strong algorithm is not available at both ends at an acceptable performance level.

3des — configures the **3-des** algorithm for encryption. This is a modified application of the **des** algorithm that uses multiple **des** operations for more security.

aes128 — configures the **aes** algorithm with a block size of 128 bits. This is the mandatory implementation size for **aes**.

aes192 — configures the **aes** algorithm with a block size of 192 bits. This is a stronger version of **aes**.

aes256 — configures the **aes** algorithm with a block size of 256 bits. This is the strongest available version of **aes**.

ipsec-lifetime

Syntax **ipsec-lifetime** *ipsec-lifetime*
no ipsec-lifetime

Context config>ipsec>ike-policy

Description This parameter specifies the lifetime of a phase 2 SA.

The **no** form of the command returns the *ipsec-lifetime* value to the default.

Default 3600 (1 hr)

Parameters *ipsec-lifetime* — specifies the lifetime of the phase 2 IKE key, in seconds

Values 1200 to 172800

isakmp-lifetime

Syntax **isakmp-lifetime** *isakmp-lifetime*
no isakmp-lifetime

Context config>ipsec>ike-policy

Description This command specifies the lifetime of a phase 1 SA. ISAKMP stands for Internet Security Association and Key Management Protocol.

The **no** form of the command returns the *isakmp-lifetime* value to the default value.

Default 86400

Parameters *isakmp-lifetime* — specifies the lifetime of the phase 1 IKE key, in seconds

Values 1200 to 172800

nat-traversal

Syntax **nat-traversal** [**force**] [**keep-alive-interval** *keep-alive-interval*] [**force-keep-alive**]
no nat-traversal

Context config>ipsec>ike-policy

Description This command specifies whether NAT-T (Network Address Translation Traversal) is enabled, disabled, or in force mode. Enabling NAT-T enables the NAT detection mechanism. If a NAT device is detected in the path between the 7705 SAR and its IPSec peer, then UDP encapsulation is done on the IPSec packet to allow the IPSec traffic to traverse the NAT device.

When **nat-traversal** is used without any parameters, NAT-T is enabled and sending keepalive packets is disabled (*keep-alive-interval* is 0 s).

When the **force** keyword is used, the IPSec tunnel always uses a UDP value in its header, regardless of whether a NAT device is detected.

The **force-keep-alive** keyword specifies whether keepalive packets are sent only when a NAT device is detected or are always sent (regardless of detection of a NAT device). When **force-keep-alive** is used, packets are always sent and the "Behind NAT Only" field in the **show>ipsec>ike-policy** *ike-policy-id* indicates False. When **force-keep-alive** is not used, packets are may or may not be sent, depending on the whether NAT-T is enabled or disabled. In this case, the "Behind NAT Only" field indicates True.

The **keep-alive-timer** keyword defines the frequency, where "0" means that keepalives are disabled.

The **no** form of the command returns the parameters to the default values (NAT-T is disabled, *keep-alive-interval* is 0 s, and **force-keep-alive** is True).

Default no nat-traversal

Parameters **force —** when specified, forces NAT-T to be enabled

keep-alive-interval **—** specifies the keepalive interval for NAT-T. If the value is 0 s, then keepalive messages are disabled.

Values 120 to 600 s

Default 0 s

force-keep-alive — specifies that NAT-T keepalive packets are always sent, regardless of NAT detection results

own-auth-method

Syntax **own-auth-method psk**
no own-auth-method

Context config>ipsec>ike-policy

Description This command specifies the authentication method used by the 7705 SAR OS to self-authenticate. This command (**own-auth-method**) applies only to IKEv2.

The default self-authentication method used by the 7705 SAR OS is symmetric, which means the self-authentication method is the same as the authentication method used by this IKE policy for the remote peer (that is, the **own-auth-method** is the same as **auth-method**).

The **no** form of the command returns the parameter to the default value (symmetric).

Default no own-auth-method

Parameters **psk —** specifies the use of a pre-shared key to self-authenticate

pfs

Syntax **pfs** [**dh-group** {**1** | **2** | **5**}]
no pfs

Context config>ipsec>ike-policy

Description This command enables Perfect Forward Secrecy (PFS) on the IPSec tunnel using this policy. PFS provides for a new Diffie-Hellman key exchange each time the SA key is renegotiated. After each SA expires, the key is forgotten and another key is generated (if the SA remains up). This means that an attacker who cracks part of the exchange can only read the part that used the key before the key changed. Thus, there is no advantage to cracking the other parts of the exchange if an attacker has already cracked one.

When **pfs** is used without the **dh-group** command, the default DH group (Group 2) is used.

The **no** form of the command disables PFS. If **pfs** is turned off during an active SA, then when the SA expires and it is time to re-key the session, the original Diffie-Hellman primes is used to generate the new keys.

Default no pfs

Parameters **dh-group {1 | 2 | 5} —** when **dh-group** is used, specifies which Diffie-Hellman group to use for calculating session keys. Higher dh-group values translate to higher level of security, but require more processing. Three groups are supported:

- Group 1: 768 bits
- Group 2: 1024 bits
- Group 5: 1536 bits

ipsec-transform

Syntax **ipsec-transform** *transform-id* [**create**]
no ipsec-transform *transform-id*

Context config>ipsec

Description This command enables the context to create an **ipsec-transform** policy. IPSec transform policies can be shared between IPSec tunnels by using the transform command.

IPSec transform policy assignments to a tunnel require the tunnel to be shut down.

The **no** form of the command removes the transform ID from the configuration.

Parameters *transform-id* — specifies a policy ID value to identify the IPSec transform policy

Values 1 to 2048

create — mandatory keyword required when creating an **ipsec-transform** policy. The **create** keyword requirement can be enabled/disabled in the **environment**>**create** context.

esp-auth-algorithm

Syntax **esp-auth-algorithm** {**null** | **md5** | **sha1** | **sha256** | **sha384** | **sha512**}
no esp-auth-algorithm

Context config>ipsec>transform

Description This command specifies which hashing algorithm should be used for the authentication function Encapsulating Security Payload (ESP). Both ends of a tunnel must share the same configuration parameters in order for the IPSec tunnel to enter the operational state.

The **null** keyword in this command and the **null** keyword in the **esp-encryption-algorithm** command are mutually exclusive.

The **no** form of the command returns the parameter to its default value.

Default sha1

Parameters **null —** a very fast algorithm specified in RFC 2410, which provides no authentication

md5 — configures ESP to use the **hmac-md5** algorithm for authentication

sha1 — configures ESP to use the **hmac-sha1** algorithm for authentication

sha256 — configures ESP to use the **sha256** algorithm for authentication

sha384 — configures ESP to use the **sha384** algorithm for authentication

sha512 — configures ESP to use the **sha512** algorithm for authentication

esp-encryption-algorithm

Syntax **esp-encryption-algorithm** {**null** | **des** | **3des** | **aes128** | **aes192** | **aes256**}
no esp-encryption-algorithm

Context config>ipsec>transform

Description This command specifies the encryption algorithm to use for the IPSec session. Encryption only applies to Encapsulating Security Payload (ESP) configurations.

For IPSec tunnels to come up, both ends of the IPSec tunnel (both private-side endpoints) must be configured with the same encryption algorithm. That is, the configuration for **vprn>if>sap> ipsec-tunnel transform** must match at both nodes.

The **null** keyword in this command and the **null** keyword in the **esp-auth-algorithm** command are mutually exclusive.

The **no** form of the command returns the parameter to its default value.

Default aes128

Parameters **null —** configures the high-speed null algorithm, which does nothing. This is the same as not having encryption turned on.

des — configures the 56-bit **des** algorithm for encryption. This is an older algorithm, with relatively weak security. Although slightly better than no encryption, it should only be used when a strong algorithm is not available at both ends at an acceptable performance level.

3des — configures the **3-des** algorithm for encryption. This is a modified application of the **des** algorithm that uses multiple **des** operations to make things more secure.

aes128 — configures the **aes** algorithm with a block size of 128 bits. This is the mandatory implementation size for **aes**. This is a very strong algorithm choice.

aes192 — configures the **aes** algorithm with a block size of 192 bits. This is a stronger version of **aes**.

aes256 — configures the **aes** algorithm with a block size of 256 bits. This is the strongest available version of **aes**.

Service Configuration Commands

ipsec

Syntax **ipsec**

Context config>service>vprn

Description This command enables the context to configure IPSec policies.

Default n/a

security-policy

Syntax **security-policy** *security-policy-id* [**create**]
no security-policy *security-policy-id*

Context config>service>vprn>ipsec

Description This command configures a security policy to use for an IPSec tunnel. An **entry** specifying local and remote IP addresses must be defined before the policy can be used.

The **no** form of the command removes the policy. Policy entries must be deleted before the policy can be removed.

Default n/a

Parameters *security-policy-id* — specifies an identifier value to be assigned to a security policy

Values 1 to 8192

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

entry

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

Context config>service>vprn>ipsec>sec-plcy

Description This command configures an IPSec security policy entry.

The **no** form of the command removes the entry.

Default n/a

Parameters *entry-id* — specifies an identifier value for the IPSec security policy entry

Values 1 to 16

create — mandatory keyword used to create the security policy entry. The **create** keyword requirement can be enabled/disabled in the **environment**>**create** context.

local-ip

Syntax **local-ip** {*ip-prefix I prefix-length* | *ip-prefix netmask* | **any**}
no local-ip

Context config>service>vprn>ipsec>sec-plcy>entry

Description This command configures the local (from the VPN) IP prefix/mask for the policy parameter entry.

Only one entry is necessary to describe a potential traffic flow. The **local-ip** and **remote-ip** commands can be defined only once. The system will evaluate the local IP as the source IP when traffic is examined in the direction of the VPN to the tunnel and as the destination IP when traffic flows from the tunnel to the VPN. The remote IP will be evaluated as the source IP when traffic flows from the tunnel, and as the destination IP to the VPN when traffic flows from the VPN to the tunnel.]

The **no** form of the command clears the IP entry.

Default no local-ip

Parameters *ip-prefix* — the destination address of the aggregate route in dotted-decimal notation

Values a.b.c.d (host bits must be 0) (0.0.0.0 is not allowed)
prefix-length: 1 to 32

netmask — the subnet mask in dotted-decimal notation

Values a.b.c.d (network bits all 1 and host bits all 0) (0.0.0.0 is not allowed

any — keyword to specify that it can be any address

local-v6-ip

Syntax **local-v6-ip** {*ipv6-prefix I prefix-length* | **any**}
no local-v6-ip

Context config>service>vprn>ipsec>sec-plcy>entry

Description This command configures the local (from the VPN) IPv6 address for the policy parameter entry.

Only one entry is necessary to describe a potential traffic flow. The **local-v6-ip** and **remote-v6-ip** commands can be defined only once. The system will evaluate the local IPv6 address as the source IPv6 address when traffic is examined in the direction of the VPN to the tunnel and as the destination IPv6 address when traffic flows from the tunnel to the VPN. The remote IPv6 address will be evaluated as the source IPv6 address when traffic flows from the tunnel to the VPN and as the destination IPv6 address when traffic flows from the VPN to the tunnel.

The **no** form of the command clears the IPv6 address entry.

Default no local-v6-ip

Parameters *ipv6-prefix / prefix-length* — the local IPv6 address

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 to FFFF]H
	d: [0 to 255]D
prefix-length	0 to 128

any — keyword to specify that it can be any address

remote-ip

Syntax **remote-ip** {*ip-prefix / prefix-length* | *ip-prefix netmask* | **any**}
no remote-ip

Context config>service>vprn>ipsec>sec-plcy>entry

Description This command configures the remote (from the tunnel) IP prefix/mask for the policy parameter entry.

Only one entry is necessary to describe a potential traffic flow. The **local-ip** and **remote-ip** commands can be defined only once. The system will evaluate the local IP as the source IP when traffic is examined in the direction of the VPN to the tunnel and as the destination IP when traffic flows from the tunnel to the VPN. The remote IP will be evaluated as the source IP when traffic flows from the tunnel to the VPN and as the destination IP when traffic flows from the VPN to the tunnel.

The **no** form of the command clears the IP entry.

Default no remote-ip

Parameters *ip-prefix* — specifies the destination address of the aggregate route in dotted-decimal notation

Values a.b.c.d (host bits must be 0) (0.0.0.0 is not allowed)
prefix-length: 1 to 32

netmask — the subnet mask in dotted-decimal notation

Values a.b.c.d (network bits all 1 and host bits all 0) (0.0.0.0 is not allowed

any — keyword to specify that it can be any address

remote-v6-ip

Syntax **remote-v6-ip** {*ipv6-prefix / prefix-length* | **any**}
no remote-v6-ip

Context config>service>vprn>ipsec>sec-plcy>entry

Description This command configures the remote (from the tunnel) IPv6 address for the policy parameter entry.

Only one entry is necessary to describe a potential traffic flow. The **local-v6-ip** and **remote-v6-ip** commands can be defined only once. The system will evaluate the local IPv6 address as the source IPv6 address when traffic is examined in the direction of the VPN to the tunnel and as the destination IPv6 address when traffic flows from the tunnel to the VPN. The remote IPv6 address will be evaluated as the source IPv6 address when traffic flows from the tunnel to the VPN and as the destination IPv6 address when traffic flows from the VPN to the tunnel.

The **no** form of the command clears the IPv6 address entry.

Default no remote-v6-ip

Parameters *ipv6-prefix* — the remote IPv6 address

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 to FFFF]H
	d: [0 to 255]D
prefix-length	0 to 128

any — keyword to specify that it can be any address

Service Interface Tunnel Commands

interface

Syntax **interface** *ip-int-name* [**tunnel**] [**create**]
no interface *ip-int-name*

Context config>service>vprn
config>service>ies

Description This command creates a logical IP routing interface.

The **tunnel** keyword is not used for IES interfaces. An IES public tunnel SAP is created when the *sap-id* includes the **tunnel** and **public** keywords (see sap below). For VPRN interfaces, **tunnel** is used to create an IP interface that supports a private tunnel SAP. The VPRN private tunnel SAP allows provisioning of an IPSec tunnel, and is created when the VPRN sap-id includes the **tunnel** and **private** keywords.

Default n/a

Parameters *ip-int-name* — specifies an IP interface name up to 32 characters in length

tunnel — specifies that the interface is a private tunnel

create — mandatory keyword required when creating an IP interface. The **create** keyword requirement can be enabled/disabled in the **environment**>**create** context.

sap

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

Context config>service>vprn>if
config>service>ies>if

Description This command creates a SAP. For IES and VPRN services using tunnel interfaces, the *sap-id* for private and public tunnel interfaces are shown below. See sap In the VLL Services Command Reference for details on configuring all SAPs.

Default n/a

Parameters *sap-id* — specifies the port identifier portion of the SAP definition. For a tunnel interface, the *sap-id* is as follows:

Values tunnel-*id*.[private | public]:*tag*

tunnel	keyword
id	1 to 16 (only the value 1 is allowed)
private	keyword
public	keyword
tag	0 to 4094

create — mandatory keyword required when creating a SAP. The **create** keyword requirement can be enabled/disabled in the **environment**>**create** context.

ipsec-tunnel

Syntax **ipsec-tunnel** *ipsec-tunnel-name* [**create**]
no ipsec-tunnel *ipsec-tunnel-name*

Context config>service>vprn>if>sap

Description This command specifies an IPSec tunnel name. Configuring the commands under the **ipsec-tunnel** context defines where the IPSec tunnel originates and terminates, and how it is secured.

Default n/a

Parameters *ipsec-tunnel-name* — specifies an IPSec tunnel name up to 32 characters in length

create — mandatory keyword required when creating an IPSec tunnel instance. The **create** keyword requirement can be enabled/disabled in the **environment**>**create** context.

bfd-designate

Syntax [**no**] **bfd-designate**

Context config>service>vprn>if>sap>ipsec-tunnel

Description This command specifies whether this IPSec tunnel is the BFD-designated tunnel.

A BFD-designated tunnel is the tunnel over which a BFD session is established. A BFD-designated tunnel does not go down when BFD goes down. Other tunnels that use that BFD-designated tunnel's BFD session will go down based on the state of the BFD session.

Default no bfd-designate

bfd-enable

Syntax **bfd-enable service** *service-id* **interface** *interface-name* **dst-ip** *ip-address*
no bfd-enable

Context config>service>vprn>if>sap>ipsec-tunnel

Description This command assigns a BFD session to provide the heart-beat mechanism for the specified IPSec tunnel. There can be only one BFD session assigned to any given IPSec tunnel, but there can be multiple IPSec tunnels using same BFD session. BFD controls the state of the associated tunnel; if the BFD session goes down, the system will also bring down the associated non-designated IPSec tunnel.

Default n/a

Parameters *service-id* — specifies the service ID or name where the BFD session resides

Values *service-id*: 1 to 2147483647 or *svc-name* (up to 64 characters)

interface — the name of the interface used by the BFD session

interface-name — specifies the interface name

Values 1 to 32 characters (must start with a letter)

ip-address — specifies the IPv4 destination address to be used for the BFD session

dst-ip — the IPv4 or IPv6 destination address to be used for the BFD session

ip-address **—** the IPv4 or IPv6 destination address

Values

IPv4 address:	a.b.c.d
IPv6 address:	x:x:x:x:x:x:x:x
	x - [0..FFFF]H
	d - [0..255] D

clear-df-bit

Syntax [**no**] **clear-df-bit**

Context config>service>vprn>if>sap>ipsec-tunnel

Description This command clears the do-not-fragment (DF) bit on incoming unencrypted IP traffic, allowing traffic to be fragmented, if necessary, before it enters the tunnel.

The **no** form of the command, corresponding to the default behavior, leaves the DF bit unchanged.

Default no clear-df-bit

copy-df-bit

Syntax [**no**] **copy-df-bit**

Context config>service>vprn>if>sap>ipsec-tunnel

Description This command specifies whether to copy the do-not-fragment (DF) bit from the customer clear traffic and insert it into the IPSec tunnel header of the outgoing packet. When disabled, the DF bit of the IPSec tunnel header is always set to 1 (do not copy the DF bit).

The **no** form of the command, corresponding to the default behavior, does not copy the customer DF bit to the IPSec tunnel header.

Default no copy-df-bit

dynamic-keying

Syntax [**no**] **dynamic-keying**

Context config>service>vprn>if>sap>ipsec-tunnel

Description This command enables dynamic keying for the IPSec tunnel. Dynamic keying means that the IKE protocol is used to dynamically exchange keys and establish IPSec-SAs. When IKE is used, a tunnel will have ISAKMP-SA for phase 1 (used by IKE) and IPSEC-SA for phase 2 (used for traffic encryption).

The **dynamic-keying** and **manual-keying** commands are mutually exclusive. One of these commands must be configured to make the tunnel operational.

The **no** form of the command returns the SA keying type to its default value.

Default no dynamic-keying

auto-establish

Syntax [**no**] **auto-establish**

Context config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying

Description This command specifies whether to attempt to establish a phase 1 exchange automatically. The **auto-establish** command should only be enabled on one side of the tunnel. A tunnel with **auto-establish** enabled acts as an IKE initiator and does not respond to a new phase 1 request.

The **no** form of the command disables the automatic attempts to establish a phase 1 exchange.

Default no auto-establish

ike-policy

Syntax **ike-policy** *ike-policy-id*
no ike-policy

Context config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying

Description This command configures the IKE policy for dynamic keying, which will be used by the tunnel.

The **no** form of the command removes the IKE policy.

Default no ike-policy

Parameters *ike-policy-id* — specifies the IKE policy ID

Values 1 to 2048

local-id

Syntax **local-id type** {**ipv4** | **fqdn**} **value** *value*
no local-id

Context config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying

Description This command allows the specification of the IKEv2 local ID value for a dynamic keyed IPSec tunnel. The allowed local ID types are a valid IPv4 address or an FQDN (fully qualified domain name) string.

If **local-id** is configured, the tunnel's local ID is set to the explicit **type** and **value** specified by the **local-id** command. If **local-id** is not configured, the tunnel's local gateway IP address is used in the ID field of IKEv2 (see local-gateway-address).

The **no** form of the command removes the local ID.

Default no local-id

Parameters **type —** specifies the type of local ID payload

Values **ipv4**: specifies IPv4 as the local ID type. The default value is the local gateway IP address.

fqdn: specifies FQDN as the local ID type. A value must be configured.

value — specifies an IPv4 address or an FQDN value. An address or value must be configured.

Values **ipv4**: a.b.c.d

fqdn: a fully qualified domain name value (for example, "myhost.sample.com"), up to 255 characters maximum

pre-shared-key

Syntax **pre-shared-key** *key*
no pre-shared-key

Context config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying

Description This command specifies the pre-shared key (PSK), or secret passphrase, that will be used to initiate the tunnel IKE session.

The **no** form of the command removes the pre-shared key.

Default no pre-shared-key

Parameters *key* — specifies a pre-shared-key for dynamic-keying, where the key is up to 64 characters

Values any ASCII character, including "!@#$%^&*()_+-:{}[]\|", when the entire key is enclosed in parentheses

transform

Syntax **transform** *transform-id* [*transform-id*...(up to 4 max)]
no transform

Context config>service>vprn>if>sap>ipsec-tunnel>dynamic-keying

Description This command associates the IPSec transform set allowed for this tunnel. A maximum of four transforms can be specified. The transforms are listed in decreasing order of preference (the first one specified is the most preferred). The list of *transform-ids* is overwritten each time the command is issued. Transforms are defined using the ipsec-transform command.

The **no** form of the command returns the command to its default state.

Default no transform

Parameters *transform-id* — specifies the value used for transforms for dynamic keying

Values 1 to 2048

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>vprn>if>sap>ipsec-tunnel

Description This command configures the IP maximum transmit unit (MTU) (packet) for this interface.

The **ip-mtu** command instructs the 7705 SAR to perform IP packet fragmentation prior to IPSec encryption and encapsulation, based on the configured MTU value.

On the 7705 SAR, unencrypted IP packets arriving on a VPRN access interface and destined for an IPSec uplink will be fragmented if the incoming packet is larger than:

- the VPRN private interface MTU
- the IPSec tunnel MTU
- the difference between the uplink MTU and the IPSec overhead (uplink interface MTU minus IPSec overhead), where the IPSec overhead values are calculated as follows:
  - → IPSec overhead if NAT-T is enabled

    IPSec overhead = outer IPSec (20) + UDP (8) + ESP (24) + trailer (17) + ICV (32)
    = 101 bytes
  - → IPSec overhead if NAT-T is disabled (`no nat-t`)

    IPSec overhead = outer IP (20) + ESP (24) + trailer (17) + ICV (32)
    = 93 bytes

The actual overhead depends on the payload size, and the encryption and authentication algorithms used.

The **no ip-mtu** command, corresponding to the default behavior, disables fragmentation of IP packets by the 7705 SAR; all IP packets, regardless of size or DF bit setting, are allowed into the tunnel.

Default no ip-mtu

Parameters *octets* — specifies the MTU for the IP packet, expressed as the number of octets

Values 512 to 9732

local-gateway-address

Syntax **local-gateway-address** *ip-address* **peer** *ip-address* **delivery-service** *service-id*
no local-gateway-address

Context config>service>vprn>if>sap>ipsec-tunnel

Description This command specifies the local gateway address used by the tunnel and the remote gateway address at the other end of the tunnel.

The local gateway address is the source address of the outgoing encrypted packet and the peer gateway address is the destination address. The delivery service is the IES service that has the corresponding public tunnel interface configured under it.

The local gateway address must be in the same subnet as the public tunnel interface.

Parameters *ip-address* — IPv4 address of the local and peer ends of the tunnel

Values a.b.c.d

service-id — specifies the ID of the IES or VPRN (front-door) delivery service of this tunnel. Use this *service-id* to find the VPRN used for delivery.

Values *service-id*: 1 to 2147483647 or *svc-name*, which specifies an existing service name up to 64 characters in length

manual-keying

Syntax **[no] manual-keying**

Context config>service>vprn>if>sap>ipsec-tunnel

Description This command allows manual configuration of tunnel Security Associations. Manual keying can be used in lieu of dynamic keying and IKE.

The **dynamic-keying** and **manual-keying** commands are mutually exclusive. One of these commands must be configured to make the tunnel operational.

When manual keying is used, both encryption and authentication must be entered manually for inbound and outbound SAs. Encryption and authentication modes, along with associated keys, must match on both sides of the tunnel. Inbound SA configuration on the near-end system must match outbound SA configuration on the far-end system, and vice versa. Make sure to use the correct key length, based on the **ipsec-transform** configuration.

A configuration example for manual keying is shown below:

Example:
```
ipsec-transform 2 create
   esp-auth-algorithm sha512
   esp-encryption-algorithm aes128
exit
ipsec-tunnel "privateTunnel" create
   security-policy 4
   local-gateway-address 10.1.1.2 peer 10.3.3.2
 delivery-service 100
   manual-keying
      security-association 8 direction inbound spi 500
        transform 2 encryption-key 5253c408a123817358
        authentication-key 0x1c4a94f71e5366f3760863
      security-association 8 direction outbound spi 600
        transform 2 encryption-key 0xe9ffb43d2ddd
        authentication-key 0x1db443f855693f0fe45d
      exit
      no shutdown
   exit
```

The **no** form of the command returns the SA keying type to its default value.

