



7705 Service Aggregation Router

Release 25.4.R1

Services Guide

3HE 21352 AAAA TQZZA

Edition: 01

April 2025

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

© 2025 Nokia.

Table of contents

List of tables.....	27
List of figures.....	37
1 Preface.....	45
1.1 Audience.....	45
1.2 Technical support.....	45
2 7705 SAR services configuration process.....	46
3 Services overview.....	47
3.1 Introduction to services on the 7705 SAR.....	47
3.1.1 Overview.....	47
3.1.2 Service types.....	49
3.1.3 Service policies.....	50
3.2 Nokia service model.....	51
3.3 Service entities.....	52
3.3.1 Customers.....	52
3.3.2 Service types.....	52
3.3.2.1 Service names.....	53
3.3.3 Service access points.....	53
3.3.3.1 SAP encapsulation types and identifiers.....	54
3.3.3.2 SAP configuration considerations.....	57
3.3.4 Service destination points.....	58
3.3.4.1 SDP binding.....	59
3.3.4.2 Spoke and mesh SDPs.....	59
3.3.4.3 SDPs and BGP route tunnels.....	60
3.3.4.4 SDP encapsulation types.....	60
3.3.4.5 Spoke-SDP terminations.....	66
3.3.4.6 SDP ping.....	67
3.3.4.7 SDP keepalives.....	67
3.3.4.8 Mixed-LSP SDPs.....	68
3.3.4.9 Multiple load-balancing LSPs under a single SDP.....	69
3.4 High-speed download packet access offload.....	69

3.4.1	HSDPA offload.....	69
3.4.1.1	Failure detection.....	70
3.5	ETH-CFM (802.1ag and Y.1731).....	71
3.5.1	802.1ag and Y.1731 terminology.....	72
3.5.1.1	MDs, MD levels, MAs, and MEPs (802.1ag).....	73
3.5.1.2	MEG levels, MEGs, and MEPs (Y.1731).....	74
3.5.2	ETH-CFM frame format.....	74
3.5.2.1	ETH-CFM OAMPDU.....	76
3.5.2.2	CFM frame processing.....	76
3.5.2.3	MEG-ID and ICC-based format.....	78
3.5.3	ETH-CFM functions and tests.....	79
3.5.3.1	ETH-CFM Ethernet OAM tests.....	79
3.5.4	MEP support (802.1ag and Y.1731).....	84
3.5.4.1	802.1ag MEP support on Ethernet SAPs.....	84
3.5.4.2	802.1ag and Y.1731 MEP support on Ethernet spoke SDPs and mesh SDPs.....	85
3.5.4.3	Y.1731 MEP support on Ethernet SAPs.....	86
3.5.5	Priority mapping (802.1ag and Y.1731).....	87
3.5.5.1	Priority mapping for SAP Up MEPs.....	87
3.5.5.2	Priority mapping for SAP Down MEPs.....	88
3.6	G.8032 Ethernet ring protection switching.....	89
3.6.1	Overview of G.8032 operation.....	90
3.6.2	Ethernet ring sub-rings.....	95
3.6.2.1	Virtual and non-virtual channels.....	95
3.6.2.2	OAM considerations.....	100
3.6.2.3	Support service and solution combinations.....	101
3.7	QinQ support.....	101
3.7.1	Overview of QinQ.....	101
3.7.2	QinQ support with forced c-tag forwarding (VPLS only).....	103
3.7.2.1	Example 1: general QinQ implementation.....	103
3.7.2.2	Example 2: QinQ using VPLS with Ethernet SAPs.....	104
3.7.2.3	Example 3: QinQ using VPLS with ATM SAPs.....	104
3.7.3	QinQ support on Ethernet ports.....	105
3.7.3.1	Special QinQ SAP identifiers.....	105
3.7.3.2	QinQ dot1p match behavior.....	107
3.7.3.3	QinQ top-only mark option.....	107
3.7.3.4	Maximum number of VLAN tags.....	107