Default no manual-keying

security-association

Syntax **security-association** *security-entry-id* **authentication-key** *authentication-key* **encryption-key** *encryption-key* **spi** *spi* **transform** *transform-id* **direction** {**inbound** | **outbound**}
no security-association *security-entry-id* **direction** {**inbound** | **outbound**}

Context config>service>vprn>if>sap>ipsec-tunnel>manual-keying

Description This command configures the information required for manual keying SA creation.

Default n/a

Parameters *security-entry-id* — specifies the ID of an SA entry

Values 1 to 16

authentication-key — specifies the key used for the authentication algorithm

Values none or 0x0..0xFFFFFFFF...(max 128 hex nibbles)

encryption-key — specifies the key used for the encryption algorithm

Values none or 0x0..0xFFFFFFFF...(max 64 hex nibbles)

spi — specifies the SPI (Security Parameter Index) used to look up the instruction to verify and decrypt the incoming IPSec packets when the direction is inbound. When the direction is outbound, the SPI will be used in the encoding of the outgoing packets. The remote node can use this SPI to look up the instruction to verify and decrypt the packet.

Values 256 to 16383

transform-id — specifies the transform entry that will be used by this SA entry. This object should be specified for all the entries created that are manual SAs.

Values 1 to 2048

direction {inbound | outbound} — specifies the direction of the IPSec tunnel

security-policy

Syntax **security-policy** *security-policy-id*
no security-policy

Context config>service>vprn>if>sap>ipsec-tunnel

Description This command identifies an IPSec security policy (defined under the **vprn>ipsec** context) that is to be used for this IPSec tunnel.

The **no** form of the command returns the **security-policy** to its default state (n/a).

Default n/a

Parameters *security-policy-id* — specifies the IPSec security policy that the tunnel will use

Values 1 to 8192

PKI Configuration Commands

- X.509 and Certificate Commands
- PKI Infrastructure Commands
- IPSec PKI Commands
- IKE PKI Commands

X.509 and Certificate Commands

clear-ocsp-cache

Syntax **clear-ocsp-cache** [*entry-id*]

Context admin>certificate

Description This command clears the current OCSP response cache. If the optional issuer and serial number are not specified, then all current cached results are cleared.

Parameters *entry-id* — the local cache entry identifier of the certificate to clear

Values 1 to 2000

cmpv2

Syntax **cmpv2**

Context admin>certificate

Description This command enables the context to configure CMPv2 parameters. Changes are not allowed when the CA profile is enabled (**no shutdown**).

cert-request

Syntax **cert-request ca** *ca-profile-name* **current-key** *key-filename* **current-cert** *cert-filename* [**hash-alg** *hash-algorithm*] **newkey** *key-filename* **subject-dn** *subject-dn* **save-as** *save-path-of-result-cert*

Context admin>certificate>cmpv2

Description This command requests an additional certificate after the system has obtained the initial certificate from the CA.

The request is authenticated by a signature signed by the **current-key**, along with the **current-cert**. The hash algorithm used for the signature depends on the key type:

- DSA key: SHA1
- RSA key: MD5 | SHA1 | SHA224 | SHA256 | SHA384 | SHA512; the default is SHA1

In some cases, the CA may not return a certificate immediately, due to reasons such as the request processing needs manual intervention. In such cases, the **admin certificate cmpv2 poll** command can be used to poll the status of the request.

Default n/a

Parameters *ca-profile-name* — specifies a certificate authority profile name that includes CMP server information, up to 32 characters

current-key *key-filename* **—** the corresponding certificate issued by the CA, up to 95 characters

cert-filename — the filename of an imported certificate that is attached to the certificate request, up to 95 characters

newkey *key-filename* **—** the filename of the imported key, up to 95 characters.

hash-algorithm **—** the hash algorithm for the RSA key

Values md5, sha1, sha224, sha256, sha384,s ha512

dn — the subject of the requesting certificate, up to 256 characters

Values attr1=val1,attr2=val2... where: attrN = {C | ST | O | OU | CN}

save-path-of-result-cert — the full path name to save the result certificate to, up to 200 characters

clear-request

Syntax **clear-request ca** *ca-profile-name*

Context admin>certificate>cmpv2

Description This command clears current pending CMPv2 requests toward the specified CA. If there are no pending requests, it will clear the saved results of prior requests.

Default n/a

Parameters *ca-profile-name* — a CA profile name, up to 32 characters

initial-registration

Syntax **initial-registration ca** *ca-profile-name* **key-to-certify** *key-filename* **protection-alg** {**password** *password* **reference** *ref-number* | **signature** [**cert** *cert-file-name* [**send-chain** [**with-ca** *ca-profile-name*]]] [**protection-key** *key-file-name*] [**hash-alg** {**md5** | **sha1** | **sha224** | **sha256** | **sha384** | **sha512**}]} **subject-dn** *dn* **save-as** *save-path-of-result-cert*

Context admin>certificate>cmpv2

Description This command requests the initial certificate from the CA by using the CMPv2 initial registration procedure.

The **ca** keyword specifies a **ca-profile** that includes CMP server information.

The **key-to-certify** keyword is an imported key file to be certified by the CA.

The **protection-key** keyword is an imported key file used to for message protection if **protection-alg** is configured as **signature**.

The request is authenticated using either of the following methods:

- a password and a reference number that is predistributed by the CA using out-of-band means. The specified password and reference number are not necessarily in the CMP **key-list** configured in the corresponding **ca-profile**.
- a signature signed by the **protection-key** or **key-to-certify**, optionally along with the corresponding certificate. If the **protection-key** is not specified, the system will use the **key-to-certify** keyword for message protection. The hash algorithm used for the signature depends on the key type:
  - → DSA key: SHA1
  - → RSA key: MD5 | SHA1 | SHA224 | SHA256 | SHA384 | SHA512; the default is SHA1

Optionally, the system could also send a certificate or a chain of certificates in the extraCerts field. The certificate is specified by the **cert** *cert-file-name* parameter; it must include the public key of the key used for message protection.

Sending a chain is enabled by specifying the **send-chain** keyword.

The **subject-dn** keyword specifies the subject of the requesting certificate.

The **save-as** keyword specifies the full path name to save the result certificate to.

In some cases, the CA may not return the certificate immediately; for example, because the request processing requires manual intervention. In such cases, the **admin certificate cmpv2** poll command could be used to poll the status of the request. If the **key-list** command is not configured in the corresponding **ca-profile**, then the system will use the existing password to authenticate the CMPv2 packets from the server if it is in password protection.

If **key-list** is configured in the corresponding **ca-profile** and the server does not send a SenderKID message, then the system will use the lexicographical first key in the **key-list** to authenticate the CMPv2 packets from the server in case it is in password protection.

Default n/a

Parameters **ca** *ca-profile-name* **—** specifies a certificate authority profile name that includes CMP server information, up to 32 characters

key-to-certify *key-filename* **—** the filename of the key to certify, up to 95 characters

password — an ASCII string, up to 64 characters

ref-number — the reference number for this CA initial authentication key, up to 64 characters

cert-file-name — specifies the certificate filename, up to 95 characters

send-chain with-ca *ca-profile-name* **—** sends the chain

protection-key *key-file-name* **—** the protection key associated with the action on the CA profile

hash-alg — the hash algorithm for the RSA key

Values md5, sha1, sha224, sha256, sha384, sha512

dn — the subject of the requesting certificate, up to 256 characters

Values attr1=val1,attr2=val2... where: attrN = {C | ST | O | OU | CN}

save-path-of-result-cert — the full path name to save the result certificate to, up to 200 characters

key-update

Syntax **key-update ca** *ca-profile-name* **newkey** *key-filename* **oldkey** *key-filename* **oldcert** *cert-filename* [**hash-alg** *hash-algorithm*] **save-as** *save-path-of-result-cert*

Context admin>certificate>cmpv2

Description This command requests a new certificate from the certificate authority to update an existing certificate.

In some cases, the CA may not return a certificate immediately; for example, because the request processing requires manual intervention. In such cases, the **admin>certificate>cmpv2>poll** command can be used to poll the status of the request.

Parameters *ca-profile-name* — specifies a certificate authority profile name that includes CMP server information, up to 32 characters

newkey *key-filename* **—** the key file of the requesting certificate, up to 95 characters

oldkey *key-filename* **—** the key to be replaced, up to 95 characters

cert-filename — the filename of an imported certificate to be replaced, up to 95 characters

hash-algorithm — the hash algorithm for the RSA key

Values md5, sha1, sha224, sha256, sha384, sha512

save-path-of-result-cert — the full path name to save the result certificate to, up to 200 characters

poll

Syntax **poll ca** *ca-profile-name*

Context admin>certificate>cmpv2

Description This command polls the status of the pending CMPv2 request toward the specified CA.

If the response is ready, this command will resume the CMPv2 protocol exchange with the server as the original command would do. If the request is still pending, then this command could be used again to poll the status.

The 7705 SAR allows only one pending CMP request per CA, which means that no new request is allowed when a pending request is present.

Default n/a

Parameters *ca-profile-name* — specifies a CA profile name, up to 32 characters

show-request

Syntax **show-request** [**ca** *ca-profile-name*]

Context admin>certificate>cmpv2

Description This command displays the current CMPv2 pending request toward the specified CA. If there is no pending request, the last pending request is displayed including the status (one of success, fail, or rejected) and the receive time of the last CMPv2 message from the server.

The following information is included in the output:

- request type
- original input parameter (password is not displayed)
- checkAfter and reason of last PollRepContent
- time of original command input

Default n/a

Parameters *ca-profile-name* — specifies a CA profile, up to 32 characters. If not specified, the system will display the pending requests of all CA profile.

display

Syntax **display type** {**cert** | **key** | **crl** | **cert-request**} *url-string* **format** {**pkcs10** | **pkcs12** | **pkcs7-der** | **pkcs7-pem** | **pem** | **der**} [**password** *password*]

Context admin>certificate

Description This command displays the contents of an input file in plain text. When displaying the key file contents, only the key size and type are displayed.

The following list summarizes the formats supported by this command:

- System
  - → system format
  - → PKCS #12
  - → PKCS #7 PEM encoded
  - → PKCS #7 DER encoded
  - → RFC4945
- Certificate Request
  - → PKCS #10
- Key
  - → system format
  - → PKCS #12

- CRL
  - → system format
  - → PKCS #7 PEM encoded
  - → PKCS #7 DER encoded
  - → RFC4945

Default n/a

Parameters *url-string* — the local compact flash URL of the input file

Values url-string <local-url> - [99 chars max]
local-url <cflash-id>/<file-path>
cflash-id cf3:

type — the type of input file; possible values are cert, key, crl, or cert-request

Values cert, key, crl, cert-request

format — the format of the input file

Values pkcs10, pkcs12, pkcs7-der, pkcs7-pem, pem, der

password — the password to decrypt the input file if it is an encrypted PKCS# 12 file, up to 32 characters

export

Syntax **export type** {**cert** | **key** | **crl**} **input** *filename* **output** *url-string* **format** *output-format* [**password** *password*] [**pkey** *pkey-filename*]

Context admin>certificate

Description This command performs certificate operations.

gen-keypair

Syntax **gen-keypair** *url-string* [**size** {**512** | **1024** | **2048**}] [**type** {**rsa** | **dsa**}]

Context admin>certificate

Description This command generates an RSA or DSA private key/public key pair and stores it in a local file in the cf3:\system-pki\key directory.

Parameters *url-string* — the name of the key file

Values url-string <local-url> - [99 chars max]
local-url <cflash-id>/<file-path>
cflash-id cf3:

size — the key size in bits (the minimum key size is 1024 bits when running in FIPS-140-2 mode)

Values 512, 1024, or 2048

Default 2048

type — the type of key

Values rsa dsa

Default rsa

gen-local-cert-req

Syntax **gen-local-cert-req keypair** *url-string* **subject-dn** *subject-dn* [**domain-name** *domain-name*] [**ip-addr** {*ip-address* | *ipv6-address*}] **file** *url-string* [**hash-alg** *hash-algorithm*] [**use-printable**]

Context admin>certificate

Description This command generates a PKCS# 10 formatted certificate request by using a local existing key pair file.

Default n/a

Parameters *url-string* — the name of the key file in cf3:\system-pki\key that is used to generate a certificate request

Values url-string<local-url> - [99 chars max]
local-url <cflash-id>/<file-path>
cflash-id cf3:

subject-dn — the distinguishing name that is used as the subject in a certificate request, including:

- C – Country
- ST – State
- O – Organization name
- OU – Organization Unit name
- CN – common name

This parameter is formatted as a text string including any of the above attributes. The attribute and its value are linked by using "=", and "," is used to separate different attributes.

For example: C=US,ST=CA,O=ALU,CN=SR12

Values attr1=val1,attr2=val2... where: attrN = {C | ST | O | OU | CN}, up to 256 characters

domain-name — optionally, a domain name string can be specified and included as the dNSName in the Subject Alternative Name extension of the certificate request, up to 255 characters

ip-address | *ipv6-address* — optionally, an IPv4 or IPv6 address string can be specified and included as the ipAddress in the Subject Alternative Name extension of the certificate request

url-string — a local compact flash path and filename to save the certificate request to, or an FTP URL to upload the certificate request

hash-algorithm — the hash algorithm to be used in a certificate request

Values sha1, sha224, sha256, sha384, sha512

use-printable — encodes the certificate in printable text format instead of in UTF8

import

Syntax **import type** {**cert** | **key** | **crl**} **input** *url-string* **output** *filename* **format** *input-format* [**password** *password*]

Context admin>certificate

Description This command converts an input file (either key, certificate, or CRL) to a system format file. The following list summarizes the formats supported by this command.

- Certificate
  - → PKCS #12
  - → PKCS #7 PEM encoded
  - → PKCS #7 DER encoded
  - → PEM
  - → DER
- Key
  - → PKCS #12
  - → PEM
  - → DER
- CRL
  - → PKCS #7 PEM encoded
  - → PKCS #7 DER encoded
  - → PEM
  - → DER

If there are multiple objects with same type in the input file, only the first object will be extracted and converted.

Default n/a

Parameters **input** *url-string* **—** the URL for the input file. This URL could be either a local compact flash URL file or an FTP URL to download the input file.

output *filename* **—** the name of output file, up to 95 characters in length. The output directory depends on the file type:

- Key: cf3:\system-pki\key
- Cert: cf3:\system-pki\cert

- CRL: cf3:\system-pki\CRL

Values url-string<local-url> - [99 chars max]
local-url <cflash-id>/<file-path>
cflash-id cf3:

type — the type of input file

Values cert, key, crl

input-format — the format of the input file

Values pkcs12, pkcs7-der, pkcs7-pem, pem, der

password — the password to decrypt the input file if it is an encrypted PKCS# 12 file, up to 32 characters

reload

Syntax **reload type** {**cert** | **key**} *filename* [**key-file** *filename*]

Context admin>certificate

Description This command reloads an imported certificate or key file or both at the same time. This command is typically used to update a certificate and/or key file without shutting down the IPSec tunnel, **cert-profile**, or **ca-profile**.

- If the new file exists and is valid, then for each tunnel using it:
  - → if the key matches the certificate, then the new file will be downloaded to the 7705 SAR to be used the next time. Tunnels currently up are not affected.
  - → if the key does not match the certificate:
    - if the **cert** and **key** configuration is used instead of **cert-profile**, then the tunnel will be brought down
    - if the **cert-profile** is used, then **cert-profile** will be brought down. The next authentication will fail but the established tunnels are not affected.

If the new file does not exist or is invalid, then this command will abort.

Default n/a

Parameters **cert —** reload a certificate file

key — reload a key file

filename — the filename of the imported certificate or key

key-file *filename* **—** the imported key file

PKI Infrastructure Commands

pki

Syntax **pki**

Context config>system>security

Description This command enables the context to configure certificate parameters.

Default n/a

ca-profile

Syntax **ca-profile** *name* [**create**]
no ca-profile *name*

Context config>system>security>pki

Description This command creates a new certificate authority profile or enters the configuration context of an existing certificate authority profile. Up to 128 CA profiles can be created in the system. A **shutdown** of the **ca-profile** will not affect the current up and running **ipsec-tunnel** associated with the **ca-profile**; however, subsequent authentication will fail.

Executing a **no shutdown** command in this context will cause the system to reload the configured **cert-file** and **crl-file**.

A **ca-profile** can be applied under the **ipsec-tunnel** configuration.

The **no** form of the command removes the name parameter from the configuration. A CA profile cannot be removed until all the associations (IPSec tunnels) have been removed.

Parameters *name* — the name of the **ca-profile**, a string up to 32 characters

create — the keyword used to create a new **ca-profile**. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

cert-file

Syntax **cert-file** *filename*
no cert-file

Context config>system>security>pki>ca-profile

Description This command specifies the name of a file in the cf3:\system-pki\cert directory as the CA's certificate of the CA profile.

The system performs the following checks against a configured **cert-file** when a **no shutdown** command is issued.

- The configured **cert-file** is a DER-formatted X.509v3 certificate file.
- All mandatory fields defined in section 4.1 of RFC 5280 exist and conform to the RFC 5280-defined format.
- The Version field has a value of 0x2.
- The Validity field indicates that the certificate is still valid.
- The X.509 basic constraints extension exists, and the CA Boolean is true.
- If the Key Usage extension exists, then at least keyCertSign and cRLSign are asserted.
- If the certificate is not a self-signing certificate, then the system will look for the issuer's CA's certificate to verify that this certificate is signed by the issuer's CA. If there is no such CA profile configured, then the system will just proceed with a warning message.
- If the certificate is not a self-signing certificate, then the system will look for the issuer's CA's CRL to verify that it has not been revoked. If there is no such CA profile configured or there is no such CRL, then the system will just proceed with a warning message.

If any of above checks fails, then the **no shutdown** command will fail.

Changing or removing the **cert-file** is only allowed when the **ca-profile** is in a shutdown state.

The **no** form of the command removes the filename from the configuration.

Parameters *filename* — the local compact flash file URL

cmpv2

Syntax **cmpv2**

Context config>system>security>pki>ca-profile

Description This command enables the context to configure CMPv2 parameters. Changes are not allowed when the CA profile is enabled (**no shutdown**).

accept-unprotected-errormsg

Syntax [**no**] **accept-unprotected-errormsg**

Context config>system>security>pki>ca-profile>cmpv2

Description This command enables the system to accept both protected and unprotected CMPv2 error messages. Without this command, the system will accept only protected error messages.

The **no** form of the command causes the system to accept only protected PKI error messages.

Default no accept-unprotected-errormsg

accept-unprotected-pkiconf

Syntax [**no**] **accept-unprotected-pkiconf**

Context config>system>security>pki>ca-profile>cmpv2

Description This command enables the system to accept both protected and unprotected CMPv2 PKI confirmation messages. Without this command, the system will accept only protected PKI confirmation messages.

The **no** form of the command causes the system to accept only protected PKI confirmation messages.

Default n/a

always-set-sender-for-ir

Syntax [**no**] **always-set-sender-for-ir**

Context config>system>security>pki>ca-profile>cmpv2

Description This command specifies to always set the sender field in the CMPv2 header of all Initial Registration (IR) messages with the subject name. By default, the sender field is only set if an optional certificate is specified in the CMPv2 request.

Default no always-set-sender-for-ir

http-response-timeout

Syntax **http-response-timeout** *timeout*
no http-response-timeout

Context config>system>security>pki>ca-profile>cmpv2

Description This command specifies the timeout value for the HTTP response that is used by CMPv2.

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

Default 30 s

Parameters *timeout* — the HTTP response timeout in seconds

Values 1 to 3600

http-version

Syntax **http-version** {**1.0** | **1.1**}

Context config>system>security>pki>ca-profile>cmpv2

Description This command configures the HTTP version for CMPv2 messages.

Default 1.1

key-list

Syntax **key-list**

Context config>system>security>pki>ca-profile>cmpv2

Description This command enables the context to configure pre-shared key list parameters.

key

Syntax **key** *password* [**hash** | **hash2**] **reference** *reference-number*
no key reference *reference-number*

Context config>system>security>pki>ca-profile>cmpv2>key-list

Description This command specifies a pre-shared key used for CMPv2 initial registration. Multiples of key commands are allowed to be configured under this context.

The *password* and *reference-number* parameters are distributed by the CA using out-of-band means.

The configured password is stored in a configuration file in an encrypted form by using a 7705 SAR hash2 algorithm.

The **no** form of the command removes the parameters from the configuration.

Default n/a

Parameters *password* — a printable ASCII string, up to 64 characters

hash — specifies that the given password is already hashed using hashing algorithm version 1. A semantic check is performed on the given password field to verify that it is a valid hash 1 key to store in the database.

hash2 — specifies that the given password is already hashed using hashing algorithm version 2. A semantic check is performed on the given password field to verify that it is a valid hash 2 key to store in the database.

reference-number — Specifies a printable ASCII string, up to 64 characters.

response-signing-cert

Syntax **response-signing-cert** *filename*
no response-signing-cert

Context config>system>security>pki>ca-profile>cmpv2

Description This command specifies an imported certificate that is used to verify the CMP response messages if they are protected by a signature. If this command is not configured, then the CA's certificate is used.

Default n/a

Parameters *filename* — the filename of the imported certificate

same-recipnonce-for-pollreq

Syntax [**no**] **same-recipnonce-for-pollreq**

Context config>system>security>pki>ca-profile>cmpv2

Description This command enables the system to use the same recipNonce as the last CMPv2 response for a poll request.

Default n/a

url

Syntax **url** *url-string* [**service-id** *service-id*]
no url

Context config>system>security>pki>ca-profile>cmpv2

Description This command specifies the HTTP URL of the CMPv2 server. The URL must be unique across all configured CA profiles.

The URL will be resolved by the DNS server configured (if configured) in the corresponding router context.

If the *service-id* is 0 or omitted, then the system will try to resolve the FQDN using the DNS server configured in bof.cfg. After resolution, the system will first connect to the address in the management routing instance, then to the address in the base routing instance.

If the service is VPRN, then the system only allows HTTP ports 80 and 8080.

Default n/a

Parameters *url-string* — Specifies the HTTP URL of the CMPv2 server, up to 180 characters.

service-id *service-id* **—** the service or router instance that is used to reach the CMPv2 server

Values service-id: 1 to 2147483647
base-router: 0

crl-file

Syntax **crl-file** *filename*
no crl-file

Context config>system>security>pki>ca-profile

Description This command specifies the name of a file in the cf3:\system-pki\crl directory as the Certification Revoke List file of the **ca-profile**.

The system performs the following checks against a configured **crl-file** when a **no shutdown** command is issued.

- A valid **cert-file** of the **ca-profile** is already configured.
- A configured **crl-file** is a DER-formatted CRLv2 file.
- All mandatory fields defined in section 5.1 of RFC 5280 exist and conform to the RFC 5280-defined format.
- The version field has a value of 0x1.
- The delta CRL Indicator does not exist (delta CRL is not supported).
- The CRL's signature is verified by using the **cert-file** of the **ca-profile**.

If any of above checks fail, the **no shutdown** command will fail.

Changing or removing the **crl-file** is only allowed when the **ca-profile** is in a **shutdown** state.

The **no** form of the command removes the filename from the configuration.

Default n/a

Parameters *filename* — the name of CRL file stored in cf3:\system-pki\crl

description

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

Context config>system>security>pki>ca-profile

Description This command configures a description of the specified CA profile.

Default n/a

Parameters *description-string* — describe the CA profile, up to 80 characters

ocsp

Syntax **ocsp**

Context config>system>security>pki>ca-profile

Description This command enables the context to configure OCSP parameters.

responder-url

Syntax **responder-url** *url-string*
no responder-url

Context config>system>security>pki>ca-profile>ocsp

Description This command specifies the HTTP URL of the OCSP responder for the CA. This URL will only be used if there is no OCSP responder defined in the AIA extension of the certificate to be verified.

Default no responder-url

Parameters *url-string* — the HTTP URL of the OCSP responder

service

Syntax **service** *service-id*
no service

Context config>system>security>pki>ca-profile>ocsp

Description This command specifies the service or routing instance that is used to contact the OCSP responder. This applies to OCSP responders that are either configured in the CLI or defined in the AIA extension of the certificate to be verified.

The **responder-url** is resolved by using the DNS server configured in the configured routing instance.

For a VPRN service, the system verifies that the specified *service-id* or *service-name* is an existing VPRN service at the time of CLI configuration; if it is not, the configuration will fail.

Parameters *service-id* — specifies an existing service or router instance to be used in the match criteria

Values service-id: 1 to 2147483647
base-router: 0

shutdown

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

Context config>system>security>pki>ca-profile>
config>ipsec>cert-profile

Description This command enables or disables the **ca-profile**. The system will verify the configured **cert-file** and **crl-file**. If the verification fails, then the **no shutdown** command will fail.

A **ca-profile** in a **shutdown** state cannot be used in certificate authentication.

In the **config>ipsec>cert-profile** context, this command enables or disables the certificate profile.

Default shutdown

certificate-display-format

Syntax **certificate-display-format** {**ascii** | **utf8**}

Context config>system>security>pki

Description This command specifies the display format used for the Certificates and Certificate Revocation Lists.

Default ascii

Parameters **ascii —** use the ASCII format for the Certificates and Certificate Revocation Lists

utf8 — use the UTF8 format for the Certificates and Certificate Revocation Lists

maximum-cert-chain-depth

Syntax **maximum-cert-chain-depth** *level*
no maximum-cert-chain-depth

Context config>system>security>pki

Description This command defines the maximum depth of certificate chain verification. This value is applied system-wide.

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

Default 7

Parameters *level* — specifies the maximum depth of certificate chain verification. The certificate under verification is not counted in the chain. For example, if this parameter is set to 1, then the certificate under verification must be directly signed by the trust anchor CA.

Values 1 to 7

IPSec PKI Commands

cert-profile

Syntax **cert-profile** *profile-name* [**create**]
no cert-profile *profile-name*

Context config>ipsec

Description This command creates a new certificate profile or enters the configuration context of an existing certificate profile.

The **no** form of the command removes the profile name from the **cert-profile** configuration.

Default n/a

Parameters *profile-name* — the name of the certificate profile, up to 32 characters in length

entry

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

Context config>ipsec>cert-profile

Description This command configures an entry for the specified certificate profile.

The **no** form of the command removes the specified entry from the specified **cert-profile**.

Default n/a

Parameters *entry-id* — the entry ID

Values 1 to 8

cert

Syntax **cert** *cert-filename*
no cert

Context config>ipsec>cert-profile>entry

Description This command configures an imported certificate for the **cert-profile** entry.

The **no** form of the command removes the *cert-filename* from the entry configuration.

Default n/a

Parameters *cert-filename* — the name of the imported certificate, up to 32 characters in length

key

Syntax **key** *key-filename*
no key

Context config>ipsec>cert-profile>entry

Description This command configures an imported key for the **cert-profile** entry.

The **no** form of the command removes the *key-filename* from the entry configuration.

Default n/a

Parameters *key-filename* — the filename of an imported key

send-chain

Syntax [**no**] **send-chain**

Context config>ipsec>cert-profile>entry

Description This command enters the configuration context of send-chain in the **cert-profile** entry.

This command is optional. By default, the system sends the certificate specified by the **cert** command in the selected entry to the peer. This command allows the system to send additional CA certificates to the peer. These additional CA certificates must be in the certificate chain of the certificate specified by the **cert** command in the same entry.

ca-profile

Syntax [**no**] **ca-profile** *name*

Context config>ipsec>cert-profile>entry>send-chain

Description This command specifies that a certificate authority (CA) certificate in the specified **ca-profile** is to be sent to the peer.

Multiple configurations (up to seven) of this command are allowed in the same entry.

Default n/a

Parameters *name* — the profile name, up to 32 characters in length

ike-policy

Syntax **ike-policy** *ike-policy-id* [**create**]
no ike-policy *ike-policy-id*

Context config>ipsec

Description This command enables the context to configure an IKE policy.

The **no** form of the command deletes the IKE policy.

Parameters *ike-policy-id* — specifies a policy ID value to identify the IKE policy

Values 1 to 2048

auth-method

Syntax **auth-method** {**psk** | **cert-auth**}
no auth-method

Context config>ipsec>ike-policy

Description This command specifies the authentication method used with this IKE policy.

The **no** form of the command removes the parameter from the configuration.

Default no auth-method

Parameters **psk —** both the client and gateway authenticate each other by a hash derived from a pre-shared secret. Both client and gateway must have the PSK. This works with both IKEv1 and IKEv2.

cert-radius — use the certificate, public/private key and RADIUS to authenticate. This parameter applies to IKEv2 remote-access tunnel only.

own-auth-method

Syntax **own-auth-method** {**psk** | **cert**}
no own-auth-method

Context config>ipsec>ike-policy

Description This command configures the authentication method used with this IKE policy on its own side.

trust-anchor-profile

Syntax **trust-anchor-profile** *name* [**create**]
no trust-anchor-profile *name*

Context config>ipsec

Description This command specifies the **trust-anchor-profile** for the IPSec tunnel. This command will override the **trust-anchor-profile** configuration in the **config>service>vprn>if>sap>ipsec-tunnel>cert** context.

Default no trust-anchor-profile

Parameters *profile-name* — the **trust-anchor-profile** name

IKE PKI Commands

cert-profile

Syntax **cert-profile** *profile-name* [**create**]
no cert-profile *profile-name*

Context config>service>vprn>if>sap>ipsec-tunnel>cert

Description This command creates a new certificate profile or enters the configuration context of an existing certificate profile.

The **no** form of the command removes the profile name from the **cert-profile** configuration.

Default n/a

Parameters *profile-name* — the name of the certificate profile, up to 32 characters in length

status-verify

Syntax **status-verify**

Context config>service>vprn>if>sap>ipsec-tunnel>cert

Description This command enters the context to configure verification parameters for certificate revocation status.

Default n/a

default-result

Syntax **default-result** {**revoked** | **good**}
no default-result

Context config>service>vprn>if>sap>ipsec-tunnel>cert>status-verify

Description This command specifies the default result when both the primary and secondary methods fail to provide an answer.

Default revoked

Parameters **good —** the certificate is considered acceptable

revoked — the certificate is considered revoked

primary

Syntax **primary** {**crl** | **ocsp**}
no primary

Context config>service>vprn>if>sap>ipsec-tunnel>cert>status-verify

Description This command configures the primary method used to verify the revocation status of the peer's certificate. The method can be either CRL or OCSP.

To verify the revocation status of the peer's certificate, the CRL or OCSP uses the corresponding configuration in the CA profile of the issuer of the certificate in question.

Default crl

Parameters **crl —** the CRL file is configured in the corresponding CA profile

ocsp — the OCSP server is configured in the corresponding CA profile

secondary

Syntax **secondary** {**crl** | **ocsp**}
no secondary

Context config>service>vprn>if>sap>ipsec-tunnel>cert>status-verify

Description This command specifies the secondary method used to verify the revocation status of the peer's certificate. The method can be either CRL or OCSP.

To verify the revocation status of the peer's certificate, the CRL or OCSP uses the corresponding configuration in the CA profile of the issuer of the certificate in question.

The secondary method is used only when the primary method fails to provide an answer.

- CRL: CRL expired
- OCSP — unreachable / any answer other than "good" or "revoked" / OCSP is not configured in ca-profile/ OCSP response is not signed / Invalid nextUpdate

Default no secondary

Parameters **crl —** the CRL file is configured in the corresponding CA profile

ocsp — the OCSP server is configured in the corresponding CA profile

trust-anchor-profile

Syntax **trust-anchor-profile** *profile-name*
no trust-anchor-profile

Context config>service>vprn>if>sap>ipsec-tunnel>cert

Description This command configures the **trust-anchor-profile** for the specified IPSec tunnel. This command overrides the **trust-anchor-profile** configured in the **config>ipsec** context.

Default no trust-anchor-profile

Parameters *profile-name* — the name of the **trust-anchor-profile**

Show Commands

Note: The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

ca-profile

Syntax **ca-profile** *name* [**association**]

Context show>certificate

Description This command displays IPSec certificate profile information for root and subordinate CAs.

Parameters *name* — specifies an existing CA profile name, up to 32 characters

association — displays information for which this CA profile is associated

Output The following output is an example of CA profile information.

Sample Output

```
*A:Dut-A# show>certificate# ca-profile "test"
===============================================================================
PKI CA-Profile Information
===============================================================================
CA Profile     : test                           Admin State    : down
Description    : (Not Specified)
CRL File       : (Not Specified)
Cert File      : (Not Specified)
Oper State     : down
Oper Flags     : adminDown
CMPv2
-------------------------------------------------------------------------------
HTTP Timeout   : 30 secs                        Router         : base
CA URL         : (Not Specified)
Sign Cert URL  : (Not Specified)
Unprot Err Msg : disabled                       Unprot Pki Conf: disabled
Same RecipNonce: disabled
for Poll-reqs
Set Sndr for IR: True
HTTP version   : 1.1
OCSP
-------------------------------------------------------------------------------
Responder URL  : (Not Specified)
Router         : base
===============================================================================
*A:Dut-A# show>certificate#
```

ocsp-cache

Syntax **ocsp-cache** [*entry-id*

Context show>certificate

Description This command displays OCSP cache information.

Parameters *entry-id* — specifies the ID of an entry in the OCSP cache, from 1 to 2000

statistics

Syntax **statistics**

Context show>certificate

Description This command displays certificate-related statistics.

Output The following output is an example of certificate-related statistics information.

Sample output

```
*A:Dut-A# show>certificate# statistics
===============================================================================
Certificate Statistics
===============================================================================
Auth Failed          : 0                    Auth Passed      : 4
Total Auth Req       : 4
===============================================================================
*A:Dut-A# show>certificate#
```

trust-anchor-profile

Syntax **trust-anchor-profile** *trust-anchor-profile* **association**
trust-anchor-profile [*trust-anchor-profile*]

Context show>ipsec

Description This command displays trust anchor profile information. Specifying a trust anchor profile shows the CA certificates associated with that trust anchor profile. When a trust anchor profile is not specified, the command shows all trust anchor profiles configured on the system and the number of CAs that are down in each profile. When a trust anchor profile is specified along with the association **keyword**, the command displays the names of the IPSec tunnels that are using a particular trust anchor profile.

Parameters *trust-anchor-profile* — specifies a trust anchor profile name, up to 32 characters

Output The following output is an example of trust anchor profile information.

Sample Output

```
*A:7705:Dut-A# show>ipsec# trust-anchor-profile trustAnchorProfile_11
===============================================================================
Trust Anchor CA-Profile List
===============================================================================

CA Profile                      Admin/Oper State
-------------------------------------------------------------------
caProfile_11                    down/down
===================================================================
A:7705:Dut-A# show>ipsec>trust-anchor-profile#

A:7705:Dut-A# show>ipsec# trust-anchor-profile
===================================================================
Trust Anchor Profile Information
===================================================================
Name                            CA Profiles Down
-------------------------------------------------------------------
trustAnchorProfile_1            0
trustAnchorProfile_11           0
===================================================================
*A:7705:Dut-A# show>ipsec#

*A:7705:Dut-A# show>ipsec# trust-anchor-profile "trustAnchorProfile_1" association
===============================================================================
IPsec tunnels using trust-anchor-profile
===============================================================================
SvcId      Type   SAP                          Tunnel
-------------------------------------------------------------------------------
2          vprn   tunnel-1.private:1           tunnelPrivateSide_1
===============================================================================
Number of tunnel entries: 1
===============================================================================
===============================================================================
*A:7705:Dut-A# show>ipsec#
```

cert-profile

Syntax **cert-profile** *name* **association**
cert-profile [*name*]
cert-profile *name* **entry** [**1..8**]

Context show>ipsec

Description This command displays IPSec certificate profile information.

Parameters *name* — specifies an existing certificate profile name

association — displays information for which this IPSec certificate profile is associated

entry [1..8] — displays information for the specified entry

Output The following output is an example of IPSec certificate profile information.

Sample Output

```
*A:Dut-A# show ipsec cert-profile cert "cert-1.der"
===============================================================================
Certificate Profile Entry
===============================================================================
Id Cert                     Key                      Status Flags
-------------------------------------------------------------------------------
1  cert-1.der               key-1.der
===============================================================================
*A:Dut-A#


*A:Dut-A# show ipsec cert-profile "cert-1.der" entry 1
===============================================================================
IPsec Certificate Profile: cert-1.der Entry: 1 Detail
===============================================================================
Cert File        : cert-1.der
Key File         : key-1.der
Status Flags     : (Not Specified)
Comp Chain       : complete

Compute Chain CA Profiles
-------------------------------------------------------------------------------
CA10
CA9
CA8
CA7
CA6
===============================================================================
*A:Dut-A# exit
```

ike-policy

Syntax **ike-policy**
ike-policy *ike-policy-id*

Context show>ipsec

Description This command displays provisioning parameters for a given IKE policy. When an *ike-policy-id* is not specified then a summary display showing all IKE policies is displayed. When an *ike-policy-id* is specified then a detailed display showing IKE policy settings for the specific IKE policy is displayed.

Parameters *ike-policy-id* — specifies the ID of an IKE policy entry

Values 1 to 2048

Output The following output is an example of IPSec security policy information, and Table 144 describes the fields.

Sample Output

```
*A:7705custDoc:Sar18>show>ipsec# ike-policy
===============================================================================
IPsec IKE Policies
===============================================================================
Id   Ike  Ike DH Pfs   Pfs Auth   Encr   Isakmp IPsec  Auth     DPD     NAT
     Mode Ver          DH  Alg    Alg    Life-  Life-  Method
                                         time   time
-------------------------------------------------------------------------------
1    Main  2  2  False 2   Sha1   Aes128 86400  3600   psk      disable disable
2    Main  2  14 True  5   Sha384 Aes192 60000  48000  psk      enable  enable
-------------------------------------------------------------------------------
No. of IPsec IKE Policies: 2
===============================================================================
*A:7705custDoc:Sar18>show>ipsec#

*A:7705custDoc:Sar18>show>ipsec# ike-policy 1
===============================================================================
IPsec IKE policy Configuration Detail
===============================================================================
Policy Id        : 1                    IKE Mode         : main
DH Group         : Group2               Auth Method      : psk
PFS              : False                PFS DH Group     : Group2
Auth Algorithm   : Sha1                 Encr Algorithm   : Aes128
ISAKMP Lifetime  : 86400                IPsec Lifetime   : 3600
NAT Traversal    : Disabled
NAT-T Keep Alive : 0                    Behind NAT Only  : True
DPD              : Disabled
DPD Interval     : 30                   DPD Max Retries  : 3
Description      : (Not Specified)
IKE Version      : 2                    Own Auth Method  : symmetric
```

Table 144: Show IPSec IKE-Policy Command Output Fields

Label	Description
IPsec IKE Policies	
Id	The IKE policy identifier
Ike Mode	The IKE mode
Ike Ver	The IKE version
DH	The Diffie-Hellman group (DH) used for the IKE policy
Pfs	Displays whether or not perfect forward secrecy (PFS) is used on the IPSec tunnel using this policy
Pfs DH	The Diffie-Hellman group (DH) used for calculating PFS keys
Auth Alg	The hashing algorithm used for the IKE authentication function
Encr Alg	The encryption algorithm used for the IKE session

Table 144: Show IPSec IKE-Policy Command Output Fields (Continued)

Label	Description
Isakmp Life-time	The lifetime of a phase 1 IKE key, in seconds
IPsec Life-time	The lifetime of a phase 2 IKE key, in seconds
Auth Method	The authentication method
DPD	The state of the dead peer detection (DPD) mechanism: Enabled or Disabled
NAT	The state of Network Address Translation Traversal (NAT-T)
No. of IPsec IKE Policies:	The number of IPSec IKE policies
IPsec IKE Policy Configuration Detail	
Policy Id	The IKE policy identifier
IKE Mode	The IKE mode
DH Group	The Diffie-Hellman group (DH) used for the IKE policy
Auth Method	The authentication method
PFS	Displays whether or not perfect forward secrecy (PFS) is used on the IPSec tunnel using this policy
PFS DH Group	The Diffie-Hellman group (DH) used for calculating PFS keys
Auth Algorithm	The hashing algorithm used for the IKE authentication function
Encr Algorithm	The encryption algorithm used for the IKE session
ISAKMP Lifetime	The lifetime of a phase 1 IKE key, in seconds
IPsec Lifetime	The lifetime of a phase 2 IKE key, in seconds
NAT Traversal	The state of Network Address Translation Traversal (NAT-T): Enabled, Disabled, or Force
NAT-T Keep Alive	Displays the configured NAT-T keepalive interval, in seconds
Behind NAT Only	Indicates when NAT-T keepalive messages are sent `True`—keepalive messages are sent if a NAT device is detected. Detection is done by each IKE session, for each IPSec tunnel. `False`—keepalive messages are always sent Note that when **force-keep-alive** is specified, the state of Behind NAT Only is False, otherwise it is True.
DPD	The state of the Dead Peer Detection (DPD) mechanism: Enabled or Disabled
DPD Interval	The interval used to test connectivity to the tunnel peer

Table 144: Show IPSec IKE-Policy Command Output Fields (Continued)

Label	Description
DPD Max Retries	The maximum number of retries before the tunnel is removed
Description	A user-configured description of the IKE policy
IKE Version	The IKE version
Own Auth Method	Indicates the authentication method used with this IKE policy to authenticate on the local side of the tunnel

security-policy

Syntax **security-policy service** *service-id* [**security-policy-id** *security-policy-id*]
security-policy

Context show>ipsec

Description This command displays the provisioning parameters for a given security policy.

Parameters *service-id* — specifies the service ID or name of the tunnel delivery service

Values 1 to 2147483690 or *service-name*

security-policy-id — specifies the IPSec security policy entry that this service will use

Values 1 to 8192

Output The following output is an example of IPSec security policy information, and Table 145 describes the fields.

Sample Output

```
*A:7705custDoc:Sar18>show>ipsec# security-policy
===============================================================================
IPsec Security Policies
===============================================================================
ServiceId               SecurityPolicyId           Security Policy Params
                                                   Entry count
-------------------------------------------------------------------------------
20                      1                          2
20                      17                         0
-------------------------------------------------------------------------------
No. of IPsec Security Policies: 2
===============================================================================
```

```
*A:7705custDoc:Sar18>show>ipsec# security-policy 20
===============================================================================
Security Policy Param Entries
===============================================================================
SvcId      Security   Policy     LocalIp             RemoteIp
           PlcyId     ParamsId
-------------------------------------------------------------------------------
20         1          1          any                 any
20         1          2          11.11.11.11/32      10.10.10.10/32
-------------------------------------------------------------------------------
No. of IPsec Security Policy Param Entries: 2
===============================================================================
===============================================================================
Security Policy Param Entries
===============================================================================
SvcId      Security   Policy     LocalIp             RemoteIp
           PlcyId     ParamsId
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of IPsec Security Policy Param Entries: 0
===============================================================================

*A:7705custDoc:Sar18>show>ipsec# security-policy 20 1
===============================================================================
Security Policy Param Entries
===============================================================================
SvcId      Security   Policy     LocalIp             RemoteIp
           PlcyId     ParamsId
-------------------------------------------------------------------------------
20         1          1          any                 any
20         1          2          11.11.11.11/32      10.10.10.10/32
-------------------------------------------------------------------------------
No. of IPsec Security Policy Param Entries: 2
===============================================================================
*A:7705custDoc:Sar18>show>ipsec#
```

Table 145: Show IPSec Security Policy Command Output Fields

Label	Description
IPsec Security Policies	
ServiceId	The service identifier
SecurityPolicyId	The security policy identifier applied to the service
Security Policy Params Entry count	The number of entries in the security policy
No. of IPsec Security Policies:	The number of IPSec security policies on the router
Security Policy Param Entries	
SvcId	The service identifier

Table 145: Show IPSec Security Policy Command Output Fields (Continued)

Label	Description
Security PlcyId	The security policy identifier applied to the service
Policy ParamsId	The parameter entry number for the security policy
LocalIp	The IP address of the local IP interface
RemoteIp	The IP address of the remote IP interface
No. of IPsec Security Policy Param Entries:	The number of parameter entries for the IPSec security policy

transform

Syntax **transform** [*transform-id*]

Context show>ipsec

Description This command displays IPSec transforms.

Parameters *transform-id* — specifies an IPSec transform entry

Values 1 to 2048

Output The following output is an example of IPSec transform information, and Table 146 describes the fields.

Sample Output

```
*A:7705custDoc:Sar18>show>ipsec# transform
===============================================================
IPsec Transforms
===============================================================
TransformId    EspAuthAlgorithm    EspEncryptionAlgorithm
---------------------------------------------------------------
1              Sha1                Aes128
2              Md5                 3Des
---------------------------------------------------------------
No. of IPsec Transforms: 2
===============================================================
*A:7705custDoc:Sar18>show>ipsec#
```

Table 146: Show IPSec Transform Command Output Fields

Label	Description
IPsec Transforms	
TransformId	The identifier of the IPSec transform policy

Table 146: Show IPSec Transform Command Output Fields (Continued)

Label	Description
EspAuthAlgorithm	Displays the type of Encapsulating Security Payload (ESP) authorization algorithm defined in the transform policy
EspEncryptionAlgorithm	Displays the type of Encapsulating Security Payload (ESP) encryption algorithm defined in the transform policy
No. of IPsec Transforms:	The number of IPSec transform policies

tunnel

Syntax **tunnel**
tunnel *ipsec-tunnel-name*
tunnel count

Context show>ipsec

Description This command displays the IPSec tunnel information for existing tunnels.

Parameters *ipsec-tunnel-name* — specifies the configured name of the IPSec tunnel to be displayed, 32 characters maximum

count — displays the total number of IPSec tunnels

Output The following output is an example of IPSec tunnel information, and Table 147 describes the fields.

Sample Output

```
*A:7705custDoc:Sar18>show>ipsec# tunnel
===============================================================================
IPsec Tunnels
===============================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
-------------------------------------------------------------------------------
vprn_ipsec_tunnel                0.0.0.0           20           Down   Manual
  tunnel-1.private:1               0.0.0.0           None         Down   None
-------------------------------------------------------------------------------
IPsec Tunnels: 1
===============================================================================
*A:7705custDoc:Sar18>show>ipsec#

*A:7705custDoc:Sar18>show>ipsec# tunnel vprn_ipsec_tunnel
===============================================================================
IPsec Tunnel Configuration Detail
===============================================================================
Service Id       : 20                   Sap Id           : tunnel-1.private:1
Tunnel Name      : vprn_ipsec_tunnel
Description      : None
```

```
Local Address    : 0.0.0.0
Remote Address   : 0.0.0.0
Delivery Service : None                 Security Policy  : None
Admin State      : Down                 Oper State       : Down
Last Oper Change : 05/29/2015 15:10:01
Keying Type      : Manual               Replay Window    : None
Clear DF Bit     : false                IP MTU           : max
Copy DF Bit      : false                I
Oper Flags       : unresolvedLocalIp tunnelAdminDown sapDown
                   unresolvedPublicSvc
-------------------------------------------------------------------------------
BFD Interface
-------------------------------------------------------------------------------
BFD Designate    : no
===============================================================================
*A:7705custDoc:Sar18>show>ipsec#

*A:7705custDoc:Sar18>show>ipsec# tunnel count
===============================================================================
IPsec Tunnel Count
===============================================================================
Total IPsec Tunnels                  : 1
===============================================================================
*A:7705custDoc:Sar18>show>ipsec#

*A:7705custDoc:Sar18>show>ipsec# tunnel ipsec_tunnel_tag1
===============================================================================
IPsec Tunnel Configuration Detail
===============================================================================
Service Id       : 20                   Sap Id           : tunnel-1.private:1
Tunnel Name      : ipsec_tunnel_tag1
Description      : None
Local Address    : 10.10.10.1
Remote Address   : 11.11.11.11
Delivery Service : 10                   Security Policy  : 1
Admin State      : Down                 Oper State       : Down
Last_Oper_Change : 05/29/2015 15:10:01
Keying Type      : Dynamic              Replay Window    : None
TrustAnchor Prof : certChainTrustAnchorProfile
Match TrustAnchor: CA.Level6
Cert Profile     : certChainProfile
Local Id Type    : none
Clear DF Bit     : false                IP MTU           : max
Copy DF Bit      : false
Oper Flags       : unresolvedLocalIp tunnelAdminDown sapDown
                   unresolvedPublicSvc
-------------------------------------------------------------------------------
BFD Interface
-------------------------------------------------------------------------------
BFD Designate    : no
-------------------------------------------------------------------------------
Dynamic Keying Parameters
-------------------------------------------------------------------------------
Transform Id1    : 1                    Transform Id2    : 2
Transform Id3    : None                 Transform Id4    : None
Ike Policy Id    : 1                    Auto Establish   : disabled
PreShared Key:12345abc!def%67890
Selected Cert    : depth6.cer
```

```
Selected Key     : depth6.key
Send Chain Prof  : CA.Level0
                 : CA.Level1
                 : CA.Level2
                 : CA.Level3
                 : CA.Level4
                 : CA.Level5
                 : CA.Level6
Certificate Status Verify
-------------------------------------------------------------------------------
Primary          : crl                  Secondary        : none
Default Result   : revoked
-------------------------------------------------------------------------------
ISAKMP-SA
-------------------------------------------------------------------------------
State            : Up
Established      : 12/02/2015 20:01:54  Lifetime         : 86400
Expires          : 12/03/2015 20:01:54
ISAKMP Statistics
--------------------
Tx Packets       : 2                    Rx Packets       : 2
Tx Errors        : 0                    Rx Errors        : 0
Tx DPD           : 0                    Rx DPD           : 0
Tx DPD ACK       : 0                    Rx DPD ACK       : 0
DPD Timeouts     : 0                    Rx DPD Errors    : 0
===============================================================================
===============================================================================
*A:7705custDoc:Sar18>show>ipsec#
```

Table 147: Show IPSec Tunnel Command Output Fields

Label	Description
IPsec Tunnels	
TunnelName	The specified name of the IPSec tunnel
LocalAddress	The IPv4 address of the local router
SvcId	The service identifier
Admn	The administrative state of the IPSec tunnel
Keying	The type of security keying for the tunnel: None, Manual, or Dynamic
SapId	The SAP identifier
RemoteAddress	The IPv4 address of the remote router
DlvrySvcId	The service identifier of the delivery service
Oper	The operational state of the IPSec tunnel
Sec Plcy	The identifier of the security policy used
IPsec Tunnels:	The number of IPSec tunnels

Table 147: Show IPSec Tunnel Command Output Fields (Continued)

Label	Description
IPsec Tunnel Configuration Detail	
Service Id	The service identifier
Sap Id	The SAP identifier
Tunnel Name	The specified name of the IPSec tunnel
Description	The description configured for the IPSec tunnel
Local Address	The IPv4 address of the local router
Remote Address	The IPv4 address of the remote router
Delivery Service	The service identifier of the delivery service
Security Policy	The identifier of the security policy used
Admin State	The administrative state of the IPSec tunnel
Oper State	The operational state of the IPSec tunnel
Last Oper Change	The timestamp indicating the last operational status change for the IPSec tunnel
Keying Type	The type of security keying for the tunnel: None, Manual, or Dynamic
Replay Window	The size of the replay window used for anti-replay
TrustAnchor Prof	The trust anchor profile that is being used
Match TrustAnchor	The actual CA certificate that has been selected from the trust anchor profile
Cert Profile	The certification profile
Clear DF Bit	Indicates whether the tunnel is clearing the DF bit: true (clearing) or false (not clearing)
Copy DF Bit	Indicates whether the tunnel is copying the DF bit: true (copying) or false (not copying)
IP MTU	The interface IP MTU. The value “max” indicates that the tunnel will receive whatever IP payload is sent to it.
Oper Flags	Displays the operational flags currently in effect
BFD Interface	
BFD Designate	Displays whether a BFD designate has been specified: yes or no
Dynamic Keying Parameters	

Table 147: Show IPSec Tunnel Command Output Fields (Continued)

Label	Description
Transform Id1 Transform Id2 Transform Id3 Transform Id4	The ipsec-transform IDs that are assigned under the VPRN **ipsec-tunnel** context
Ike Policy Id	The IKE policy ID
Auto Establish	Displays whether automatic establishing of an IPSec tunnel has been specified: yes or no
PreShared Key	The PSK or shared secret used with dynamic keying as defined under the VPRN **ipsec-tunnel** context
Selected Cert	The actual certificate being used, selected from the **cert-profile**
Selected Key	The actual key being used, selected from the **cert-profile**
Send Chain Prof	The send chain, if configured, under the **cert-profile**
Certificate Status Verify	
Primary	The primary method used to verify the revocation status of the peer's certificate, either CRL or OCSP
Secondary	The secondary method used to verify the revocation status of the peer's certificate, either CRL or OCSP
Default Result	The default result when both the primary and secondary methods fail to verify the revocation status of the peer's certificate, either good or revoked
Isakmp State	The state of ISAKMP: Up or Down
ISAKMP Statistics	ISAKMP statistics are for traffic sent and received by the IKE protocol
Tx Packets	The number of IKE packets transmitted
Rx Packets	The number of IKE packets received
Tx Errors	The number of IKE packet errors transmitted
Rx Errors	The number of IKE packet errors received
Tx DPD	The number of IKE Dead Peer Detection (DPD) packets transmitted
Rx DPD	The number of IKE DPD packets received
Tx DPD ACK	The number of IKE DPD acknowledged packets transmitted
Rx DPD ACK	The number of IKE DPD acknowledged packets received

Table 147: Show IPSec Tunnel Command Output Fields (Continued)

Label	Description
DPD Timeouts	The number of IKE DPD timeouts
Rx DPD Errors	The number of IKE DPD packet errors received
IPsec Tunnel Count	
Total IPsec Tunnels	The total number of IPSec tunnels on the local router

Clear Commands

mda

Syntax **mda** {*slot/mda* | **all**}
mda all statistics
mda *slot/mda* **statistics security** [**encryption**]

Context clear

Description This command clears statistics.