3.7.4	QinQ configuration overview.....	107
3.8	Raw socket IP transport service.....	107
3.8.1	Remote host manual TCP connection check.....	113
3.8.2	QoS requirements for IP transport.....	113
3.9	Service creation overview.....	113
3.10	Port and SAP CLI identifiers.....	114
3.11	Configuring global service entities with CLI.....	115
3.12	Service model entities.....	115
3.13	Basic configuration.....	116
3.14	Common configuration tasks.....	117
3.14.1	Configuring customer accounts.....	117
3.14.2	Configuring SDPs.....	118
3.14.2.1	Enabling IP fragmentation for GRE SDPs.....	122
3.14.3	Configuring service names.....	122
3.15	ETH-CFM (802.1ag and Y.1731) tasks.....	122
3.15.1	Configuring ETH-CFM parameters (802.1ag and Y.1731).....	123
3.15.2	Applying ETH-CFM parameters.....	125
3.16	Service management tasks.....	126
3.16.1	Modifying customer accounts.....	126
3.16.2	Deleting customers.....	127
3.16.3	Modifying SDPs.....	127
3.16.4	Deleting SDPs.....	128
3.16.5	Deleting LSP associations.....	128
3.17	Global service command reference.....	129
3.17.1	Command hierarchies.....	129
3.17.1.1	Global service configuration commands.....	129
3.17.1.2	Show commands.....	132
3.17.1.3	Monitor commands.....	132
3.17.2	Command descriptions.....	133
3.17.2.1	Global service configuration commands.....	133
3.17.2.2	Show commands.....	164
3.17.2.3	Monitor commands.....	179
4	VLL services.....	189
4.1	ATM VLL (Apipe) services.....	189
4.1.1	ATM VLL for end-to-end ATM service.....	189

4.1.2	ATM virtual trunk over an IP/MPLS packet-switched network.....	190
4.1.3	ATM SAP-to-SAP service.....	191
4.1.4	ATM traffic management support.....	192
4.1.4.1	Network ingress classification.....	192
4.1.4.2	ATM access egress queuing and shaping.....	192
4.1.5	Control word.....	193
4.2	Circuit emulation VLL (Cpipe) services.....	193
4.2.1	Cpipe service overview.....	193
4.2.2	TDM SAP-to-SAP service.....	194
4.2.3	Cpipe service modes.....	194
4.2.3.1	Unstructured mode (SAToP).....	194
4.2.3.2	SAToP serial.....	194
4.2.3.3	SAToP teleprotection interface.....	198
4.2.3.4	Structured mode (CESoPSN).....	198
4.2.4	TDM PW encapsulation.....	201
4.2.5	Circuit emulation parameters and options.....	204
4.2.5.1	Unstructured.....	206
4.2.5.2	Structured DS1/E1 CES without CAS.....	206
4.2.5.3	Structured T1/E1 CES with CAS.....	209
4.2.5.4	Packet payload size.....	211
4.2.5.5	Jitter buffer.....	211
4.2.5.6	Asymmetric delay control.....	212
4.2.5.7	Cpipe network latency measurement.....	217
4.2.5.8	RTP header.....	218
4.2.5.9	Control word.....	218
4.2.6	Transparent SDH/SONET over packet (TSoP).....	219
4.2.7	Error situations.....	219
4.3	Ethernet VLL (Epipe) services.....	220
4.3.1	Epipe service overview.....	220
4.3.2	Ethernet access egress queuing and scheduling.....	221
4.3.3	Ethernet SAP-to-SAP.....	222
4.3.4	Epipe with ATM SAPs.....	222
4.3.5	MEF 8.....	223
4.3.5.1	Epipe service modes.....	225
4.3.6	Ethernet OAM.....	227
4.3.7	Control word.....	227

4.3.8	MTU.....	227
4.3.9	Raw and tagged modes.....	228
4.3.9.1	Raw mode.....	228
4.3.9.2	Tagged mode.....	228
4.3.9.3	VLAN translation.....	228
4.3.9.4	Tagging rules for Epipe.....	229
4.3.10	IP filters (Epipe).....	234
4.3.11	MPLS entropy label.....	235
4.3.12	Security zones and Epipes.....	235
4.4	Frame relay VLL (Fpipe) services.....	237
4.4.1	Fpipe service overview.....	237
4.4.2	Frame relay SAP-to-SAP service.....	238
4.4.3	Frame relay traffic management.....	238
4.4.3.1	Ingress SAP classification and marking.....	238
4.4.4	Frame relay encapsulation.....	239
4.4.4.1	Frame relay PW control word.....	240
4.4.5	Status signaling and OAM propagation.....	241
4.5	HDLC VLL (Hpipe) services.....	241
4.5.1	Hpipe service overview.....	241
4.5.2	HDLC VLL for end-to-end HDLC service.....	242
4.5.3	HDLC SAP-to-SAP service.....	242
4.5.4	HDLC encapsulation.....	243
4.5.4.1	HDLC PW control word and payload size.....	243
4.5.5	Status signaling.....	244
4.5.5.1	OAM propagation.....	245
4.6	IP interworking VLL (Ipipe) services.....	245
4.6.1	Ipipe service overview.....	245
4.6.2	IP interworking VLL datapath.....	246
4.6.3	CE IP address discovery and distribution.....	247
4.6.3.1	Manual IP configuration.....	247
4.6.3.2	CE IP address discovery using ARP.....	248
4.6.3.3	CE IP address discovery for frame relay.....	248
4.6.3.4	CE IP address discovery for cHDLC.....	248
4.6.3.5	CE IP address discovery for PPP/MLPPP.....	248
4.6.3.6	CE IP address distribution to remote PE nodes.....	248
4.6.4	IP SAP-to-SAP service.....	249

4.6.5	Hardware support for interworking IP PWs.....	249
4.6.6	Control word.....	251
4.6.7	Termination at access.....	251
4.6.7.1	PPP and MLPPP termination.....	251
4.6.7.2	FR termination.....	251
4.6.7.3	cHDLC termination.....	252
4.6.7.4	PPP and MLPPP termination on 2-port OC3/STM1 Channelized Adapter cards.....	253
4.6.8	Traffic management.....	254
4.6.8.1	Traffic management for FR IW SAPs.....	254
4.6.9	Status signaling.....	254
4.6.9.1	FR LMI.....	254
4.6.10	IP filters (lpipe).....	254
4.7	Pseudowire switching.....	255
4.7.1	Overview.....	255
4.7.2	Pseudowire switching with pseudowire redundancy.....	258
4.7.3	Pseudowire switching behavior.....	258
4.7.3.1	Pseudowire switching with IP tunnels.....	260
4.7.4	Pseudowire switching TLV.....	260
4.8	VLL service considerations.....	261
4.8.1	Service support.....	262
4.8.2	SDPs.....	264
4.8.2.1	SDP statistics for VLL services.....	264
4.8.3	SAP encapsulations and pseudowire types.....	264
4.8.4	ATM PWE3 N-to-1 cell mode encapsulation.....	266
4.8.4.1	N-to-1 cell mode encapsulation (N = 1).....	267
4.8.4.2	N-to-1 cell mode encapsulation (N > 1).....	268
4.8.5	SAP aggregation groups.....	271
4.8.5.1	Configuration.....	271
4.8.5.2	QoS and traffic descriptor profiles (N > 1).....	272
4.8.5.3	Statistics and counters.....	273
4.8.5.4	Fault management.....	273
4.8.6	QoS policies.....	274
4.8.7	IP filter policies.....	274
4.8.8	MTU settings.....	275
4.8.8.1	Targeted LDP and MTU.....	280