Parameters *slot/mda* **—** the port or module identifier

all — resets all ports or modules on the node

all statistics — clears all security statistics on the node

encryption — specifies the security type

statistics security — clears only security statistics for the specified port or module

Debug Commands

cmpv2

Syntax [**no**] **cmpv2**

Context debug

Description This command enables the context to perform CMPv2 debug operations.

ca-profile

Syntax [**no**] **ca-profile** *profile-name*

Context debug>cmpv2

Description This command debugs the output from the specified CA profile.

- The protection method of each message is logged.
- All HTTP messages are logged. The format allows offline analysis using Wireshark.
- In the event of failed transactions, saved certificates are not deleted from the file system in order to allow for further debug and analysis.
- The system allows CMPv2 debugging for multiple CA profiles at the same time.

certificate

Syntax [**no**] **certificate** *filename*

Context debug>ipsec

Description This command enables debug for certificate chain computation in cert-profile.

Parameters *filename* — Displays the filename of imported certificate.

tunnel

Syntax **tunnel** [*ipsec-tunnel-name*] [**detail**]
no tunnel [*ipsec-tunnel-name*]

Context debug>ipsec

Description This command can be used to facilitate debugging related to IPSec tunnels.

Parameters *ipsec-tunnel-name* — specifies an IPSec tunnel name up to 32 characters in length

keyword - optional to

detail — enables detailed debug information

Network Group Encryption

In This Chapter

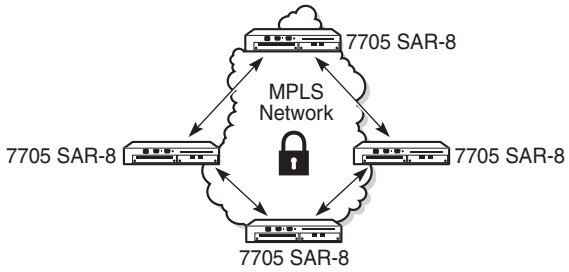
This chapter provides information to configure network group encryption (NGE). Topics in this chapter include:

- NGE Overview
- NGE Global Encryption Label
- Key Groups
- Services Encryption
- QoS for NGE Traffic
- Statistics
- Remote Network Monitoring (RMON) Support
- Configuration Notes
- Configuring NGE with CLI
- NGE Command Reference

NGE Overview

The network group encryption (NGE) feature provides the ability to configure end-to-end encryption of MPLS-based user data traffic without the need to create and manage a mesh of IPSec tunnels to achieve an any-to-any connectivity model, as used in VPRN or VPLS services for example. Figure 111 illustrates any-to-any encryption of MPLS-based services

Figure 111: Any-to-Any Encryption of MPLS-based Services



The 7705 SAR provides two main types of encryption to secure MPLS services in an IP/MPLS network:

- SDP encryption for a variety of services (VLL, VPLS, VPRN)
- encryption for MP-BGP-based VPRN services

The NGE solution on the 7705 SAR is fully redundant. When an activity switch occurs, there is no traffic loss for any encrypted services. Operators that enable NGE encryption can choose the specific MPLS services they would like encrypted depending on their security requirements and needs per service.

NGE is supported on the following adapter cards and platforms:

- 7705 SAR-8/7705 SAR-18: 8-port Gigabit Ethernet Adapter card
- 7705 SAR-8/7705 SAR-18: 6-port Ethernet 10Gbps Adapter card
- 7705 SAR-18: 10-port 1GigE/1-port 10GigE X-Adapter card
- 7705 SAR-H (NGE is also supported over T1/E1 interfaces using the 4-port T1/E1 and RS-232 Combination module)
- 7705 SAR-Hc
- 7705 SAR-W
- 7705 SAR-Wx
- 7705 SAR-X

This section contains information on the following topics:

- NGE Security Domains and Encryption Partitions
- 5620 SAM Management

NGE Security Domains and Encryption Partitions

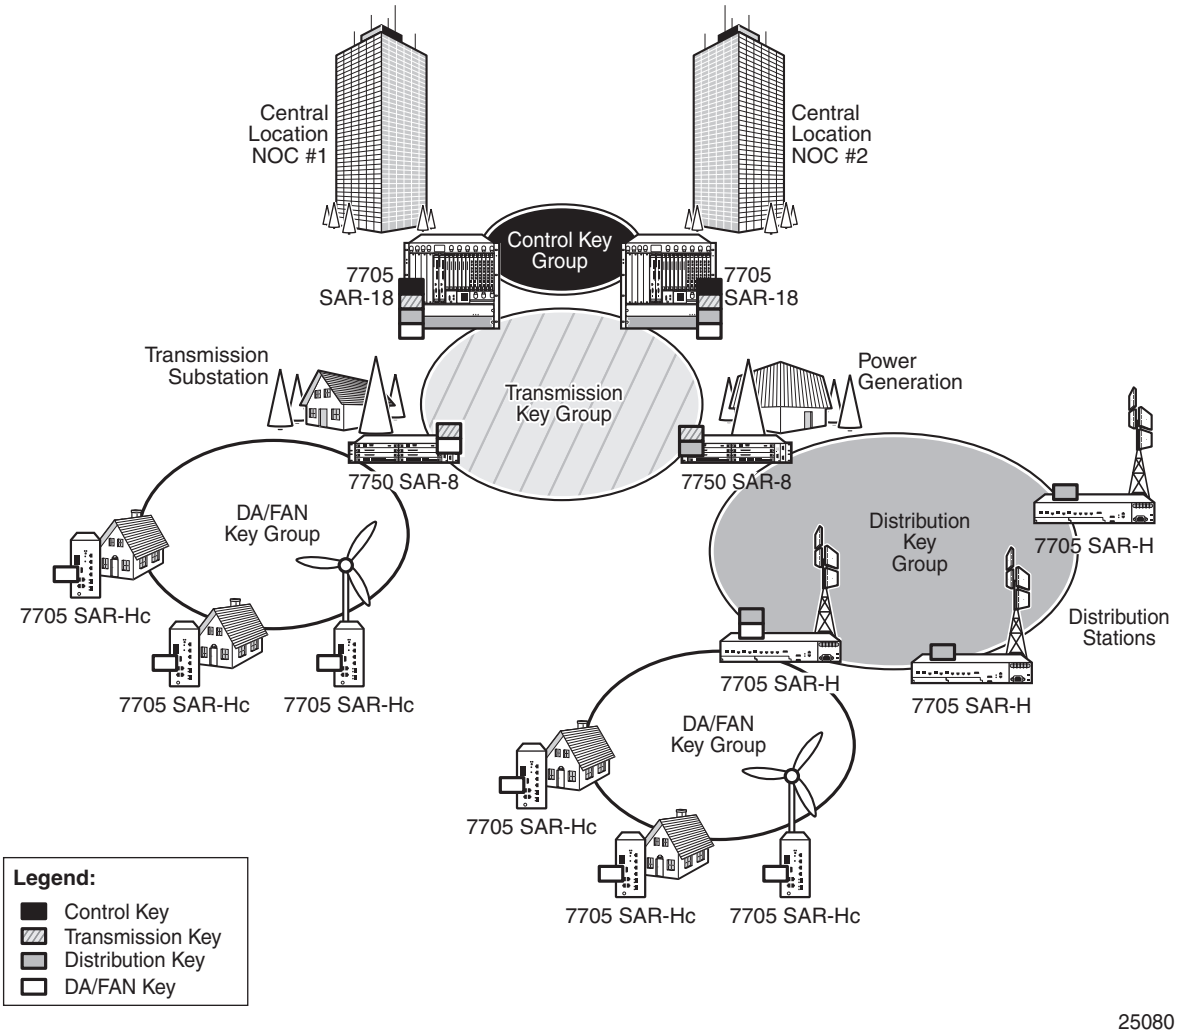
The NGE feature allows a tiered approach to managing encryption keys in a network using key groups. This is accomplished by the configuration of services to use specific key groups depending on security policies for the service and network topology.

Figure 112 illustrates a typical application of NGE key-group partitioning.

In the Figure 112 scenario, there are several critical levels (tiers) of security that need to be considered. In this example, the protection of Distribution Automation and Field Area Network (DA/FAN) equipment might be considered less critical than the Transmission or Distribution Substation network equipment. It is ideal to ensure that nodes more at risk of a security breach do not contain more critical information than is necessary. To this end, encryption keys for the sensitive portions of the network should not be available on nodes that are at risk. The 7705 SAR NGE feature enables operators to partition and distribute encryption keys between different domains or groups in a network. NGE partitions are enabled by configuring different key groups per security partition.

For example, assume an attacker attempts to gain access to the network from a DA/FAN location where the physical installation of equipment in remote locations and a large number of devices could make those locations more prone to attack, and NGE has partitioned this part of the network. Because those devices do not contain any highly sensitive key-group information used for other parts of the network, the attacker is limited in the potential scope of any attack.

Figure 112: Key-Group Partitioning



Another application of key-group partitioning allows different parts of an organization to have a method of end-to-end communication without the need to share keys or key security policies across the entire network. If two domains need to communicate between domains, gateway nodes could be configured with both key groups and gateway functionality between various encryption domains, as needed.

5620 SAM Management

The NGE feature is tightly integrated with the 5620 SAM for key-group management associated with MPLS services that are configured on nodes and for encryption-service provisioning on a network-wide basis. The following functions are provided by the 5620 SAM:

- managing and synchronizing keys within key groups on a network-wide basis
- configuring NGE by downloading relevant key groups to nodes with NGE enabled
- managing network-wide re-keying of key values within key groups
- managing services that are encrypted by ensuring that only the nodes that need specific key groups have those key groups downloaded

The 5620 SAM acts as the key server for the nodes and provides the keys in key groups that are used to perform encryption and authentication. The 5620 SAM ensures that all nodes in a key group are kept in synchronization and that only the key groups that are relevant to a particular node are downloaded with sensitive keys.

The 5620 SAM also provides the function of re-keying a key group with zero outage time during the re-key procedure. Different key groups can be re-keyed at different times, if desirable, or all key groups can be re-keyed network-wide at the same time.

For more information on 5620 SAM management, refer to the “Service Management” section in the Alcatel-Lucent 5620 SAM User Guide.

NGE Global Encryption Label

The global encryption label must be a unique network-wide label (that is, the same label must be used on all nodes in the network that require NGE services). The label must be configured on individual nodes before NGE can become operational on that node.

The label is used for identifying packets that have an NGE encrypted payload.

Once the global encryption label is set, it should not be changed. If the label must be changed without impacting traffic, then all key groups in the system should first be deleted, next, the label should be changed, and finally, all key groups should be reconfigured.

The 5620 SAM helps coordinate the distribution of the global encryption label and ensures that all nodes in the network are using the same global encryption label.

Key Groups

Key groups are used to organize encryption keys into distinct groups that allow a user to partition the network based on security requirements. A key group contains the following elements:

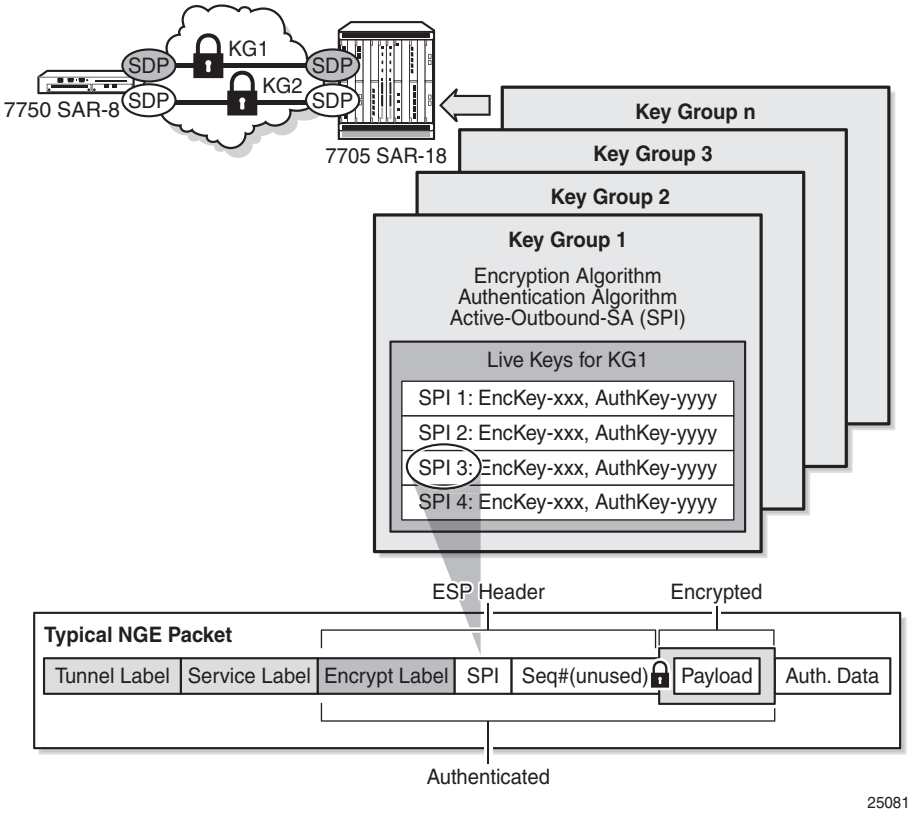
- an encryption algorithm—see Key Group Algorithms
- an authentication algorithm—see Key Group Algorithms
- a list of security associations (SAs)—see Security Associations (SAs)
- an active outbound SA—see Active Outbound SA

Figure 113 illustrates the use of key groups (KGs), security associations (SAs), and security parameter indices (SPIs). The 7705 SAR-8 Shelf V2 and 7705 SAR-18 both have the same set of key groups configured. One path uses key group 1 (KG1) and the other uses key group 2 (KG2). Each key group contains the elements listed above. Key group 1 has four live keys, SPI_1 through SPI_4, and SPI_3 is the active outbound SA. The active outbound SA is identified by its SPI, and this SPI is embedded in the NGE packet.

Each SA listed in a key group, indexed by an SPI, specifies a single key for encryption and a single key for authentication. Packets transmitted or received that reference a particular SPI use the keys in the SA for that SPI when performing encryption and authentication.

Before enabling encryption, key groups must be configured on the node. Only after a key group is configured can it be assigned to an SDP or VPRN services.

Figure 113: Key Groups and a Typical NGE Packet



Key Group Algorithms

All SAs configured in a key group share the same encryption algorithm and the same authentication algorithm. The size and values required by a particular key depend on the requirements of the algorithms selected (see lists below). One encryption algorithm and one authentication algorithm must be selected per key group.

Encryption algorithms available per key group include:

- AES128 (a 128-bit key, requiring a 32-digit ASCII hexadecimal string)
- AES256 (a 256-bit key, requiring a 64-digit ASCII hexadecimal string)

Authentication algorithms available per key group include:

- HMAC-SHA-256 (a 256-bit key, requiring a 64-digit ASCII hexadecimal string)
- HMAC-SHA-512 (a 512-bit key, requiring a 128-digit ASCII hexadecimal string)

Encryption and authentication strengths can be mixed depending on the requirements of the application. For example, 256-bit strength encryption can be used with 512-bit strength authentication.

The configured algorithms cannot be changed when there is an existing SA configured for the key group. All SAs in a key group must be deleted before a key group algorithm can be modified.

Key values are not visible in CLI or retrievable using SNMP. Each node calculates a 32-bit CRC checksum for the keys configured against the SPI. The CRC can be displayed in the CLI or read by SNMP. The purpose of the CRC is to provide a tool to check consistency between nodes, thereby verifying that each node is set with the same key values while keeping the actual key values hidden.

Encapsulating Security Payload (ESP)

The NGE feature uses the Encapsulating Security Payload (ESP) protocol according to IETF RFC 4303. ESP maintains data integrity, ensuring privacy and confidentiality for encrypted traffic.

The ESP protocol used by NGE relies on symmetric ciphers, meaning that the same key is used for encryption and decryption. The 7705 SAR supports Cipher Block Chaining (CBC) encryption mode. Block ciphers used by NGE include:

- AES128 with a 128-bit key using 128-bit blocks
- AES256 with a 256-bit key using 128-bit blocks

For authentication, the integrity check value (ICV) size is as follows:

- HMAC-SHA-256 (16 bytes or 128 bits)
- HMAC-SHA-512 (32 bytes or 256 bits)

Security Associations (SAs)

Each key group has a list of up to four SAs. An SA is a reference to a particular pair of encryption and authentication keys that can used to decrypt and authenticate packets received by the node, and to encrypt packets leaving the node.

For encrypted ingress traffic, any of the four SAs in the key group can be used for decryption as long as there is a match between the SPI in the traffic and the SPI in the SA. For egress traffic, only one of the SAs can be used for encryption and is designated as the active outbound SA. Figure 113 illustrates these relationships.

As shown in Figure 113, each SA is referenced by a particular SPI value, which is included in packets during encryption and authentication. SPI values must be numerically unique throughout all SAs in all key groups. If an SPI value is configured in one key group and an attempt is made to configure the same SPI value in another key group, the configuration is blocked.

Note: Keys are entered in clear text using the `security-association` command. Once entered, they are never displayed in their original, clear text form. Keys are displayed in a 7705 SAR OS-encrypted form, which is indicated by the system-appended `crypto` keyword when an `info` command is run. The 7705 SAR OS also includes the `crypto` keyword with an `admin>save` operation so that the 7705 SAR OS can decrypt the keys when reloading a configuration database. For security reasons, keys encrypted on one node are not usable on other nodes (that is, keys are not exchangeable between nodes).

Active Outbound SA

The active outbound SA is specified by the SPI referencing the specific SA used to encrypt and authenticate packets egressing the node for the SDP or service using the key group. The SPI value for the active outbound SA is included in the Encapsulating Security Payload (ESP) header of packets being encrypted and authenticated.

Services Encryption

The NGE feature provides the ability to encrypt MPLS services using key groups that are configured against these services. These services include:

- VLL service (Epipe and Cpipe)
- VPRN service using Layer 3 spoke-SDP termination
- IES service using Layer 3 spoke-SDP termination
- VPLS service using spoke and mesh SDPs
- routed VPLS service into a VPRN or IES
- MP-BGP-based VPRNs

For services that use SDPs, all tunnels may be either MPLS LSPs (RSVP-TE, LDP, or static LSP) or GRE tunnels.

For MP-BGP services, `auto-bind` is supported using LDP, GRE, RSVP-TE, or MPLS (LDP or RSVP-TE).

This section contains information on the following topics:

- Services Encryption Overview
- Assigning Key Groups to Services
- Pseudowire Switching for NGE Traffic
- Pseudowire Control Word for NGE Traffic
- VPRN Layer 3 Spoke-SDP Encryption and MP-BGP-based VPRN Encryption Interaction
- NGE and RFC 3107
- NGE Packet Overhead and MTU Considerations

Services Encryption Overview

NGE adds a global encryption label to the label stack for encrypting MPLS services. The global encryption label is used to identify NGE packets that are encrypted and is added to the bottom of the stack. This allows network elements such as LSRs, ABRs, ASBRs, and RRs to forward NGE packets without needing to understand NGE or to know that the contents of these MPLS packets are encrypted. Only when a destination PE (7705 SAR) receives a packet that needs to be understood at the service layer does the PE check for an encryption label, and then decrypt the packet.

Figure 114 illustrates the NGE MPLS/GRE label stack. Figure 115 illustrates VPRN and PW (with control word) packet formats using NGE encryption.

Figure 114: NGE MPLS/GRE Label Stack

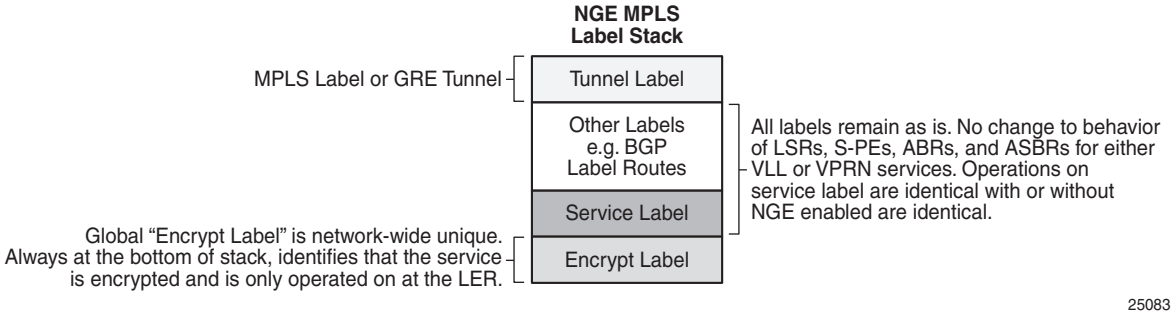
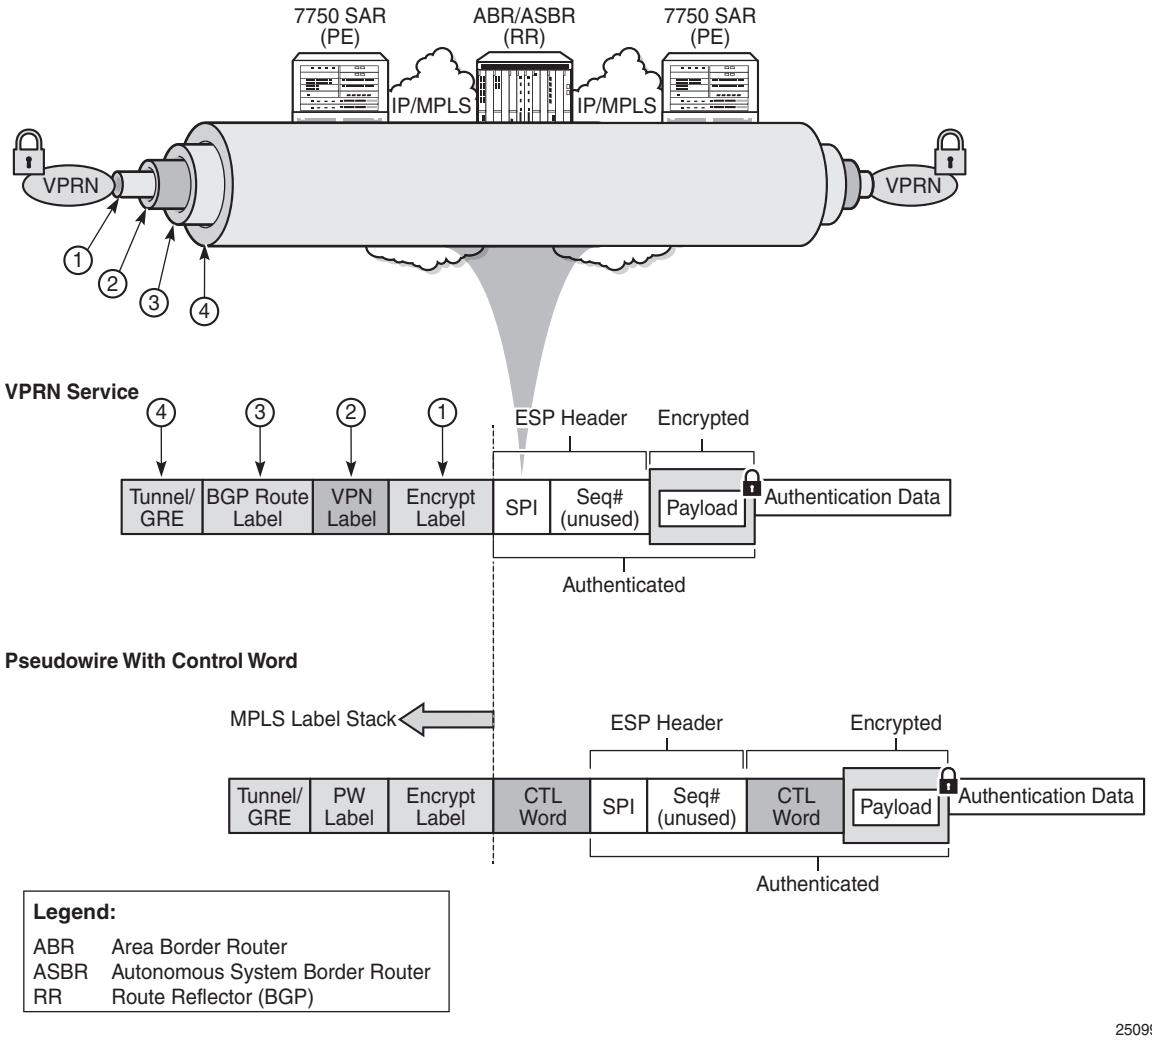


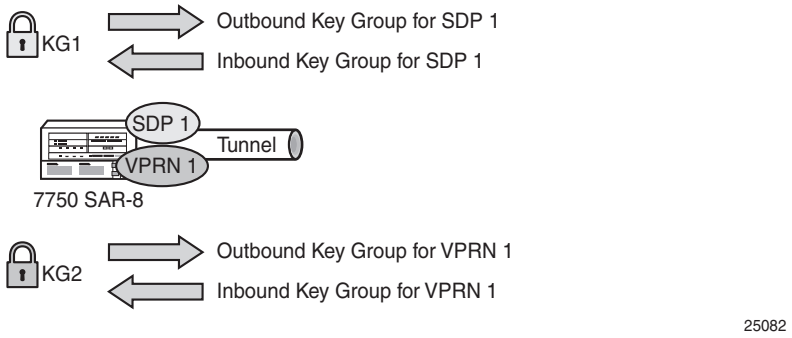
Figure 115: NGE Encryption and Packet Formats



Assigning Key Groups to Services

Assigning key groups to services requires configuring an inbound and outbound key group on a per-service basis, for directional processing (see Figure 116).

Figure 116: Inbound and Outbound Key-Group Assignments



The outbound key group identifies which key group to use for traffic that egresses the node for the service. The inbound key group ensures that ingress traffic is using the correct key group for the service.

If the inbound key group is not set, then the node ensures that packets are either unencrypted or are using one of the valid key groups configured in the system.

In the majority of deployment scenarios, the inbound and outbound key groups are identical. However, it is possible to configure different key groups as the outbound and the inbound key groups, as this is not checked by the node. Using different key groups in this way is not typical.

Including an inbound and outbound direction when assigning key groups to services allows users to:

- easily enable NGE for existing services
- move services from one key-group domain to another domain without halting encryption

The NGE feature makes use of the 5620 SAM to help manage the assignment of key groups to services on a network-wide basis. Refer to the Alcatel-Lucent 5620 SAM User Guide for more information.

Pseudowire Switching for NGE Traffic

For VLL services, the 7705 SAR supports pseudowire (PW) switching of encrypted traffic from one PW to another. There are three scenarios that are supported on the 7705 SAR with regard to PW switching of traffic:

- PW switch using the same key group

  When a PW is using an encrypted SDP, the PW may be switched to another PW that is also using an encrypted SDP, and both SDPs are in the same key group. For this case, to perform the PW switch, the 7705 SAR leaves the encrypted payload unchanged and swaps the labels as needed for passing traffic between PWs.
- PW switch using different key groups

  When a PW is using an encrypted SDP, the PW may be switched to another PW that is also using an encrypted SDP, but both SDPs are in different key groups. For this case, the 7705 SAR decrypts the traffic from the first SDP by using the configured key group for that SDP, and then re-encrypts the traffic by using the egress SDP's key group egress SPI ID.
- PW switch between an encrypted and unencrypted PW

  When traffic is switched from an encrypted PW to an unencrypted PW, the traffic is decrypted before it is sent. The converse occurs in the reverse direction (that is, traffic from an unencrypted PW to an encrypted PW gets encrypted before it is sent).

See Pseudowire Switching for more information.

Pseudowire Control Word for NGE Traffic

The control word is a configurable option for PWs and is included in PW packets when it (the control word) is configured. When `control-word` is enabled and NGE is used, the datapath creates two copies of the CW. One CW is both encrypted and authenticated, and is inserted after the ESP header. The other CW is not encrypted (clear form) and is inserted before the ESP header.

For cases where PW switching is configured, the 7705 SAR ensures—in the CLI and with SNMP—that both segments of the PW have consistent configuration of the control word when encryption is being used.

See Pseudowire Control Word for more information.

VPRN Layer 3 Spoke-SDP Encryption and MP-BGP-based VPRN Encryption Interaction

The encryption configured on an SDP used to terminate the Layer 3 spoke SDP of a VPRN always overrides any VPRN-level configuration for encryption:

- When VPRN encryption is enabled, all routes resolved via MP-BGP are encrypted or decrypted using the VPRN key group.
- When Layer 3 spoke-SDP encryption is enabled, all routes resolved via the Layer 3 interface are encrypted or decrypted using the SDP's key group.

Some examples are as follows.

- If a VPRN is enabled for encryption while a Layer 3 spoke SDP for the same VPRN is using an SDP that is not enabled for encryption, then traffic egressing the spoke SDP is not encrypted.
- If a VPRN is disabled for encryption while a Layer 3 spoke SDP for the same VPRN is using an SDP that is enabled for encryption, then traffic egressing the spoke SDP is encrypted.
- If a VPRN is enabled for encryption using key group X, while a Layer 3 spoke SDP for the same VPRN is using key group Y, then traffic egressing the spoke SDP is encrypted using key group Y.

The commands used for these scenarios are `config>service>sdp>encryption-keygroup` and `config>service>vprn>encryption-keygroup`.

NGE and RFC 3107

When RFC 3107 is enabled on the node and NGE traffic is crossing the RR between two VPRN domains, the same key group must be used between the two domains.



Note: It is the responsibility of the network operator to ensure key group consistency across the RR.

NGE Packet Overhead and MTU Considerations

NGE adds overhead packets to services. Table 148 shows the additional overhead for the worst-case scenario. Additional overhead depends on which encryption and authentication algorithms are chosen.

Table 148: NGE Overhead

Item	Number of Bytes
Encryption label	4
ESP	24
ICV	32
Padding	17
Control word copy	4
Total	**81**

For MP-BGP-based VPRNs, the total is 77 bytes because the control word copy is not required.

The overhead values in Table 148 must be considered for services that are supported by NGE.

Note: Currently, the port MTU has a default value of 1572 bytes. This value is too low for outbound traffic when NGE is enabled. Users must configure new MTU values to adjust for the overhead associated with NGE, as outlined below for MPLS-based and GRE-based services. For details on configuring MTU, refer to the “MTU Configuration Guidelines” section in the 7705 SAR OS Interface Configuration Guide.

The sample calculations in Table 149 show how NGE overhead affects SDP MTU and service MTU values for MPLS-based, GRE-based, and VPRN-based services. The calculations are with and without NGE enabled.

Table 149: Accounting for NGE Overhead SDP and Service MTU — Sample Calculations

Service Type	MTU Values With and Without NGE Enabled
MPLS-based services	**SDP MTU (MPLS)** = 1572 (network port MTU) – 14 (Ethernet header) – 4 (outer label) – 4 (inner label) = 1550 bytes (without NGE enabled) => 1469 bytes (with NGE enabled)
	Service MTU (MPLS) considerations with NGE enabled • Layer 3 spoke IP MTU (MPLS) = 1469 – 14 (inner Ethernet header) = 1455 bytes • PW spoke SDP MTU (MPLS) = SDP MTU = 1469 bytes
GRE-based services	**SDP MTU (GRE)** = 1572 (network port MTU) – 14 (Ethernet header) – 20 (IP header) – 4 (GRE header) – 4 (inner label) = 1530 bytes (without NGE enabled) => 1449 bytes (with NGE enabled)
	Service MTU (GRE) considerations with NGE enabled • Layer 3 Spoke IP MTU (GRE) = 1449 – 14 (inner Ethernet header) = 1435 bytes • PW Spoke MTU (GRE) = SDP MTU = 1449 bytes
VPRN-based services	Each interface inherits its MTU from the SAP or spoke SDP to which it is bound and can be manually changed using the `ip-mtu` command.
	MP-BGP-based VPRN services The MTU of the egress IP interface is used. When NGE is enabled on a VPRN service, customers must account for the additional 77 bytes of overhead needed by NGE for any egress IP interface used by the VPRN service.

QoS for NGE Traffic

The 7705 SAR provides priority and scheduling for traffic into the encryption and decryption engines on nodes that support NGE. This is described for the following traffic directions:

- Network Ingress
- Network Egress

For information on QoS and QoS policies, refer to the 7705 SAR OS Quality of Service Guide. Also, since QoS for NGE is similar to QoS for IPSec, see QoS for IPSec for more information.

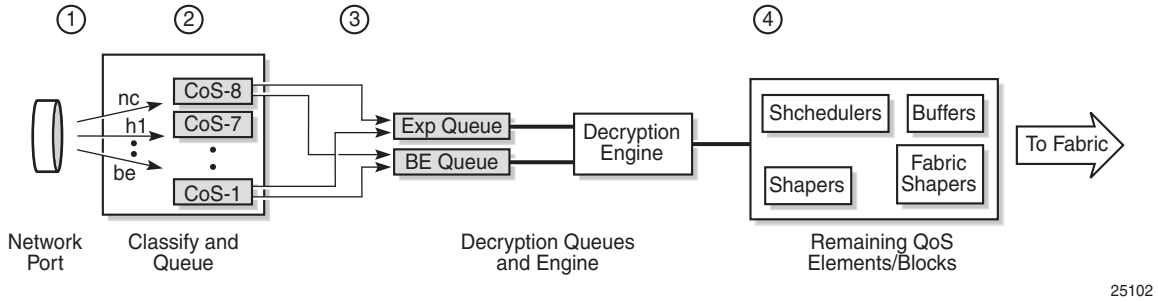
Network Ingress

NGE traffic arriving on network ingress is classified based on the network QoS policy that is assigned to the network router interface. This classification is done using the EXP bits of the outer MPLS tunnel LSP or the DSCP markings on the IP packet of GRE packets.

There are two queues provided for mapping ingress NGE traffic into the decryption engine—an expedited queue and a best-effort queue. The encrypted traffic maps to one of these queues based on the queue-type configuration of the network ingress queue and the FC-to-queue (classification) mapping. Specifically, at network ingress, the following occurs (as shown in Figure 117):

- the network QoS policy determines the forwarding class (FC) for the packet (item 1)
- the ingress network queue QoS policy maps the FC to a queue, where the queue-type has been configured as expedited, best-effort, or auto-expedite (item 2)
- the queue-type is used again to map to the appropriate queue into the decryption engine (that is, the expedited and best-effort queue) (item 3)
- after decryption, normal QoS processing takes place as if NGE was not enabled for the service (item 4).

Figure 117: QoS for NGE Traffic (Network Ingress)



For more information on QoS for network ingress, refer to the "Network Ingress" section in the 7705 SAR OS Quality of Service Guide.

Network Egress

Traffic egressing a network interface typically has a known FC based on the traffic management (TM) configuration at SAP ingress.

There are two queues provided for mapping egress NGE traffic into the encryption engine—an expedited queue and a best-effort queue. The traffic maps to one of these queues based on the FC of the packet, as determined at SAP ingress. Specifically, the following occurs:

- the SAP ingress QoS policy maps the FC to a queue-type, where the queue-type of the SAP-ingress queue has been configured as expedited, best-effort, or auto-expedite
- on the network egress side of the fabric, the ingress queue-type is used to map to the appropriate queue into the encryption engine (that is, the expedited and best-effort queue)

For more information on QoS for network egress, refer to the "Network Egress" section in the 7705 SAR OS Quality of Service Guide.

Statistics

Statistics specific to NGE are counted for the following main areas:

- key group
- SPI
- MDA
- service

Remote Network Monitoring (RMON) Support

Remote network Monitoring (RMON) can be used in conjunction with NGE statistics to provide event and alarm reporting. This can be used by customers to detect security breaches of NGE traffic flows and provide real-time reporting of such events.

Threshold crossing alerts and alarms using RMON are supported for SNMP MIB objects, including NGE.

Configuration Notes

This section describes NGE configuration caveats.

- To enable NGE for an SDP or VPRN service, follow these steps:
  i. install the outbound direction key group on each node for the service
  ii. install the inbound direction key group on each node for the service
- To change NGE from one key group to another key group for an SDP or VPRN service, follow these steps in order:
  i. remove the inbound direction key group from each node for the service
  ii. change the outbound direction key group on each node for the service
  iii. install the new inbound direction key group on each node for the service
- To disable NGE for an SDP or VPRN service, follow these steps in order:
  i. remove the inbound direction key group from each node providing the service
  ii. remove the outbound direction key group from each node for the service
- The authentication and encapsulation keys must contain the exact number of hexadecimal characters required by the algorithm used. For example, using sha256 requires 64 hexadecimal characters.
- The key group bound to an SDP or service must be unbound from that SDP or service before the active outgoing SA for the key group can be removed.
- The active outgoing SA must be removed (deconfigured) before the SPI can be deleted from the SA list in the key group.
- The encryption or authentication algorithm for a key group cannot be changed if there are any SAs in the key group.
- The encryption configured on an SDP used to terminate the Layer 3 spoke SDP of a VPRN (enabled or disabled) always overrides any VPRN-level configuration for encryption. See VPRN Layer 3 Spoke-SDP Encryption and MP-BGP-based VPRN Encryption Interaction for more information.
- The 5620 SAM provides configuration parameters that are not configurable using the CLI. See 5620 SAM Management for more information.

Reference Sources

For information on supported IETF drafts and standards as well as standard and proprietary MIBS, refer to Standards and Protocol Support.

Configuring NGE with CLI

This section provides information about configuring NGE using the command line interface.

Topics in this chapter include:

- Basic NGE Configuration Overview
- Configuring NGE Components
- Service Management Tasks

Basic NGE Configuration Overview

Use the following steps to configure NGE for a particular service. The steps must be performed in order.

1. Configure the group encryption label. The label must be unique, and the same label must be used on all nodes in the network group.
2. Create a key group, duplicating this configuration on all nodes participating in this key group.
   i. Configure the encryption and authentication algorithms for the group.
   ii. Configure a security association (SA) that contains the encryption and authentication keys.
   iii. Configure the active outbound SA for the group.
3. Select the SDPs or VPRN services that require encryption.
   i. For each SDP or VPRN service, configure the outbound direction key group.
   ii. For each SDP or VPRN service, configure the inbound direction key group.

Configuring NGE Components

Use the CLI syntax below to configure the following NGE parameters:

- Configuring the Global Encryption Label
- Configuring a Key Group
- Assigning a Key Group to an SDP or Service

Configuring the Global Encryption Label

The global encryption label is the network-wide, unique MPLS encryption label used for all nodes in the network group. The same encryption label must be configured on each node in the group.

Use the following CLI syntax to configure the global encryption label:

CLI Syntax:
```
config>group-encryption
     group-encryption-label encryption-label
```

The following example displays global encryption label usage:

Example:
```
config# group-encryption
config>grp-encryp# group-encryption-label 34
```

The following example displays the global encryption label configuration:

```
ALU-1>config>grp-encryp# info
----------------------------------------------------
     group-encryption-label 34
----------------------------------------------------
ALU-1>config>grp-encryp#
```

Configuring a Key Group

To configure a key group, set the following parameters:

- encryption and authentication algorithms
- security association
- active outbound SA

The authentication and encapsulation keys must contain the exact number of hexadecimal characters required by the algorithm used. For example, using sha256 requires 64 hexadecimal characters.

Keys are entered in clear text using the `security-association` command. Once entered, they are never displayed in their original, clear text form. Keys are displayed in a 7705 SAR OS-encrypted form, which is indicated by the system-appended `crypto` keyword when an `info` command is run (see the CLI Syntax, Example, and sample output below). The 7705 SAR OS also includes the `crypto` keyword with an `admin>save` operation so that the 7705 SAR OS can decrypt the keys when reloading a configuration database. For security reasons, keys encrypted on one node are not usable on other nodes (that is, keys are not exchangeable between nodes).

Use the following CLI syntax to configure key group options:

CLI Syntax:

```
config# group-encryption
     encryption-keygroup keygroup-id [create]
          description description-string
          esp-auth-algorithm {sha256|sha512}
          esp-encryption-algorithm {aes128|aes256}
          keygroup-name keygroup-name
          security-association spi spi authentication-
            key authentication-key encryption-key
            encryption-key [crypto]
          active-outbound-sa spi
```

The following example displays key group command usage:

Example:

```
config>grp-encryp# encryption-keygroup KG1_secure
config>grp-encryp>encryp-keygrp# description
 Main_secure_KG
config>grp-encryp>encryp-keygrp# esp-auth-algorithm
 sha256
config>grp-encryp>encryp-keygrp# esp-encryption-
 algorithm aes128
config>grp-encryp>encryp-keygrp# keygroup-name
 KG1_secure
```

```
config>grp-encryp>encryp-keygrp# security-association
 spi 2 authentication-key
 0x88433A6DB4FA4F8A490EF661CBE69F010BFAE9C2784BED7059E5
 ADAAB1A225C6 encryption-key
 0x63DCDD501B66F85441E4A55B597DA617
config>grp-encryp>encryp-keygrp# security-association
 spi 6 authentication-key
 0x88433A6DB4FA4F8A490EF661CBE69F010BFAE9C2784BED7059E5
 ADAAB1A225C5 encryption-key
 0x63DCDD501B66F85441E4A55B597DA616
config>grp-encryp>encryp-keygrp# active-outbound-sa 6 ]
```

The following example displays the key group configuration:

```
ALU-1>config>grp-encryp# info detail
----------------------------------------------
        group-encryption-label 34
        encryption-keygroup 2 create
            description "Main_secure_KG"
            keygroup-name "KG1_secure"
            esp-auth-algorithm sha256
            esp-encryption-algorithm aes128
            security-association spi 2 authentication-
key 0x78d9e66a6669bd17454fe3184 ee161315b67adb8912949ceda20b6b741eb63604abe17de478e2
4723a7d1d5f7b6ffafc encryption-
key 0x8d51db8f826239f672457442cecc73665f52cbe00aedfb4eda6166001247b4eb crypto
            security-association spi 6 authentication-key 0x7fb9fc5553630924ee29973f
7b0a48f801b0ae1cb38b7666045274476a268e8d694ab6aa7ea050b7a43cdf8d80977625 encryption-
key 0x72bd9b87841dbebcb2d114031367ab5d9153a41b7c79c8f889ac56b950d8fffa crypto
            active-outbound-sa 6
        exit
----------------------------------------------
ALU-1>config>grp-encryp#
```

Assigning a Key Group to an SDP or Service

A key group can be assigned to the following entities:

- VLL services (Epipe and Cpipe)
- VPRN service using Layer 3 spoke-SDP termination
- IES service using Layer 3 spoke-SDP termination
- VPLS service using spoke and mesh SDPs
- routed VPLS service into a VPRN or IES
- MP-BGP-based VPRNs

For services that use SDPs, all tunnels may be either MPLS LSPs (RSVP-TE, LDP, or static LSP) or GRE tunnels.

For MP-BGP services, `auto-bind` is supported using LDP, GRE, RSVP-TE, or MPLS (LDP or RSVP-TE).

Use the following CLI syntax to assign a key group to an SDP or a VPRN service:

CLI Syntax:
```
config>service# sdp sdp-id [create]
     encryption-keygroup keygroup-id direction
        {inbound|outbound}
```

CLI Syntax:
```
config>service# vprn service-id
     encryption-keygroup keygroup-id direction
        {inbound|outbound}
```

The following examples display a key group assigned to an SDP or a VPRN service:

Example:
```
config>service# sdp 61 create
config>service>sdp# encryption-keygroup 4 direction
 inbound
config>service>sdp# encryption-keygroup 4 direction
 outbound
```

Example:
```
config>service# vprn 22
config>service>vprn# encryption-keygroup 2 direction
 inbound
config>service>vprn# encryption-keygroup 2 direction
 outbound
```

The following example displays key group configuration for an SDP or a VPRN service.

```
ALU-1:Sar18>config>service# info
----------------------------------------------
...
        sdp 61 create
            shutdown
            far-end 10.10.10.10
            exit
            encryption-keygroup 4 direction inbound
            encryption-keygroup 4 direction outbound
        exit
...
        vprn 22 customer 1 create
            shutdown
            encryption-keygroup 2 direction inbound
            encryption-keygroup 2 direction outbound
        exit
...
----------------------------------------------
```

Service Management Tasks

This section discusses the following service management tasks:

- Modifying a Key Group
- Removing a Key Group from an SDP or VPRN Service
- Changing the Key Group for an SDP or VPRN Service
- Deleting a Key Group from a 7705 SAR

Modifying a Key Group

When modifying a key group, observe the following conditions.

- The encryption or authentication algorithm for a key group cannot be changed if there are any SAs in the key group.
- The active outgoing SA must be removed (deconfigured) before the SPI can be deleted from the SA list in the key group.
- Before the outgoing SA can be deconfigured, the key group must be removed from all services on the node that use the key group

In the following example, the active outgoing SA is deconfigured, the SAs are removed, and the encryption algorithm is changed. Then the SAs are reconfigured, followed by reconfiguration of the active outgoing SA. The output display shows the new configuration based on those shown in Configuring a Key Group.

Use the following CLI syntax to modify a key group. The first syntax deconfigures the key-group items and the second syntax reconfigures them.

CLI Syntax:

```
config# group-encryption
     encryption-keygroup keygroup-id
          no active-outbound-sa
          no security-association spi spi
     exit
```

CLI Syntax:

```
config# group-encryption
     encryption-keygroup keygroup-id
          security-association spi spi authentication-
            key auth-key encryption-key encrypt-key
          esp-encryption-algorithm {aes128|aes256}
     exit
```

Example:
```
config>grp-encryp# encryption-keygroup KG1_secure
config>grp-encryp>encryp-keygrp# no active-outbound-sa
config>grp-encryp>encryp-keygrp# no security-association
 spi 2
config>grp-encryp>encryp-keygrp# no security-association
 spi 6
```

Example:
```
config>grp-encryp# encryption-keygroup KG1_secure
config>grp-encryp>encryp-keygrp# esp-encryption-
 algorithm aes256
config>grp-encryp>encryp-keygrp# security-association
 spi 2 authentication-key
 0x01234567890123456789012345678901234567890123456789012
 34567890123 encryption-key
 0x01234567890123456789012345678901234567890123456789012
 34567890123
config>grp-encryp>encryp-keygrp# security-association
 spi 6 authentication-key
 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123
 456789ABCDEF encryption-key
 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123
 456789ABCDEF [crypto]
config>grp-encryp>encryp-keygrp# active-outbound-sa 2
```

The following example displays the commands used to modify a key group. The first example deconfigures the key-group items and the second example reconfigures them. The encryption algorithm is changed from 128 to 256, the keys are changed, and the active outbound SA is changed to SPI 2.

```
ALU-1>config>grp-encryp# info detail
----------------------------------------------
        group-encryption-label 34
        encryption-keygroup 2 create
            description "Main_secure_KG"
            keygroup-name "KG1_secure"
            esp-auth-algorithm sha256
            esp-encryption-algorithm aes128
            no security-association spi 2
            no security-association spi 6
            no active-outbound-sa
        exit
----------------------------------------------
ALU-1>config>grp-encryp#
```

```
ALU-1>config>grp-encryp# info detail
----------------------------------------------
        group-encryption-label 34
        encryption-keygroup 2 create
            description "Main_secure_KG"
            keygroup-name "KG1_secure"
            esp-auth-algorithm sha256
            esp-encryption-algorithm aes256
            security-association spi 2 authentication-
key 0x01234567890123456789012345678901234567890123456789012345678901234567890123 encryption-
key 0x01234567890123456789012345678901234567890123456789012345678901234567890123
            security-association spi 6 authentication-
key 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF encryption-
key 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF crypto
            active-outbound-sa 2
        exit
----------------------------------------------
ALU-1>config>grp-encryp#
```

Removing a Key Group from an SDP or VPRN Service

Both inbound and outbound direction key groups must be deconfigured before the key group can be removed (unbound). The inbound and outbound key groups must be deconfigured individually. Including *keygroup-id* is optional.