4.8.9	QinQ (VLL service).....	280
4.8.10	Pseudowire control word.....	280
4.8.11	Pseudowire redundancy.....	281
4.8.11.1	PW redundancy operation.....	282
4.8.11.2	Selecting the active spoke SDP for PW redundancy configuration.....	284
4.8.11.3	PW redundancy and inter-chassis backup.....	285
4.8.11.4	AIS fault propagation.....	287
4.8.12	Active/standby mode for pseudowire redundancy (standby signaling).....	288
4.8.12.1	PW status signaling label withdrawal option.....	290
4.8.12.2	Pseudowire redundancy on serial data interface ports.....	291
4.9	Configuring a VLL service with CLI.....	292
4.10	Common configuration tasks.....	292
4.11	Configuring VLL components.....	292
4.11.1	Creating an Apipe service.....	293
4.11.1.1	Configuring Apipe SAP parameters.....	294
4.11.1.2	Configuring Apipe SDP bindings.....	295
4.11.1.3	Configuring Apipe SAP aggregation groups.....	296
4.11.1.4	Configuring Apipe SAPs as aggregation group members.....	297
4.11.2	Creating a Cpipe service.....	298
4.11.2.1	Configuring Cpipe SAP parameters.....	299
4.11.2.2	Configuring Cpipe SDP bindings.....	301
4.11.3	Creating an Epipe service.....	302
4.11.3.1	Configuring Epipe SAP parameters.....	302
4.11.3.2	Configuring an Epipe with an ATM SAP.....	304
4.11.3.3	Configuring Epipe SAP MEF 8 parameters.....	306
4.11.3.4	Configuring Epipe SAP microwave link parameters for interworking with TDM2Ethernet.....	308
4.11.3.5	Configuring ATM encapsulation under Epipe service (7705 SAR-M only).....	309
4.11.3.6	Configuring Epipe spoke SDP bindings.....	311
4.11.3.7	Configuring a security zone within an Epipe.....	313
4.11.4	Creating an Fpipe service.....	314
4.11.4.1	Configuring Fpipe SAP parameters.....	315
4.11.4.2	Configuring Fpipe SDP bindings.....	316
4.11.5	Creating an Hpipe service.....	317
4.11.5.1	Configuring Hpipe SAP parameters.....	318
4.11.5.2	Configuring Hpipe SDP bindings.....	320

4.11.6	Creating an lpipe service.....	320
4.11.6.1	Configuring lpipe SAP parameters.....	321
4.11.6.2	Configuring lpipe SDP bindings.....	322
4.11.7	Configuring PW switching.....	322
4.11.8	Configuring ingress and egress SAP parameters.....	324
4.11.9	Using the control word.....	325
4.11.10	Configuring PW redundancy.....	326
4.11.10.1	Configuring PW redundancy – standby signaling.....	327
4.11.10.2	Configuring PW redundancy – ICB.....	328
4.12	Service management tasks.....	330
4.12.1	Modifying service parameters.....	330
4.12.2	Disabling a service.....	332
4.12.3	Re-enabling a service.....	334
4.12.4	Deleting a service.....	334
4.13	VLL services command reference.....	336
4.13.1	Command hierarchies.....	336
4.13.1.1	VLL services configuration commands.....	336
4.13.1.2	Show commands.....	345
4.13.1.3	Clear commands.....	346
4.13.2	Command descriptions.....	347
4.13.2.1	VLL service configuration commands.....	347
4.13.2.2	Show commands.....	419
4.13.2.3	Clear commands.....	519
5	VPLS.....	523
5.1	VPLS overview.....	523
5.1.1	VPLS redundancy.....	524
5.1.2	Access control and traffic management.....	525
5.1.3	Split horizon.....	525
5.1.4	VPLS packet walkthrough.....	525
5.1.5	Bridged mobile backhaul.....	529
5.1.6	Multi-tenant unit termination.....	531
5.2	VPLS features.....	532
5.2.1	VPLS enhancements.....	533
5.2.2	Fabric mode.....	534
5.2.3	Subscriber VLAN.....	534

5.2.4	ATM encapsulated residential SAPs.....	535
5.2.5	VPLS over MPLS.....	535
5.2.6	VPLS MAC learning and packet forwarding.....	536
5.2.7	Pseudowire control word.....	536
5.2.8	Agent circuit ID insertion.....	536
5.2.9	MAC filters.....	537
5.2.10	FDB table management.....	538
5.2.10.1	FDB size.....	538
5.2.10.2	FDB size alarms.....	538
5.2.10.3	Local and remote aging timers.....	539
5.2.10.4	Unknown MAC discard.....	539
5.2.11	VPLS and rate limiting via QoS policy.....	539
5.2.12	MAC move.....	540
5.2.13	Split horizon groups (SAP and spoke SDP).....	541
5.2.13.1	Residential split horizon groups.....	541
5.2.14	Multicast for VPLS and routed VPLS (IGMP and MLD snooping).....	541
5.2.14.1	Application examples.....	542
5.2.14.2	Group and addressing support.....	543
5.2.14.3	IP multicast in r-VPLS.....	544
5.2.14.4	Multicast router ports.....	546
5.2.14.5	Tagged access traffic.....	546
5.2.14.6	Hardware support.....	546
5.2.15	PIM snooping for VPLS.....	546
5.2.15.1	Snooping versus proxy mode.....	548
5.2.16	MPLS entropy label.....	548
5.2.17	Ethernet OAM.....	548
5.2.18	Security zones and VPLS.....	549
5.3	Routed VPLS.....	550
5.3.1	IES or VPRN IP interface binding.....	551
5.3.1.1	Assigning a service name to a VPLS service.....	551
5.3.1.2	Binding a service name to an IP interface.....	551
5.3.1.3	Removing a bound VPLS service or service name.....	552
5.3.1.4	IP interface and VPLS operational state coordination.....	552
5.3.2	IP interface MTU and fragmentation.....	552
5.3.3	ARP/ND and VPLS FIB interactions.....	553
5.3.3.1	Routed VPLS specific ARP/ND cache behavior.....	553

5.3.4	The allow-ip-int-binding VPLS flag.....	554
5.3.4.1	VPLS feature restrictions with allow-ip-int-binding.....	554
5.3.5	DSCP marking.....	555
5.3.6	VPLS ingress IP filter override.....	555
5.3.7	Routed VPLS supported routing-related protocols.....	555
5.4	VPLS and spanning tree protocol.....	556
5.4.1	VPLS redundancy.....	556
5.4.1.1	Spoke SDP redundancy for metro interconnection.....	556
5.4.1.2	Spoke SDP-based redundant access.....	557
5.4.2	VPLS access redundancy.....	558
5.4.2.1	STP-based redundant access to VPLS.....	558
5.4.3	MAC flush message processing.....	558
5.4.3.1	Dual homing to a VPLS service.....	559
5.5	ATM PVC access and termination on a VPLS service.....	562
5.6	VPLS service considerations.....	563
5.6.1	SAP encapsulations.....	563
5.6.2	VLAN processing.....	563
5.6.2.1	Tagging rules for VPLS.....	564
5.6.3	QinQ (VPLS).....	564
5.7	Configuration notes.....	564
5.8	Configuring a VPLS service with CLI.....	564
5.9	Basic configuration.....	565
5.10	Common configuration tasks.....	566
5.11	Configuring VPLS components.....	566
5.11.1	Creating a VPLS service.....	567
5.11.2	Creating a split horizon group.....	567
5.11.3	Enabling MAC move.....	568
5.11.4	Configuring STP bridge parameters in a VPLS.....	569
5.11.4.1	Bridge STP admin state.....	569
5.11.4.2	Mode.....	569
5.11.4.3	Bridge priority.....	570
5.11.4.4	Hello time.....	570
5.11.4.5	Hold count.....	570
5.12	Configuring a VPLS SAP.....	570
5.12.1	Local VPLS SAPs.....	571
5.12.2	Distributed VPLS SAPs.....	571