Use the following CLI syntax to remove a key group from an SDP or a VPRN service.

CLI Syntax:
```
config>service# sdp sdp-id
     no encryption-keygroup keygroup-id direction
        {inbound|outbound}
```

CLI Syntax:
```
config>service# vprn service-id
     no encryption-keygroup keygroup-id direction
        {inbound|outbound}
```

The following examples display a key group removed from an SDP or a VPRN service:

Example:
```
config>service# sdp 61
config>service>sdp# no encryption-keygroup 4 direction
 inbound
config>service>sdp# no encryption-keygroup 4 direction
 outbound
```

Example:
```
config>service# vprn 22
config>service>vprn# no encryption-keygroup 2 direction
 inbound
config>service>vprn# no encryption-keygroup 2 direction
 outbound
```

The following example shows that the key group configuration has been removed from an SDP or a VPRN service.

```
ALU-1:Sar18>config>service# info
----------------------------------------------
...
        sdp 61 create
            shutdown
            far-end 10.10.10.10
            exit
        exit
...
...
        vprn 22 customer 1 create
            shutdown
        exit
...
----------------------------------------------
ALU-1:Sar18>config>service# info
```

Changing the Key Group for an SDP or VPRN Service

Use the following sequence of CLI commands to change key groups:

1. Remove the inbound direction key group on each node for the service.
2. Change the outbound direction key group on each node for the service.
3. Install the new inbound direction key group on each node for the service.

The following CLI syntax changes the key group on an SDP. The syntax for a VPRN service is similar. In the example below, the inbound and outbound key groups are changed from key group 4 to key group 6.

CLI Syntax:

```
config>service# sdp sdp-id
     no encryption-keygroup keygroup-id direction
        {inbound|outbound}
```

Example:

```
config>service# sdp 61
config>service>sdp# no encryption-keygroup 4 direction
 inbound
config>service>sdp# encryption-keygroup 6 direction
 outbound
config>service>sdp# encryption-keygroup 6 direction
 inbound
```

The following example shows that the key group configuration has been changed for the SDP or the VPRN service.

```
ALU-1:Sar18>config>service# info
----------------------------------------------
...
        sdp 61 create
            shutdown
            far-end 10.10.10.10
            exit
            encryption-keygroup 6 direction inbound
            encryption-keygroup 6 direction outbound
        exit
...
...
        vprn 22 customer 1 create
            shutdown
            encryption-keygroup 2 direction inbound
            encryption-keygroup 2 direction outbound
        exit
...
----------------------------------------------
ALU-1:Sar18>config>service# info
```

Deleting a Key Group from a 7705 SAR

To delete a key group from a 7705 SAR, the key group must be removed (unbound) from all SDPs and VPRN services that use it.

To locate the key group bindings, use the CLI command `show>group-encryption> encryption-keygroup` *`keygroup-id`*.

Use the following CLI syntax to delete a key group:

CLI Syntax: `config# group-encryption`
`no encryption-keygroup` *`keygroup-id`*

Example: `config>group-encryption# no encryption-keygroup 8`

NGE Command Reference

Command Hierarchies

- Configuration Commands
  - → NGE Commands
  - → Services Commands
- Show Commands

Configuration Commands

NGE Commands

```
config
    — group-encryption
        — encryption-keygroup keygroup-id [create]
        — no encryption-keygroup keygroup-id
            — active-outbound-sa spi
            — no active-outbound-sa
            — description description-string
            — no description
            — esp-auth-algorithm {sha256 | sha512}
            — no esp-auth-algorithm
            — esp-encryption-algorithm {aes128 | aes256}
            — no esp-encryption-algorithm
            — keygroup-name keygroup-name
            — no keygroup-name
            — security-association spi spi authentication-key authentication-key encryption-key
                encryption-key [crypto]
            — no security-association spi spi
        — group-encryption-label encryption-label
        — no group-encryption-label
```

Services Commands

```
config
    — service
        — sdp
            — encryption-keygroup keygroup-id direction {inbound | outbound}
            — no encryption-keygroup direction {inbound | outbound}
        — vprn
            — encryption-keygroup keygroup-id direction {inbound | outbound}
            — no encryption-keygroup direction {inbound | outbound}
```

See Global Service Command Reference for information on encryption key groups for an SDP and VPRN Services Command Reference for information on encryption key groups for a VPRN service.

Show Commands

```
show
    — group-encryption
        — encryption-keygroup keygroup-id
        — encryption-keygroup keygroup-id spi spi
        — summary
```

Command Descriptions

- Configuration Commands
- Show Commands

Configuration Commands

- Generic Commands
- Group Encryption Commands

Generic Commands

description

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

Context config>grp-encryption>encryp-keygrp

Description This command is used to add a description to the key group being referenced.

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

Default n/a

Parameters *description-string* — the description of the key group

Group Encryption Commands

active-outbound-sa

Syntax **active-outbound-sa** *spi*
no active-outbound-sa

Context config>grp-encryption>encryp-keygrp

Description This command specifies the Security Association, referenced by the Security Parameter Index (SPI), to use when performing encryption and authentication on NGE packets egressing the node for all services configured using this key group.

The **no** form of the command returns the parameter to its default value and is the same as removing this key group from all outbound direction key groups in all services configured with this key group (that is, all packets of services using this key group will egress the node in without being encrypted).

Default n/a

Parameters *spi* — specifies the SPI to use for packets of services using this key group when egressing the node

Values 1 to 127

encryption-keygroup

Syntax **encryption-keygroup** *keygroup-id* [**create**]
no encryption-keygroup *keygroup-id*

Context config>grp-encryption

Description This command is used to create a key group. Once the key group is created, use the command to enter the key group context or delete a key group.

The **no** form of the command removes the key group. Before using the **no** form, the key group association must be deleted from all services that are using this key group.

Default n/a

Parameters *keygroup-id* — the number of the key group being referenced

Values 1 to 15, or *keygroup-name* (up to 64 characters)

create — mandatory keyword when creating a key group

esp-auth-algorithm

Syntax **esp-auth-algorithm** {**sha256** | **sha512**}
no esp-auth-algorithm

Context config>grp-encryption>encryp-keygrp

Description This command specifies the hashing algorithm used to perform authentication on the Encapsulating Security Payload (ESP) within NGE packets for services configured using this key group. All SPI entries must be deleted before the **no** form of the command may be entered or the **esp-auth-algorithm** value changed from its current value.

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

Default sha256

Parameters **sha256 —** configures the ESP to use the HMAC-SHA-256 algorithm for authentication

sha512 — configures the ESP to use the HMAC-SHA-512 algorithm for authentication

esp-encryption-algorithm

Syntax **esp-encryption-algorithm** {**aes128** | **aes256**}
no esp-encryption-algorithm

Context config>grp-encryption>encryp-keygrp

Description This command specifies the encryption algorithm used to perform encryption on the Encapsulating Security Payload (ESP) within NGE packets for services configured using this key group. All SPI entries must be deleted before the **no** form of the command may be entered or the **esp-encryption-algorithm** value changed from its current value.

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

Default aes128

Parameters **aes128 —** configures the AES algorithm with a block size of 128 bits. This is a very strong algorithm choice.

aes256 — configures the AES algorithm with a block size of 256 bits. This is the strongest available version of AES.

group-encryption

Syntax **group-encryption**

Context config

Description This command creates the context to configure group encryption parameters.

keygroup-name

Syntax **keygroup-name** *keygroup-name*
no keygroup-name

Context config>grp-encryption>encryp-keygrp

Description This command is used to name the key group. The key-group name can be used to reference a key group when configuring services or displaying information.

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

Default n/a

Parameters *keygroup-name* — up to 64 characters

security-association

Syntax **security-association spi** *spi* **authentication-key** *authentication-key* **encryption-key** *encryption-key* [**crypto**]
no security-association spi *spi*

Context config>grp-encryption>encryp-keygrp

Description This command is used to create a security association for a specific SPI value in a key group. The command is also used to enter the authentication and encryption key values for the security association, or to delete a security association.

The SPI value used for the security association is a node-wide unique value, meaning that no two security associations in any key group on the node may share the same SPI value.

Keys are entered in clear text. Once configured, they are never displayed in their original, clear text form. Keys are displayed in a 7705 SAR OS-encrypted form, which is indicated by the system-appended `crypto` keyword when an **info** or an **admin>save** command is run. For security reasons, keys encrypted on one node are not usable on other nodes (that is, keys are not exchangeable between nodes).

The **no** form of the command removes the security association and related key values from the list of security associations for the key group. If the **no** form of the command is attempted using the same SPI value that is configured for **active-outbound-sa**, then a warning is issued and the command is blocked. If the **no** form of the command is attempted on the last SPI in the key group and the key group is configured on a service, then the command is blocked.

Default n/a

Parameters *spi* — specifies the SPI ID of the SPI being referenced for the security association

Values 1 to 127

authentication-key — specifies the authentication key for the SPI, in hexadecimal format. The number of characters in the hexadecimal string must be 64 or 128, depending on whether the authentication algorithm is set to sha256 or sha512, respectively.

encryption-key — specifies the encryption key for the SPI, in hexadecimal format. The number of characters in the hexadecimal string must be 32 or 64, depending on whether the encryption algorithm is set to aes128 or aes256, respectively.

crypto — indicates that the keys showing on the CLI **info** display are in 7705 SAR OS-encrypted form

group-encryption-label

Syntax **group-encryption-label** *encryption-label*
no group-encryption-label

Context config>grp-encryption

Description This command configures the group encryption label used to identify when an MPLS payload is encrypted. This label must be unique network-wide and must be configured consistently on all nodes participating in a network group encryption domain. The label cannot be changed or deleted when there are any key groups configured on the node.

The **no** form of the command reverts to the default setting.

Default n/a

Parameters *encryption-label* — the network-wide, unique reserved MPLS label for group encryption

Values 32 to 2047

Show Commands

group-encryption

Syntax **group-encryption**

Context show

Description This command accesses the **show>group encryption** context.

encryption-keygroup

Syntax **encryption-keygroup** *keygroup-id*
encryption-keygroup *keygroup-id* **spi** *spi*

Context show>grp-encryp

Description This command displays NGE information for a key group.

Parameters *keygroup-id* — specifies the key group identifier to use for the output display

Values 1 to 15 or *keygroup-name* (up to 64 characters)

spi — specifies the SPI to use for the output display

Output The following output is an example of encryption key group information, and Table 150 describes the fields.

Sample Output

```
*A:7705custDoc:Sar18>show>grp-encryp#  encryption-keygroup 2
===============================================================================
Encryption Keygroup Configuration Detail
===============================================================================
Keygroup Id         : 2
Keygroup Name       : KG1_secure
Description         : Most_secure_KG
Authentication Algo: sha256
Encryption Algo     : aes128
Active Outbound SA : 6
Activation Time     : 04/20/2015 20:07:31
-------------------------------------------------------------------------------
Security Associations
-------------------------------------------------------------------------------
Spi                 : 2
Install Time        : 04/20/2015 20:08:17
Key CRC             : 0x806fb970
Spi                 : 6
Install Time        : 04/20/2015 19:43:40
Key CRC             : 0xa4f2d262
```

```
-------------------------------------------------------------------------------
Encryption Keygroup Forwarded Statistics
-------------------------------------------------------------------------------
Encrypted Pkts          : 0             Encrypted Bytes         : 0
Decrypted Pkts          : 0             Decrypted Bytes         : 0
-------------------------------------------------------------------------------
Encryption Keygroup Outbound Discarded Statistics (Pkts)
-------------------------------------------------------------------------------
Total Discard           : 0             Unsupported Uplink      : 0
Enqueue Error           : 0             Other                   : 0
-------------------------------------------------------------------------------
Encryption Keygroup Inbound Discarded Statistics (Pkts)
-------------------------------------------------------------------------------
Total Discard           : 0             Invalid Spi             : 0
Authentication Failure *: 0             Control Word Mismatch   : 0
Padding Error           : 0             Enqueue Error           : 0
Other                   : 0
-------------------------------------------------------------------------------

-----------------------------------------------
SDP Keygroup Association Table
-----------------------------------------------
SDP ID         Direction
-----------------------------------------------
61             Inbound   Outbound
-----------------------------------------------
Inbound Keygroup SDP Association Count:  1
Outbound Keygroup SDP Association Count: 1

-----------------------------------------------
VPRN Keygroup Association Table
-----------------------------------------------
VPRN SVC ID    Direction
-----------------------------------------------
12             Inbound   Outbound
-----------------------------------------------
Inbound Keygroup VPRN Association Count:  1
Outbound Keygroup VPRN Association Count: 1
-----------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
A:ALU-1:Sar18>show>grp-encryp#

*A:7705:ALU-1# show group-encryption encryption-keygroup 1 spi 1
===============================================================================
Encryption Keygroup Security Association Detail
===============================================================================
Keygroup Id      : 1                    SPI Id           : 1
Install Time     : 06/16/2015 11:28:49
Key CRC          : 0x36e5af55
-------------------------------------------------------------------------------
Encryption Keygroup Security Association Forwarded Statistics
-------------------------------------------------------------------------------
Encrypted Pkts          : 1662534       Encrypted Bytes         : 837917136
Decrypted Pkts          : 1662333       Decrypted Bytes         : 837815832
```

```
-------------------------------------------------------------------------------
Encryption Keygroup Security Association Outbound Discarded Statistics (Pkts)
-------------------------------------------------------------------------------
Total Discard           : 0             Enqueue Error           : 0
Other                   : 0
-------------------------------------------------------------------------------
Encryption Keygroup Security Association Inbound Discarded Statistics (Pkts)
-------------------------------------------------------------------------------
Total Discard           : 0             Authentication Failure  : 0
Control Word Mismatch   : 0             Padding Error           : 0
Enqueue Error           : 0             Other                   : 0
===============================================================================
```

Table 150: Show Encryption Key Group Output Fields

Label	Description
Encryption Keygroup Configuration Detail	
Keygroup Id	The key group identifier
Keygroup Name	The key group name
Description	The key group description
Authentication Algo	The authentication algorithm used for the key group
Encryption Algo	The encryption algorithm used for the key group
Active Outbound SA	The active outbound SA for the key group
Activation Time	The date and time that the key group was activated
Security Associations	
Spi	The security parameter index for the SA in the key group
Install Time	The date and time that the SA was installed in the key group
Key CRC	The CRC for the key belonging to the SA
Encryption Keygroup Forwarded Statistics	
Encrypted Pkts	The number of encrypted packets forwarded by the key group
Encrypted Bytes	The number of encrypted bytes forwarded by the key group
Decrypted Pkts	The number of decrypted packets forwarded by the key group
Decrypted Bytes	The number of decrypted bytes forwarded by the key group

Table 150: Show Encryption Key Group Output Fields (Continued)

Label	Description
Encryption Keygroup Outbound Discarded Statistics (Pkts)	
Total Discard	The total number of outbound packets discarded by the key group
Unsupported Uplink	The total number of outbound packets discarded by the key group due to an unsupported uplink
Enqueue Error	The total number of outbound packets discarded by the key group due to an enqueuing error
Other	The total number of outbound packets discarded by the key group due to some other reason, such as an internal configuration error (for example, a key group that points to an SA, but the SA is not valid)
Encryption Keygroup Inbound Discarded Statistics (Pkts)	
Total Discard	The total number of inbound packets discarded by the key group
Invalid Spi	The total number of inbound packets discarded by the key group due to an invalid SPI
Authentication Failure *	The total number of inbound packets discarded by the key group due to an authorization failure
Control Word Mismatch	The total number of inbound packets discarded by the key group due to a control word (CW) mismatch between the encrypted (protected) CW in the ESP payload and the CW that is not encrypted
Padding Error	The total number of inbound packets discarded by the key group due to a padding error
Enqueue Error	The total number of inbound packets discarded by the key group due to an enqueuing error
Other	The total number of inbound packets discarded by the key group due to some other reason (for example, an incoming packet length is incorrect)
SDP Keygroup Association Table	
SDP ID	The SDP ID
Direction	The direction in which key group authentication and encryption occurs for traffic on the SDP
Inbound Keygroup SDP Association Count	The number of SDPs configured to use inbound SA

Table 150: Show Encryption Key Group Output Fields (Continued)

Label	Description
Outbound Keygroup SDP Association Count	The number of SDPs configured to use outbound SA
VPRN Keygroup Association Table	
VPRN SVC ID	The VPRN service identifier
Direction	The direction in which key group authentication and encryption occurs for traffic on the VPRN
Inbound Keygroup VPRN Association Count	The number of VPRNs configured to use inbound SA
Outbound Keygroup VPRN Association Count	The number of VPRNs configured to use outbound SA

summary

Syntax **summary**

Context show>grp-encryp

Description This command shows NGE summary information.

Output The following output is an example of NGE summary information, and Table 151 describes the fields.

Sample Output

```
A:ALU-1:Sar18>show>grp-encryp# summary
============================
Group Encryption
============================
Encryption Label : 34
============================
=======================================================
Encryption Keygroup
=======================================================
Id Name         Auth Algo    Encr Algo    Active OutSA
-------------------------------------------------------
2  KG1_secure   sha256       aes128                 6
4               sha256       aes128                 0
-------------------------------------------------------
No. of Encryption Keygroup: 2
=======================================================
A:ALU-1:Sar18>show>grp-encryp#
```

Table 151: Show Group Encryption Summary Output Fields

Label	Description
Group Encryption	
Encryption Label	The unique network-wide group encryption label
Encryption Keygroup	
Id	The key group identifier value
Name	The key group name
Auth Algo	The authentication algorithm used by the key group
Encr Algo	The encryption algorithm used by the key group
Active OutSA	The active outbound SA for the key group
No. of Encryption Keygroup	The number of encryption key groups currently configured on the node

List of Acronyms

Table 152: Acronyms

Acronym	Expansion
2G	second generation wireless telephone technology
3DES	triple DES (data encryption standard)
3G	third generation mobile telephone technology
5620 SAM	5620 Service Aware Manager
6VPE	IPv6 on Virtual Private Edge Router
7705 SAR	7705 Service Aggregation Router
7710 SR	7710 Service Router
7750 SR	7750 Service Router
9500 MPR	9500 microwave packet radio
ABR	area border router available bit rate
AC	alternating current attachment circuit
ACK	acknowledge
ACL	access control list
ACR	adaptive clock recovery
ADM	add/drop multiplexer
ADP	automatic discovery protocol
AES	advanced encryption standard
AFI	authority and format identifier
AIS	alarm indication signal

Table 152: Acronyms (Continued)

Acronym	Expansion
ANSI	American National Standards Institute
Apipe	ATM VLL
APS	automatic protection switching
ARP	address resolution protocol
A/S	active/standby
AS	autonomous system
ASAP	any service, any port
ASBR	autonomous system boundary router
ASM	any-source multicast autonomous system message
ASN	autonomous system number
ATM	asynchronous transfer mode
ATM PVC	ATM permanent virtual circuit
B3ZS	bipolar with three-zero substitution
Batt A	battery A
B-bit	beginning bit (first packet of a fragment)
Bc	committed burst size
Be	excess burst size
BECN	backward explicit congestion notification
Bellcore	Bell Communications Research
BFD	bidirectional forwarding detection
BGP	border gateway protocol
BITS	building integrated timing supply
BMCA	best master clock algorithm

Table 152: Acronyms (Continued)

Acronym	Expansion
BMU	broadcast, multicast, and unknown traffic Traffic that is not unicast. Any nature of multipoint traffic: • broadcast (that is, all 1s as the destination IP to represent all destinations within the subnet) • multicast (that is, traffic typically identified by the destination address, uses special destination address); for IP, the destination must be 224.0.0.0 to 239.255.255.255 • unknown (that is, the destination is typically a valid unicast address but the destination port/interface is not yet known; therefore, traffic needs to be forwarded to all destinations; unknown traffic is treated as broadcast)
BOF	boot options file
BPDU	bridge protocol data unit
BRAS	Broadband Remote Access Server
BSC	Base Station Controller
BSR	bootstrap router
BSTA	Broadband Service Termination Architecture
BTS	base transceiver station
CA	certificate authority
CAS	channel associated signaling
CBN	common bonding networks
CBS	committed buffer space
CC	continuity check control channel
CCM	continuity check message
CE	circuit emulation customer edge
CEM	circuit emulation

Table 152: Acronyms (Continued)

Acronym	Expansion
CES	circuit emulation services
CESoPSN	circuit emulation services over packet switched network
CFM	connectivity fault management
cHDLC	Cisco high-level data link control protocol
CIDR	classless inter-domain routing
CIR	committed information rate
CLI	command line interface
CLP	cell loss priority
CMP	certificate management protocol
CoS	class of service
CPE	customer premises equipment
Cpipe	circuit emulation (or TDM) VLL
CPM	Control and Processing Module (CPM is used instead of CSM when referring to CSM filtering to align with CLI syntax used with other SR products). CSM management ports are referred to as CPM management ports in the CLI.
CPU	central processing unit
C/R	command/response
CRC	cyclic redundancy check
CRC-32	32-bit cyclic redundancy check
CRL	certificate revocation list
CRON	a time-based scheduling service (from chronos = time)
CRP	candidate RP
CSM	Control and Switching Module
CSNP	complete sequence number PDU
CSPF	constrained shortest path first
C-TAG	customer VLAN tag

Table 152: Acronyms (Continued)

Acronym	Expansion
CV	connection verification customer VLAN (tag)
CW	control word
CWDM	coarse wavelength-division multiplexing
DC	direct current
DC-C	DC return - common
DCE	data communications equipment
DC-I	DC return - isolated
DCO	digitally controlled oscillator
DCR	differential clock recovery
DDoS	distributed DoS
DE	discard eligibility
DER	distinguished encoding rules
DES	data encryption standard
DF	do not fragment
DH	Diffie-Hellman
DHB	decimal, hexadecimal, or binary
DHCP	dynamic host configuration protocol
DHCPv6	dynamic host configuration protocol for IPv6
DIS	designated intermediate system
DLCI	data link connection identifier
DLCMI	data link connection management interface
DM	delay measurement
DNS	domain name server
DNU	do not use
DoS	denial of service

Table 152: Acronyms (Continued)

Acronym	Expansion
dot1p	IEEE 802.1p bits, in Ethernet or VLAN ingress packet headers, used to map traffic to up to eight forwarding classes
dot1q	IEEE 802.1q encapsulation for Ethernet interfaces
DPD	dead peer detection
DPI	deep packet inspection
DPLL	digital phase locked loop
DR	designated router
DSA	digital signal algorithm
DSCP	differentiated services code point
DSL	digital subscriber line
DSLAM	digital subscriber line access multiplexer
DTE	data termination equipment
DU	downstream unsolicited
DUID	DHCP unique identifier
DUS	do not use for synchronization
DV	delay variation
e911	enhanced 911 service
EAP	Extensible Authentication Protocol
EAPOL	EAP over LAN
E-bit	ending bit (last packet of a fragment)
E-BSR	elected BSR
ECMP	equal cost multipath
EE	end entity
EFM	Ethernet in the first mile
EGP	exterior gateway protocol
EIA/TIA-232	Electronic Industries Alliance/Telecommunications Industry Association Standard 232 (also known as RS-232)

Table 152: Acronyms (Continued)

Acronym	Expansion
EIR	excess information rate
EJBCA	Enterprise Java Bean Certificate Authority
eLER	egress label edge router
E&M	ear and mouth earth and magneto exchange and multiplexer
eMBMS	evolved MBMS
EPC	evolved packet core
Epipe	Ethernet VLL
EPL	Ethernet private line
EPON	Ethernet Passive Optical Network
EPS	equipment protection switching
ERO	explicit route object
ESD	electrostatic discharge
ESMC	Ethernet synchronization message channel
ESN	extended sequence number
ESP	encapsulating security payload
ETE	end-to-end
ETH-CFM	Ethernet connectivity fault management (IEEE 802.1ag)
EVDO	evolution - data optimized
EVPL	Ethernet virtual private link
EXP bits	experimental bits (currently known as TC)
FC	forwarding class
FCS	frame check sequence
FD	frequency diversity
FDB	forwarding database
FDL	facilities data link

Table 152: Acronyms (Continued)

Acronym	Expansion
FEAC	far-end alarm and control
FEC	forwarding equivalence class
FECN	forward explicit congestion notification
FeGW	far-end gateway
FF	fixed filter
FFD	fast fault detection
FIB	forwarding information base
FIFO	first in, first out
FNG	fault notification generator
FOM	figure of merit
Fpipe	frame relay VLL
FQDN	fully qualified domain name
FR	frame relay
FRG bit	fragmentation bit
FRR	fast reroute
FTN	FEC-to-NHLFE
FTP	file transfer protocol
FXO	foreign exchange office
FXS	foreign exchange subscriber
GFP	generic framing procedure
GigE	Gigabit Ethernet
GLONASS	Global Navigation Satellite System (Russia)
GNSS	global navigation satellite system (generic)
GPON	Gigabit Passive Optical Network
GPS	Global Positioning System
GRE	generic routing encapsulation

Table 152: Acronyms (Continued)

Acronym	Expansion
GRT	global routing table
GSM	Global System for Mobile Communications (2G)
HA	high availability
HCM	high capacity multiplexing
HDB3	high density bipolar of order 3
HDLC	high-level data link control protocol
HEC	header error control
HMAC	hash message authentication code
Hpipe	HDLC VLL
H-QoS	hierarchical quality of service
HSB	hot standby
HSDPA	high-speed downlink packet access
HSPA	high-speed packet access
HVPLS	hierarchical virtual private line service
IANA	internet assigned numbers authority
IBN	isolated bonding networks
ICB	inter-chassis backup
ICMP	Internet control message protocol
ICMPv6	Internet control message protocol for IPv6
ICP	IMA control protocol cells
IDS	intrusion detection system
IDU	indoor unit
IEEE	Institute of Electrical and Electronics Engineers
IEEE 1588v2	Institute of Electrical and Electronics Engineers standard 1588-2008
IES	Internet Enhanced Service

Table 152: Acronyms (Continued)

Acronym	Expansion
IETF	Internet Engineering Task Force
IGMP	internet group management protocol
IGP	interior gateway protocol
IID	instance ID
IKE	internet key exchange
iLER	ingress label edge router
ILM	incoming label map
IMA	inverse multiplexing over ATM
INVARP	inverse address resolution protocol
IOM	input/output module
IP	Internet Protocol
IPCP	Internet Protocol Control Protocol
IPIP	IP in IP
Ipipe	IP interworking VLL
IPoATM	IP over ATM
IPS	intrusion prevention system
IPSec	Internet Protocol security
ISA	integrated services adapter
ISAKMP	internet security association and key management protocol
IS-IS	Intermediate System-to-Intermediate System
IS-IS-TE	IS-IS-traffic engineering (extensions)
ISO	International Organization for Standardization
IW	interworking
JP	join prune
KG	key group
LB	loopback

Table 152: Acronyms (Continued)

Acronym	Expansion
lbf-in	pound force inch
LBM	loopback message
LBO	line buildout
LBR	loopback reply
LCP	link control protocol
LDP	label distribution protocol
LER	label edge router
LFIB	label forwarding information base
LIB	label information base
LLDP	link layer discovery protocol
LLDPDU	link layer discovery protocol data unit
LLF	link loss forwarding
LLID	loopback location ID
LM	loss measurement
LMI	local management interface
LOS	line-of-sight loss of signal
LSA	link-state advertisement
LSDB	link-state database
LSP	label switched path link-state PDU (for IS-IS)
LSR	label switch router link-state request
LSU	link-state update
LT	linktrace
LTE	long term evolution line termination equipment

Table 152: Acronyms (Continued)

Acronym	Expansion
LTM	linktrace message
LTN	LSP ID to NHLFE
LTR	link trace reply
MA	maintenance association
MAC	media access control
MA-ID	maintenance association identifier
MBB	make-before-break
MBMS	multimedia broadcast multicast service
MBS	maximum buffer space maximum burst size media buffer space
MBSP	mobile backhaul service provider
MCAC	multicast connection admission control
MC-APS	multi-chassis automatic protection switching
MC-MLPPP	multi-class multilink point-to-point protocol
MCS	multicast server multi-chassis synchronization
MCT	MPT craft terminal
MD	maintenance domain
MD5	message digest version 5 (algorithm)
MDA	media dependent adapter
MDDB	multidrop data bridge
MDL	maintenance data link
ME	maintenance entity
MED	multi-exit discriminator
MEF	Metro Ethernet Forum
MEG	maintenance entity group

Table 152: Acronyms (Continued)

Acronym	Expansion
MEG-ID	maintenance entity group identifier
MEN	Metro Ethernet network
MEP	maintenance association end point
MFC	multi-field classification
MHF	MIP half function
MIB	management information base
MI-IS-IS	multi-instance IS-IS
MIR	minimum information rate
MLD	multicast listener discovery
MLPPP	multilink point-to-point protocol
MP	merge point multilink protocol
MP-BGP	multiprotocol border gateway protocol
MPLS	multiprotocol label switching
MPLSCP	multiprotocol label switching control protocol
MPP	MPT protection protocol
MPR	see 9500 MPR
MPR-e	microwave packet radio-standalone mode
MPT	microwave packet transport
MPT-HC V2/9558HC	microwave packet transport, high capacity version 2
MPT-HLC	microwave packet transport, high-capacity long-haul cubic (ANSI)
MPT-HQAM	microwave packet transport, high capacity (MPT-HC-QAM) or extended power (MPT-XP-QAM) with 512/1024 QAM
MPT-MC	microwave packet transport, medium capacity
MPT-XP	microwave packet transport, high capacity (very high power version of MPT-HC V2/9558HC)

Table 152: Acronyms (Continued)

Acronym	Expansion
MRRU	maximum received reconstructed unit
MRU	maximum receive unit
MSDU	MAC Service Data Unit
MSO	multi-system operator
MS-PW	multi-segment pseudowire
MTIE	maximum time interval error
MTSO	mobile trunk switching office
MTU	maximum transmission unit multi-tenant unit
M-VPLS	management virtual private line service
MVR	multicast VPLS registration
MW	microwave
MWA	microwave awareness
N·m	newton meter
NAT	network address translation
NAT-T	network address translation traversal
NBMA	non-broadcast multiple access (network)
ND	neighbor discovery
NE	network element
NET	network entity title
NGE	network group encryption
NHLFE	next hop label forwarding entry
NHOP	next-hop
NLOS	non-line-of-sight
NLPID	network level protocol identifier
NLRI	network layer reachability information

Table 152: Acronyms (Continued)

Acronym	Expansion
NNHOP	next next-hop
NNI	network-to-network interface
Node B	similar to BTS but used in 3G networks — term is used in UMTS (3G systems) while BTS is used in GSM (2G systems)
NSAP	network service access point
NSP	native service processing
NSSA	not-so-stubby area
NTP	network time protocol
NTR	network timing reference
OADM	optical add/drop multiplexer
OAM	operations, administration, and maintenance
OAMPDU	OAM protocol data units
OC3	optical carrier level 3
OCSP	online certificate status protocol
ODU	outdoor unit
OIF	outgoing interface
OLT	optical line termination
OMC	optical management console
ONT	optical network terminal
OOB	out-of-band
OPX	off premises extension
ORF	outbound route filtering
OS	operating system
OSI	Open Systems Interconnection (reference model)
OSINLCP	OSI Network Layer Control Protocol
OSPF	open shortest path first

Table 152: Acronyms (Continued)

Acronym	Expansion
OSPF-TE	OSPF-traffic engineering (extensions)
OSS	operations support system
OSSP	organization specific slow protocol
OTP	one time password
OWAMP	one-way active measurement protocol
PADI	PPPoE active discovery initiation
PADR	PPPoE active discovery request
PAE	port authentication entities
PBO	packet byte offset
PBR	policy-based routing
PBX	private branch exchange
PCP	priority code point
PCR	proprietary clock recovery
PDU	power distribution unit protocol data units
PDV	packet delay variation
PDVT	packet delay variation tolerance
PE	provider edge router
PEAPv0	protected extensible authentication protocol version 0
PEM	privacy enhanced mail
PFoE	power feed over Ethernet
PFS	perfect forward secrecy
PHB	per-hop behavior
PHY	physical layer
PIC	prefix independent convergence
PID	protocol ID

Table 152: Acronyms (Continued)

Acronym	Expansion
PIM SSM	protocol independent multicast—source-specific multicast
PIR	peak information rate
PKCS	public key cryptography standards
PKI	public key infrastructure
PLAR	private line automatic ringdown
PLCP	Physical Layer Convergence Protocol
PLR	point of local repair
PoE	power over Ethernet
PoE+	power over Ethernet plus
POP	point of presence
POS	packet over SONET
PPP	point-to-point protocol
PPPoE	point-to-point protocol over Ethernet
PPS	pulses per second
PRC	primary reference clock
PSE	power sourcing equipment
PSK	pre-shared key
PSN	packet switched network
PSNP	partial sequence number PDU
PTM	packet transfer mode
PTP	performance transparency protocol precision time protocol
PuTTY	an open-source terminal emulator, serial console, and network file transfer application
PVC	permanent virtual circuit
PVCC	permanent virtual channel connection
PW	pseudowire

Table 152: Acronyms (Continued)

Acronym	Expansion
PWE	pseudowire emulation
PWE3	pseudowire emulation edge-to-edge
Q.922	ITU-T Q-series Specification 922
QL	quality level
QoS	quality of service
RADIUS	Remote Authentication Dial In User Service
RAN	Radio Access Network
RBS	robbed bit signaling
RD	route distinguisher
RDI	remote defect indication
RED	random early discard
RESV	reservation
RIB	routing information base
RIP	routing information protocol
RJ-45	registered jack 45
RMON	remote network monitoring
RNC	Radio Network Controller
RP	rendezvous point
RPF RTM	reverse path forwarding RTM
RPS	radio protection switching
RRO	record route object
RS-232	Recommended Standard 232 (also known as EIA/TIA-232)
RSA	Rivest, Shamir, and Adleman (authors of the RSA encryption algorithm)
RSHG	residential split horizon group
RSTP	rapid spanning tree protocol

Table 152: Acronyms (Continued)

Acronym	Expansion
RSVP-TE	resource reservation protocol - traffic engineering
RT	receive/transmit
RTM	routing table manager
RTN	battery return
RTP	real-time protocol
R&TTE	Radio and Telecommunications Terminal Equipment
RTU	remote terminal unit
RU	rack unit
r-VPLS	routed virtual private LAN service
SA	security association
SAA	service assurance agent
SAFI	subsequent address family identifier
SAP	service access point
SAR-8	7705 Service Aggregation Router – 8-slot chassis
SAR-18	7705 Service Aggregation Router – 18-slot chassis
SAR-A	7705 Service Aggregation Router – two variants: • passively cooled chassis with 12 Ethernet ports and 8 T1/E1 ports • passively cooled chassis with 12 Ethernet ports and no T1/E1 ports
SAR-F	7705 Service Aggregation Router – fixed form-factor chassis
SAR-H	7705 Service Aggregation Router – temperature- and EMC-hardened to the following specifications: IEEE 1613 and IEC 61850-3
SAR-Hc	7705 Service Aggregation Router – compact version of 7705 SAR-H

Table 152: Acronyms (Continued)

Acronym	Expansion
SAR-M	7705 Service Aggregation Router – four variants: • actively cooled chassis with 16 T1/E1 ports, 7 Ethernet ports, and 1 hot-insertable module slot • actively cooled chassis with 0 T1/E1 ports, 7 Ethernet ports, and 1 hot-insertable module slot • passively cooled chassis with 16 T1/E1 ports, 7 Ethernet ports, and 0 module slots • passively cooled chassis with 0 T1/E1 ports, 7 Ethernet ports, and 0 module slots
SAR-O	7705 Service Aggregation Router passive CWDM device – three variants: • 2-wavelength CWDM dual-fiber • 4-wavelength CWDM dual-fiber • 8-wavelength CWDM single-fiber Each variant has different models that are used to add and drop different wavelengths
SAR-W	7705 Service Aggregation Router – passively cooled, universal AC and DC powered unit, equipped with five Gigabit Ethernet ports (three SFP ports and two RJ-45 Power over Ethernet (PoE) ports)

Table 152: Acronyms (Continued)

Acronym	Expansion
SAR-Wx	7705 Service Aggregation Router – passively cooled, universal AC powered unit; there are six variants: • a unit that is equipped with an AC power input connector, five Gigabit Ethernet data ports (three SFP ports and two RJ-45 Ethernet ports), and an RJ-45 alarm input connector • a unit that is equipped with an AC power input connector, five Gigabit Ethernet data ports (three SFP ports and two RJ-45 Ethernet ports), a GPS receiver, and an RJ-45 alarm input connector • a unit that is equipped with an AC power input connector, five Gigabit Ethernet data ports (three SFP ports, one RJ-45 Ethernet port, and one RJ-45 PoE+ port), and an RJ-45 alarm input connector • a unit that is equipped with an AC power input connector, five Gigabit Ethernet data ports (three SFP ports, one RJ-45 Ethernet port, and one RJ-45 PoE+ port), a GPS receiver, and an RJ-45 alarm input connector • a unit that is equipped with an AC power input connector, four Gigabit Ethernet data ports (three SFP ports and one RJ-45 port), one RJ-45 4-pair xDSL port, and an RJ-45 alarm input connector • a unit that is equipped with an AC power input connector, four Gigabit Ethernet data ports (three SFP ports and one RJ-45 port), one RJ-45 4-pair xDSL port, a GPS receiver, and an RJ-45 alarm input connector
SAR-X	7705 Service Aggregation Router – fan-cooled, rack-mountable, IP20 design, available in two variants: • AC-powered variant with a single-feed AC input that can be connected to a 100 to 240 VAC, 50/60 Hz power source • DC-powered variant with a dual-feed DC input that can be connected to a +24/-48/-60 VDC power source
SAToP	structure-agnostic TDM over packet
SCADA	surveillance, control and data acquisition

Table 152: Acronyms (Continued)

Acronym	Expansion
SC-APS	single-chassis automatic protection switching
SCP	secure copy
SD	signal degrade space diversity
SDH	synchronous digital hierarchy
SDI	serial data interface
SDP	service destination point
SE	shared explicit
SeGW	secure gateway
SF	signal fail
SFP	small form-factor pluggable (transceiver)
SFTP	SSH file transfer protocol
SGT	self-generated traffic
SHA-1	secure hash algorithm
SHG	split horizon group
SIR	sustained information rate
SLA	Service Level Agreement
SLARP	serial line address resolution protocol
SLID	subscriber location identifier of a GPON module
SLM	synthetic loss measurement
SNMP	Simple Network Management Protocol
SNPA	subnetwork point of attachment
SNR	signal to noise ratio
SNTP	simple network time protocol
SONET	synchronous optical networking
S-PE	switching provider edge router

Table 152: Acronyms (Continued)

Acronym	Expansion
SPF	shortest path first
SPI	security parameter index
SPT	shortest path tree
SR	service router (includes 7710 SR, 7750 SR)
SRLG	shared risk link group
SSH	secure shell
SSM	source-specific multicast synchronization status messaging
SSU	system synchronization unit
S-TAG	service VLAN tag
STM1	synchronous transport module, level 1
STP	spanning tree protocol
SVC	switched virtual circuit
SYN	synchronize
TACACS+	Terminal Access Controller Access-Control System Plus
TC	traffic class (formerly known as EXP bits)
TCP	transmission control protocol
TDEV	time deviation
TDM	time division multiplexing
TE	traffic engineering
TEID	tunnel endpoint identifier
TFTP	trivial file transfer protocol
T-LDP	targeted LDP
TLS	transport layer security
TLV	type length value
TM	traffic management

Table 152: Acronyms (Continued)

Acronym	Expansion
ToD	time of day
ToS	type of service
T-PE	terminating provider edge router
TPID	tag protocol identifier
TPIF	IEEE C37.94 teleprotection interface
TPMR	two-port MAC relay
TPS	transmission protection switching
TRAIM	time-receiver autonomous integrity monitoring
TTL	time to live
TTLS	tunneled transport layer security
TTM	tunnel table manager
TWAMP	two-way active measurement protocol
U-APS	unidirectional automatic protection switching
UBR	unspecified bit rate
UDP	user datagram protocol
UMTS	Universal Mobile Telecommunications System (3G)
UNI	user-to-network interface
uRPF	unicast reverse path forwarding
V.11	ITU-T V-series Recommendation 11
V.24	ITU-T V-series Recommendation 24
V.35	ITU-T V-series Recommendation 35
VC	virtual circuit
VCC	virtual channel connection
VCCV	virtual circuit connectivity verification
VCI	virtual circuit identifier
VID	VLAN ID

Table 152: Acronyms (Continued)

Acronym	Expansion
VLAN	virtual LAN
VLL	virtual leased line
VoIP	voice over IP
Vp	peak voltage
VP	virtual path
VPC	virtual path connection
VPI	virtual path identifier
VPLS	virtual private LAN service
VPN	virtual private network
VPRN	virtual private routed network
VRF	virtual routing and forwarding table
VRRP	virtual router redundancy protocol
VSE	vendor-specific extension
VSO	vendor-specific option
VT	virtual trunk
WCDMA	wideband code division multiple access (transmission protocol used in UMTS networks)
WRED	weighted random early discard
WTR	wait to restore
X.21	ITU-T X-series Recommendation 21
XRO	exclude route object

Standards and Protocol Support

This chapter lists the 7705 SAR compliance with EMC, environmental, and safety standards, telecom standards, and supported protocols:

- EMC Industrial Standards Compliance
- EMC Regulatory and Customer Standards Compliance
- Environmental Standards Compliance
- Safety Standards Compliance
- Telecom Interface Compliance
- Directives, Regional Approvals and Certifications Compliance
- Telecom Standards
- Protocol Support
- Proprietary MIBs

Table 153: EMC Industrial Standards Compliance

Standard	Title	Platform								
		SAR-X	SAR-A	SAR-M	SAR-8	SAR-18	SAR-H	SAR-Hc	SAR-W	SAR-Wx
IEEE 1613:2009 + A1:2011	IEEE Standard Environmental and Testing Requirements for Communications Networking Devices Installed in Electric Power Substations	✓[1]			✓[2]	✓[1]	✓[4]	✓[4]		
IEEE 1613.1-2013	IEEE Standard Environmental and Testing Requirements for Communications Networking Devices Installed in Transmission and Distribution Facilities	✓[5]			✓[6]	✓[3]	✓[7]	✓[7]		
IEEE Std C37.90	IEEE Standard for relays and relay systems associated with Electric Power Apparatus	✓			✓	✓	✓	✓		
IEEE Std C37.90.1	Surge Withstand Capability (SWC) Tests	✓			✓	✓	✓	✓		
IEEE Std C37.90.2	Withstand Capability of Relay Systems to Radiated Electromagnetic Interference from Transceivers	✓			✓	✓	✓	✓		
IEEE Std C37.90.3	IEEE Standard Electrostatic Discharge Tests for Protective Relays	✓			✓	✓	✓	✓		
EN 50121-4: 2006	Electromagnetic Compatibility – Part 4: Emission and Immunity of the Signalling and Telecommunications Apparatus	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC 62236-4:2008	Electromagnetic Compatibility – Part 4: Emission and Immunity of the Signalling and Telecommunications Apparatus	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC 61000-6-2:2005	Generic standards – Immunity for industrial environments	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC 61000-6-4:2006	Generic standards – Emissions standard for industrial environments	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC TS 61000-6-5	Immunity for power station and substation environments	✓			✓	✓	✓	✓		
IEC 61850-3	Communication networks and systems for power utility automation - Part 3: General requirements	✓			✓	✓	✓	✓		
IEC/AS 60870.2.1	Telecontrol equipment and systems. Operating conditions. Power supply and electromagnetic compatibility	✓			✓	✓	✓	✓		

Notes:

1. Performance Class 1
2. Performance Class 1 (Class 2 w/ Optics interfaces only)

3. Zone B; Performance Class 1
4. Performance Class 2
5. Zone A; Performance Class 1
6. Zone A; Performance Class 1 (Class 2 w/Optics interfaces only)
7. Zone A; Performance Class 2

Table 154: EMC Regulatory and Customer Standards Compliance

Standard	Title	Platform								
		SAR-X	SAR-A	SAR-M	SAR-8	SAR-18	SAR-H	SAR-Hc	SAR-W	SAR-Wx
IEC 61000-4-2	Electrostatic discharge immunity test	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC 61000-4-3	Radiated electromagnetic field immunity test	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC 61000-4-4	Electrical fast transient/burst immunity test	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC 61000-4-5	Surge immunity test	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC 61000-4-6	Immunity to conducted disturbances	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC 61000-4-8	Power frequency magnetic field immunity test	✓			✓	✓	✓	✓		
IEC 61000-4-9	Pulse Magnetic field immunity test	✓			✓	✓	✓	✓		
IEC 61000-4-10	Damped Oscillatory Magnetic Field	✓			✓	✓	✓	✓		
IEC 61000-4-11	Voltage dips, short interruptions and voltage variations immunity tests	✓	✓[1]	✓[1]	✓[1]	✓[1]	✓	✓[1]	✓	✓
IEC 61000-4-12	Oscillatory wave immunity test	✓			✓	✓	✓	✓		
IEC 61000-4-16	Conducted immunity 0 Hz - 150 kHz	✓			✓	✓	✓	✓		
IEC 61000-4-17	Ripple on d.c. input power port immunity test	✓			✓	✓	✓	✓		
IEC 61000-4-18	Damped oscillatory wave immunity test	✓			✓	✓	✓	✓		
IEC 61000-4-29	Voltage dips, short interruptions and voltage variations on d.c. input power port immunity tests	✓			✓	✓	✓	✓		
IEC 61000-3-2	Limits for harmonic current emissions (equipment input current <16A per phase)	✓	✓[1]	✓[1]	✓[1]	✓[1]	✓	✓[1]	✓	✓
IEC 61000-3-3	Limits for voltage fluctuations and flicker in low-voltage supply systems for equipment with rated current <16A	✓	✓[1]	✓[1]	✓[1]	✓[1]	✓	✓[1]	✓	✓
ITU-T K.20 (DC Ports)	Resistibility of telecommunication equipment installed in a telecommunications centre to overvoltages and overcurrents	✓	✓	✓	✓	✓	✓	✓		

Table 154: EMC Regulatory and Customer Standards Compliance (Continued)

Standard	Title	Platform								
		SAR-X	SAR-A	SAR-M	SAR-8	SAR-18	SAR-H	SAR-Hc	SAR-W	SAR-Wx
ETSI 300 132-2	Power supply interface at the input to telecommunications and datacom (ICT) equipment; Part 2: Operated by -48 V direct current (dc)	✓	✓	✓	✓	✓	✓	✓	✓	
EN 300 386	Telecommunication network equipment; ElectroMagnetic Compatibility (EMC)	✓	✓	✓	✓	✓	✓	✓	✓	✓
ES 201 468	Electromagnetic compatibility and Radio spectrum Matters (ERM); Additional ElectroMagnetic Compatibility (EMC) requirements and resistibility requirements for telecommunications equipment for enhanced availability of service in specific applications	✓			✓	✓				
EN 55024	Information technology equipment - Immunity characteristics - Limits and methods of measurements	✓	✓	✓	✓	✓	✓	✓	✓	✓
Telcordia GR-1089-CORE	EMC and Electrical Safety - Generic Criteria for Network Telecommunications Equipment	✓	✓	✓	✓	✓	✓	✓	✓	✓
AS/NZS CISPR 22	Information technology equipment - Radio disturbance characteristics - Limits and methods of measurement	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[3]	✓[3]
FCC Part 15, Subpart B	Radio Frequency devices- Unintentional Radiators (Radiated & Conducted Emissions)	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[3]	✓[3]
ICES-003	Information Technology Equipment (ITE) — Limits and methods of measurement	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[3]	✓[3]
EN 55022	Information technology equipment. Radio disturbance characteristics. Limits and methods of measurement	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[3]	✓[3]
CISPR 22	Information technology equipment. Radio disturbance characteristics. Limits and methods of measurement	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[2]	✓[3]	✓[3]
KC Notice Emission (KN22) and Immunity (KN24) (South Korea)	EMS standard: NRRA notice	✓	✓	✓	✓	✓	✓	✓		

Notes:

1. With external AC/DC power supply
2. Class A

3. Class B

Table 155: Environmental Standards Compliance

Standard	Title	Platform								
		SAR-X	SAR-A	SAR-M	SAR-8	SAR-18	SAR-H	SAR-Hc	SAR-W	SAR-Wx
IEEE 1613:2009 + A1:2011	Environmental and Testing Requirements for Communications Networking Devices	✓[1]			✓[1]	✓[1]	✓	✓		
IEC 61850-3	Communication networks and systems for power utility automation - Part 3: General requirements	✓[2]			✓[2]	✓[2]	✓[2]	✓[2]		
IEC 60068-2-1	Environmental testing – Part 2-1: Tests – Test A: Cold	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC 60068-2-2	Environmental testing - Part 2-2: Tests - Test B: Dry heat	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC 60068-2-30	Environmental testing - Part 2: Tests. Test Db and guidance: Damp heat, cyclic (12 + 12-hour cycle)	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC 60255-21-2	Electrical relays - Part 21: Vibration, shock, bump and seismic tests on measuring relays and protection equipment - Section Two: Shock and bump tests	✓			✓	✓	✓	✓		
ETSI 300 753 Class 3.2	Acoustic noise emitted by telecommunications equipment	✓	✓	✓	✓	✓	✓	✓	✓	✓
Telcordia GR-63-CORE	NEBS Requirements: Physical Protection	✓	✓	✓	✓	✓	✓	✓	✓	✓
ETSI EN 300 019-2-1 v2.1.2, Class 1.2	Specification of environmental tests; Storage	✓	✓	✓	✓	✓	✓	✓	✓	✓
ETSI EN 300 019-2-2 V2.1.2, class 2.3	Specification of environmental tests; Transportation	✓	✓	✓	✓	✓	✓	✓	✓	✓
ETSI EN 300 019-2-3 V2.2.2, class 3.2	Specification of environmental tests; Stationary use at weatherprotected locations	✓	✓	✓	✓	✓	✓	✓		
ETSI EN 300 019-2-4 v2.2.2 class T4.1	Specification of environmental tests; Stationary use at non-weatherprotected locations								✓	✓
Telcordia GR-3108-CORE	Generic Requirements for Network Equipment in the Outside Plant (OSP)	✓[3]	✓[3]	✓[3]	✓[3]		✓[3]	✓[3]	✓[4]	✓[4]

Table 155: Environmental Standards Compliance (Continued)

Standard	Title	Platform								
		SAR-X	SAR-A	SAR-M	SAR-8	SAR-18	SAR-H	SAR-Hc	SAR-W	SAR-Wx
"GR-3108 Class 3 Section 6.2 IEC 60068-2-52 - Severity 3 MIL-STD-810G Method 509.5 EN 60721-3-3 Class 3C4 EN 60068-2-11: Salt Mist EN 50155 Class ST4"	Conformal Coating [5]	✓		✓	✓		✓	✓		
Telcordia GR-950-CORE	Generic Requirements for ONU Closures and ONU Systems								✓	✓

Notes:

1. Forced air system; uses fans
2. Normal environmental conditions as per IEC 61850-3 ed.2
3. Class 2
4. Class 4
5. Conformal coating is available as an orderable option

Table 156: Safety Standards Compliance

Standard	Title	Platform								
		SAR-X	SAR-A	SAR-M	SAR-8	SAR-18	SAR-H	SAR-Hc	SAR-W	SAR-Wx
UL/CSA 60950-1	Information technology equipment - Safety - Part 1: General requirements	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC/EN 60950-1	Information technology equipment - Safety - Part 1: General requirements	✓	✓	✓	✓	✓	✓	✓	✓	✓
AS/NZS 60950-1	Information technology equipment - Safety - Part 1: General requirements	✓	✓	✓	✓	✓	✓	✓	✓	✓
IEC/EN 60825-1 and 2	Safety of laser products - Part 1: Equipment classification and requirements Part 2: Safety of optical fibre communication systems (OFCS)	✓	✓	✓	✓	✓	✓	✓	✓	✓
FDA CDRH 21-CFR 1040	PART 1040 Performance Standards for Light-Emitting Products	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 156: Safety Standards Compliance (Continued)

Standard	Title	Platform								
		SAR-X	SAR-A	SAR-M	SAR-8	SAR-18	SAR-H	SAR-Hc	SAR-W	SAR-Wx
UL/CSA 60950-22	Information Technology Equipment - Safety - Part 22: Equipment to be Installed Outdoors								✓	✓
CSA–C22.2 No.94	Special Purpose Enclosures								✓	✓
UL50	Enclosures for Electrical Equipment, Non-Environmental Consideration								✓	✓
IEC/EN 60950-22	Information technology equipment. Equipment to be installed Outdoors.								✓	✓
IEC 60529	Degrees of Protection Provided by Enclosures (IP Code)	✓[1]	✓[2]	✓[1]	✓[1]	✓[1]	✓[2]	✓[2]	✓[3]	✓[3]

Notes:

1. IP20
2. IP40
3. IP65

Table 157: Telecom Interface Compliance

Standard	Title	Platform								
		SAR-X	SAR-A	SAR-M	SAR-8	SAR-18	SAR-H	SAR-Hc	SAR-W	SAR-Wx
IC CS-03 Issue 9	Compliance Specification for Terminal Equipment, Terminal Systems, Network Protection Devices, Connection Arrangements and Hearing Aids Compatibility	✓	✓	✓	✓	✓	✓			
ACTA TIA-968-B	Telecommunications - Telephone Terminal Equipment - Technical Requirements for Connection of Terminal Equipment to the Telephone Network	✓	✓	✓	✓	✓	✓			
AS/ACIF S016 (Australia)	Requirements for Customer Equipment for connection to hierarchical digital interfaces	✓	✓	✓	✓	✓	✓			
ATIS-06000403	Network and Customer Installation Interfaces- DS1 Electrical Interfaces	✓	✓	✓	✓	✓	✓			
ANSI/TIA/EIA-422-B (RS422)	Electrical Characteristics for balanced voltage digital interfaces circuits				✓	✓				

Table 157: Telecom Interface Compliance (Continued)

Standard	Title	Platform								
		SAR-X	SAR-A	SAR-M	SAR-8	SAR-18	SAR-H	SAR-Hc	SAR-W	SAR-Wx
ITU-T G.825	The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)				✓	✓				
ITU-T G.703	Physical/electrical characteristics of hierarchical digital interfaces	✓	✓	✓	✓	✓	✓			
ITU-T G.712 (E&M)	Transmission performance characteristics of pulse code modulation channels				✓	✓				
ITU-T G.957	Optical interfaces for equipments and systems relating to the synchronous digital hierarchy				✓	✓				
ITU-T V.24 (RS232)	List of definitions for interchange circuits between data terminal equipment (DTE) and data circuit-terminating equipment (DCE)				✓	✓	✓	✓		
ITU-T V.28 (V35)	Electrical characteristics for unbalanced double-current interchange circuits				✓	✓				
ITU-T V.36 (V35)	Modems for synchronous data transmission using 60-108 kHz group band circuits				✓	✓				
ITU-T V.11 / X.27 (RS422)	Electrical characteristics for balanced double current interchange circuits operating at data signalling rates up to 10 Mbit/s				✓	✓				
ITU-T X.21 (RS422)	Interface between Data Terminal Equipment and Data Circuit-terminating Equipment for synchronous operation on public data networks				✓	✓				
IEEE 802.3at (POE)	Data Terminal Equipment Power via the Media Dependent Interfaces Enhancements			✓			✓	✓	✓	✓

Table 158: Directives, Regional Approvals and Certifications Compliance

Standard	Title	Platform								
		SAR-X	SAR-A	SAR-M	SAR-8	SAR-18	SAR-H	SAR-Hc	SAR-W	SAR-Wx
EU Directive 2004/108/EC EMC	Electromagnetic Compatibility (EMC)	✓	✓	✓	✓	✓	✓	✓	✓	✓
EU Directive 2006/95/EC LVD	Low Voltage Directive (LVD)	✓	✓	✓	✓	✓	✓	✓	✓	✓
EU Directive 2012/19/EU WEEE	Waste Electrical and Electronic Equipment (WEEE)	✓	✓	✓	✓	✓	✓	✓	✓	✓
EU Directive 2011/65/EU RoHS2	Restriction of the use of certain Hazardous Substances in Electrical and Electronic Equipment (RoHS2)	✓	✓	✓	✓	✓	✓	✓	✓	✓
CE Mark		✓	✓	✓	✓	✓	✓	✓	✓	✓
CRoHS Logo; Ministry of Information Industry order No.39		✓	✓	✓	✓	✓	✓	✓	✓	✓
China (MII NAL) Network Access License			✓	✓	✓	✓	✓		✓	
South Korea (KC Mark)		✓	✓	✓	✓	✓	✓	✓		
Australia (RCM Mark)		✓	✓	✓	✓	✓	✓	✓	✓	✓
TL9000 certified		✓	✓	✓	✓	✓	✓	✓	✓	✓
ISO 14001 certified		✓	✓	✓	✓	✓	✓	✓	✓	✓
ISO 9001:2008 certified		✓	✓	✓	✓	✓	✓	✓	✓	✓

Telecom Standards

ANSI/TIA/EIA-232-C—Interface Between Data Terminal Equipment and Data Circuit-Terminating Equipment Employing Serial Binary Data Interchange

IEEE 802.1ad—IEEE Standard for Local and Metropolitan Area Networks---Virtual Bridged Local Area Networks

IEEE 802.1ag—Service Layer OAM

IEEE 802.1p/q—VLAN Tagging

IEEE 802.3—10BaseT

IEEE 802.3ab—1000BaseT

IEEE 802.3ah—Ethernet OAM

IEEE 802.3u—100BaseTX

IEEE 802.3x —Flow Control

IEEE 802.3z—1000BaseSX/LX

IEEE 802.3-2008—Revised base standard

IEEE 802.1AX-2008—Link Aggregation Task Force (transferred from IEEE 802.3ad)

IEEE C37.94-2002—N Times 64 Kilobit Per Second Optical Fiber Interfaces Between Teleprotection and Multiplexer Equipment

ITU-T G.704—Synchronous frame structures used at 1544, 6312, 2048, 8448 and 44 736 kbit/s hierarchical levels

ITU-T G.707—Network node interface for the Synchronous Digital Hierarchy (SDH)

ITU-T G.984.1—Gigabit-capable passive optical networks (GPON): general characteristics

ITU-T Y.1564—Ethernet service activation test methodology

ITU-T Y.1731—OAM functions and mechanisms for Ethernet-based networks

Protocol Support

ATM

AF-PHY-0086.001—Inverse Multiplexing for ATM (IMA)

af-tm-0121.000—Traffic Management Specification Version 4.1, March 1999

GR-1113-CORE—Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

GR-1248-CORE—Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996

ITU-T Recommendation I.432.1—B-ISDN user-network interface - Physical layer specification: General characteristics

ITU-T Recommendation I.610—B-ISDN Operation and Maintenance Principles and Functions version 11/95

RFC 2514—Definitions of Textual Conventions and OBJECT_IDENTITIES for ATM Management, February 1999

RFC 2515—Definition of Managed Objects for ATM Management, February 1999

RFC 2684—Multiprotocol Encapsulation over ATM Adaptation Layer 5

BFD

draft-ietf-bfd-mib-00.txt—Bidirectional Forwarding Detection Management Information Base

draft-ietf-bfd-base-o5.txt—Bidirectional Forwarding Detection

draft-ietf-bfd-v4v6-1hop-06.txt—BFD IPv4 and IPv6 (Single Hop)

draft-ietf-bfd-multihop-06.txt—BFD for Multi-hop Paths

BGP

RFC 1397—BGP Default Route Advertisement

RFC 1997—BGP Communities Attribute

RFC 2385—Protection of BGP Sessions via MDS

RFC 2439—BGP Route Flap Dampening

RFC 2547bis—BGP/MPLS VPNs

RFC 2918—Route Refresh Capability for BGP-4

RFC 3107—Carrying Label Information in BGP-4

RFC 3392—Capabilities Advertisement with BGP-4

RFC 4271—BGP-4 (previously RFC 1771)

RFC 4360—BGP Extended Communities Attribute

RFC 4364—BGP/MPLS IP Virtual Private Networks (VPNs) (previously RFC 2574bis BGP/MPLS VPNs)

RFC 4456—BGP Route Reflection: Alternative to Full-mesh IBGP (previously RFC 1966 and RFC 2796)

RFC 4486—Subcodes for BGP Cease Notification Message

RFC 4724—Graceful Restart Mechanism for BGP - GR Helper

RFC 4760—Multi-protocol Extensions for BGP (previously RFC 2858)

RFC 4893—BGP Support for Four-octet AS Number Space

draft-ietf-idr-add-paths-04.txt—Advertisement of Multiple Paths in BGP

draft-ietf-idr-add-paths-guidelines-00.txt—Best Practices for Advertisement of Multiple Paths in BGP

DHCP/DHCPv6

RFC 1534—Interoperation between DHCP and BOOTP

RFC 2131—Dynamic Host Configuration Protocol (REV)

RFC 2132—DHCP Options and BOOTP Vendor Extensions

RFC 3046—DHCP Relay Agent Information Option (Option 82)

RFC 3315—Dynamic Host Configuration Protocol for IPv6

RFC 3736—Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6

Differentiated Services

RFC 2474—Definition of the DS Field in the IPv4 and IPv6 Headers

RFC 2597—Assured Forwarding PHB Group

RFC 2598—An Expedited Forwarding PHB

RFC 3140—Per-Hop Behavior Identification Codes

Digital Data Network Management

V.35

RS-232 (also known as EIA/TIA-232)

X.21

DSL Modules

ITU-T G.991.2 Annex A, B, F and ITU-T G.991.2 Amendment 2 Annex G—SHDSL standards compliance

ITU-T G.991.2 Appendix F and G—Support for up to 5696 Kb/s per pair

ITU-T G.992.1 (ADSL)

ITU-T G.992.3 (G.dmt.bis), Annex A, B, J, M

ITU-T G.992.3 Annex K.2 (ADSL2)

ITU-T G.992.5, Annex A, B, J, M

ITU-T G.992.5 Annex K (ADSL2+)

ITU-T G.993.2 Amendment 1—Seamless Rate Adaptation

ITU-T G.993.2 Annex A and Annex B—xDSL Standards Compliance (ADSL2/2+ and VDSL2)

ITU-T G.993.2 Annex K.3—Supported Transport Protocol Specific Transmission Convergence functions

ITU G.994.1 (2/07) Amendment 1 and 2—G.hs Handshake

ITU-T G.998.2—SHDSL 4-pair EFM bonding

ITU-T G.998.4 G.inp—Physical layer retransmission

ITU-T Y.1564 Ethernet service activation test methodology

RFC 2684—IEEE 802.2 LLC/SNAP bridged encapsulation while operating in ATM bonded mode

TR-060—SHDSL rate and reach

TR112 (U-R2 Deutsche Telekom AG) Version 7.0 and report of Self-Test-Result (ATU-T Register#3)

ECMP

RFC 2992—Analysis of an Equal-Cost Multi-Path Algorithm

Frame Relay

ANSI T1.617 Annex D—Signalling Specification For Frame Relay Bearer Service

ITU-T Q.922 Annex A—Digital Subscriber Signalling System No. 1 (DSS1) data link layer - ISDN data link layer specification for frame mode bearer services

FRF.1.2—PVC User-to-Network Interface (UNI) Implementation Agreement

FRF.12—Frame Relay Fragmentation Implementation Agreement

RFC 2427—Multiprotocol Interconnect over Frame Relay

GRE

RFC 2784—Generic Routing Encapsulation (GRE)

IPSec

IETF draft-nourse-scep-21.txt—Cisco Systems' Simple Certificate Enrollment Protocol

ITU-T X.690 (2002)—ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)

PKCS #12 Personal Information Exchange Syntax Standard

RFC 2315—PKCS #7: Cryptographic Message Syntax

RFC 2401—Security Architecture for the Internet Protocol

RFC 2986—PKCS #10: Certification Request Syntax Specification

RFC 3706—A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers

RFC 3947—Negotiation of NAT-Traversal in the IKE

RFC 3948—UDP Encapsulation of IPsec ESP Packets

RFC 4303—IP Encapsulating Security Payload (ESP)

RFC 4210—Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)

RFC 4211—Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)

RFC 4945—The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2, and PKIX

RFC 5280—Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile

RFC 5996—Internet Key Exchange Protocol Version 2 (IKEv2)

IPv6

RFC 2460—Internet Protocol, Version 6 (IPv6) Specification

RFC 2462—IPv6 Stateless Address Autoconfiguration

RFC 2464—Transmission of IPv6 Packets over Ethernet Networks

RFC 3587—IPv6 Global Unicast Address Format

RFC 3595—Textual Conventions for IPv6 Flow Label

RFC 4007—IPv6 Scoped Address Architecture

RFC 4193—Unique Local IPv6 Unicast Addresses

RFC 4291—IPv6 Addressing Architecture

RFC 4443—Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 Specification

RFC 4649—DHCPv6 Relay Agent Remote-ID Option

RFC 4861—Neighbor Discovery for IP version 6 (IPv6)

RFC 5095—Deprecation of Type 0 Routing Headers in IPv6

RFC 5952—A Recommendation for IPv6 Address Text Representation

IS-IS

RFC 1142—OSI IS-IS Intra-domain Routing Protocol (ISO 10589)

RFC 1195—Use of OSI IS-IS for routing in TCP/IP & dual environments

RFC 2763—Dynamic Hostname Exchange for IS-IS

RFC 2966—Domain-wide Prefix Distribution with Two-Level IS-IS

RFC 2973—IS-IS Mesh Groups

RFC 3373—Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies

RFC 3567—Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication

RFC 3719—Recommendations for Interoperable Networks using IS-IS

RFC 3784—Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)

RFC 3787—Recommendations for Interoperable IP Networks

RFC 4205 for Shared Risk Link Group (SRLG) TLV

RFC 5309—Point-to-Point Operation over LAN in Link State Routing Protocols

draft-ietf-isis-igp-p2p-over-lan-05.txt

LDP

RFC 5036—LDP Specification

RFC 5283—LDP Extension for Inter-Area Label Switched Paths

LDP and IP FRR

RFC 5286—Basic Specification for IP Fast Reroute: Loop-Free Alternates

MPLS

RFC 3031—MPLS Architecture

RFC 3032—MPLS Label Stack Encoding

RFC 3815—Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)

MPLS – OAM

RFC 4379—Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

RFC 6424— Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels

Network Management

IANA-IFType-MIB

ITU-T X.721—Information technology- OSI-Structure of Management Information

ITU-T X.734—Information technology- OSI-Systems Management: Event Report Management Function

M.3100/3120—Equipment and Connection Models

RFC 1157—SNMPv1

RFC 1850—OSPF-MIB

RFC 1907—SNMPv2-MIB

RFC 2011—IP-MIB

RFC 2012—TCP-MIB

RFC 2013—UDP-MIB

RFC 2030—Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI

RFC 2096—IP-FORWARD-MIB

RFC 2138—RADIUS

RFC 2206—RSVP-MIB

RFC 2571—SNMP-FRAMEWORKMIB

RFC 2572—SNMP-MPD-MIB

RFC 2573—SNMP-TARGET-&-NOTIFICATION-MIB

RFC 2574—SNMP-USER-BASED-SMMIB

RFC 2575—SNMP-VIEW-BASED ACM-MIB

RFC 2576—SNMP-COMMUNITY-MIB

RFC 2588—SONET-MIB

RFC 2665—EtherLike-MIB

RFC 2819—RMON-MIB

RFC 2863—IF-MIB

RFC 2864—INVERTED-STACK-MIB

RFC 3014—NOTIFICATION-LOG MIB

RFC 3164—The BSD Syslog Protocol

RFC 3273—HCRMON-MIB

RFC 3411—An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks

RFC 3412—Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)

RFC 3413—Simple Network Management Protocol (SNMP) Applications

RFC 3414—User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)

RFC 3418—SNMP MIB

draft-ietf-disman-alarm-mib-04.txt

draft-ietf-mpls-ldp-mib-07.txt

draft-ietf-ospf-mib-update-04.txt

draft-ietf-mpls-lsr-mib-06.txt

draft-ietf-mpls-te-mib-04.txt

TMF 509/613—Network Connectivity Model

OSPF

RFC 1765—OSPF Database Overflow

RFC 2328—OSPF Version 2

RFC 2370—Opaque LSA Support

RFC 2740—OSPF for IPv6

RFC 3101—OSPF NSSA Option

RFC 3137—OSPF Stub Router Advertisement

RFC 3509—Alternative Implementations of OSPF Area Border Routers

RFC 3630—Traffic Engineering (TE) Extensions to OSPF

RFC 4203 for Shared Risk Link Group (SRLG) sub-TLV

PPP

RFC 1332—PPP Internet Protocol Control Protocol (IPCP)

RFC 1570—PPP LCP Extensions

RFC 1619—PPP over SONET/SDH

RFC 1661—The Point-to-Point Protocol (PPP)

RFC 1662—PPP in HDLC-like Framing

RFC 1989—PPP Link Quality Monitoring

RFC 1990—The PPP Multilink Protocol (MP)

RFC 2686—The Multi-Class Extension to Multi-Link PPP

Pseudowires

Metro Ethernet Forum—Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks

RFC 3550—RTP: A Transport Protocol for Real-Time Applications

RFC 3985—Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture

RFC 4385—Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN

RFC 4446—IANA Allocation for PWE3

RFC 4447—Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)

RFC 4448—Encapsulation Methods for Transport of Ethernet over MPLS Networks

RFC 4553—Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)

RFC 4717—Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks

RFC 4618—Encapsulation Methods for Transport of PPP/High-Level Data Link Control (HDLC) over MPLS Networks

RFC 4619—Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks

RFC 4816—Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service

RFC 5085—Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires

RFC 5086—Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)

draft-ietf-pwe3-redundancy-02.txt—Pseudowire (PW) Redundancy

RIP

RFC 1058—Routing Information Protocol

RFC 2453—RIP Version 2

RADIUS

RFC 2865—Remote Authentication Dial In User Service

RFC 2866—RADIUS Accounting

RSVP-TE and FRR

RFC 2430—A Provider Architecture for DiffServ & TE

RFC 2961—RSVP Refresh Overhead Reduction Extensions

RFC 2702—Requirements for Traffic Engineering over MPLS

RFC 2747—RSVP Cryptographic Authentication

RFC 3097—RSVP Cryptographic Authentication - Updated Message Type Value

RFC 3209—Extensions to RSVP for LSP Tunnels

RFC 3210—Applicability Statement for Extensions to RSVP for LSP Tunnels

RFC 3477—Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)

RFC 4090—Fast Reroute Extensions to RSVP-TE for LSP Tunnels

SONET/SDH

GR-253-CORE—SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000

ITU-T Recommendation G.841—Telecommunication Standardization Section of ITU, Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum1 issued in July 2002

SSH

draft-ietf-secsh-architecture.txt—SSH Protocol Architecture

draft-ietf-secsh-userauth.txt—SSH Authentication Protocol

draft-ietf-secsh-transport.txt—SSH Transport Layer Protocol

draft-ietf-secsh-connection.txt—SSH Connection Protocol

draft-ietf-secsh- newmodes.txt—SSH Transport Layer Encryption Modes

draft-ietf-secsh-filexfer-13.txt—SSH File Transfer Protocol

Synchronization

G.781—Synchronization layer functions, 2001/09/17

G.803—Architecture of transport networks based on the synchronous digital hierarchy (SDH)

G.813—Timing characteristics of SDH equipment slave clocks (SEC)

G.823—The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy, 2003/03/16

G.824—The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy, 2003/03/16

G.8261—Timing and synchronization aspects in packet networks

G.8262—Timing characteristics of synchronous Ethernet equipment slave clock

GR 1244 CORE—Clocks for the Synchronized Network: Common Generic Criteria

IEEE Std 1588-2008—IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

ITU-T G.8264—Telecommunication Standardization Section of ITU, Distribution of timing information through packet networks, issued 10/2008

ITU-T G.8265.1—Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for frequency synchronization, issued 10/2010

ITU-T G.8275.1—Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014

RFC 5905—Network Time Protocol Version 4: Protocol and Algorithms Specification

TACACS+

IETF draft-grant-tacacs-02.txt—The TACACS+ Protocol

TCP/IP

RFC 768—User Datagram Protocol

RFC 791—Internet Protocol

RFC 792—Internet Control Message Protocol

RFC 793—Transmission Control Protocol

RFC 826—Ethernet Address Resolution Protocol

RFC 854—Telnet Protocol Specification

RFC 1350—The TFTP Protocol (Rev. 2)

RFC 1812—Requirements for IPv4 Routers

TWAMP

RFC 5357—A Two-Way Active Measurement Protocol (TWAMP)

VPLS

RFC 4762—Virtual Private LAN Services Using LDP

VRRP

RFC 2787—Definitions of Managed Objects for the Virtual Router Redundancy Protocol

RFC 3768 Virtual Router Redundancy Protocol

RFC 5798 Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6

Proprietary MIBs

TIMETRA-ATM-MIB.mib

TIMETRA-CAPABILITY-7705-V1.mib

TIMETRA-CHASSIS-MIB.mib

TIMETRA-CLEAR-MIB.mib

TIMETRA-FILTER-MIB.mib

TIMETRA-GLOBAL-MIB.mib

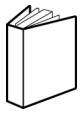
TIMETRA-LAG-MIB.mib

TIMETRA-LDP-MIB.mib

TIMETRA-LOG-MIB.mib

TIMETRA-MPLS-MIB.mib
TIMETRA-OAM-TEST-MIB.mib
TIMETRA-PORT-MIB.mib
TIMETRA-PPP-MIB.mib
TIMETRA-QOS-MIB.mib
TIMETRA-ROUTE-POLICY-MIB.mib
TIMETRA-RSVP-MIB.mib
TIMETRA-SAP-MIB.mib
TIMETRA-SDP-MIB.mib
TIMETRA-SECURITY-MIB.mib
TIMETRA-SERV-MIB.mib
TIMETRA-SYSTEM-MIB.mib
TIMETRA-TC-MIB.mib
TIMETRA-VRRP-MIB.mib

Customer documentation and product support



Customer documentation

http://documentation.alcatel-lucent.com



Technical support

http://support.alcatel-lucent.com



Documentation feedback

documentation.feedback@alcatel-lucent.com



© 2015-2016 Alcatel-Lucent. All rights reserved.
3HE 10157 AAAC TQZZA Edition 01