5.12.3	Configuring SAP-specific STP parameters.....	572
5.12.3.1	SAP STP administrative state.....	572
5.12.3.2	SAP virtual port number.....	573
5.12.3.3	SAP priority.....	573
5.12.3.4	SAP path cost.....	574
5.12.3.5	SAP edge port.....	574
5.12.3.6	SAP auto edge.....	575
5.12.3.7	SAP link type.....	575
5.12.4	STP SAP operational states.....	575
5.12.4.1	Operationally disabled.....	576
5.12.4.2	Operationally discarding.....	576
5.12.4.3	Operationally learning.....	576
5.12.4.4	Operationally forwarding.....	576
5.12.5	Configuring VPLS SAPs with split horizon.....	576
5.13	Configuring SDP bindings.....	577
5.13.1	Configuring mesh SDP bindings.....	578
5.13.2	Configuring spoke SDPs.....	578
5.13.2.1	Configuring spoke SDP bindings.....	579
5.13.2.2	Configuring spoke SDP-specific STP parameters.....	580
5.13.3	Configuring VPLS spoke SDPs with split horizon.....	583
5.13.4	Configuring selective MAC flush.....	583
5.14	Configuring routed VPLS.....	584
5.15	Configuring IP multicast in VPLS.....	584
5.16	Configuring IP multicast in r-VPLS.....	585
5.17	Configuring multicast parameters for VPLS and r-VPLS.....	587
5.18	Configuring a static multicast group.....	588
5.19	Configuring PIM snooping for VPLS.....	589
5.20	Configuring a security zone within a VPLS.....	590
5.21	Service management tasks.....	591
5.21.1	Modifying VPLS service parameters.....	591
5.21.2	Modifying management VPLS parameters.....	592
5.21.3	Deleting a management VPLS.....	592
5.21.4	Disabling a management VPLS.....	592
5.21.5	Deleting a VPLS service.....	592
5.21.6	Disabling a VPLS service.....	593
5.21.7	Re-enabling a VPLS service.....	593

5.22	VPLS command reference.....	594
5.22.1	Command hierarchies.....	594
5.22.1.1	VPLS service configuration commands.....	594
5.22.1.2	Show commands.....	603
5.22.1.3	Clear commands.....	604
5.22.1.4	Debug commands.....	605
5.22.2	Command descriptions.....	606
5.22.2.1	VPLS service configuration commands.....	606
5.22.2.2	Routed VPLS configuration commands.....	688
5.22.2.3	VPLS security configuration commands.....	690
5.22.2.4	VPLS show commands.....	698
5.22.2.5	VPLS clear commands.....	794
5.22.2.6	VPLS debug commands.....	806
6	Internet enhanced service.....	818
6.1	IES for in-band management.....	818
6.1.1	Setting up connections between the NSP NFM-P and the 7705 SAR.....	819
6.1.2	Encapsulation.....	819
6.1.3	Layer 2 and Layer 3 traffic management.....	820
6.1.4	Troubleshooting and fault detection services.....	820
6.1.5	IP ECMP load balancing.....	820
6.2	IES for customer traffic.....	821
6.2.1	DHCP relay and DHCPv6 relay.....	823
6.2.1.1	DHCP relay.....	824
6.2.1.2	DHCPv6 relay.....	825
6.2.2	IPCP.....	825
6.2.3	IPSec support.....	826
6.2.4	Security zones and IES.....	826
6.2.5	Proxy ARP.....	826
6.2.6	Configurable ARP retry timer.....	827
6.2.7	Unnumbered interfaces.....	827
6.2.8	Troubleshooting and fault detection services.....	827
6.2.9	VRRP on IES interfaces.....	828
6.2.10	SAPs.....	828
6.2.10.1	Encapsulations.....	828
6.2.10.2	Routing protocols.....	829

6.2.10.3	QoS policies.....	829
6.2.10.4	QinQ (IES).....	829
6.2.10.5	IP filter policies on an IES SAP.....	829
6.2.11	Spoke-SDP termination to IES.....	829
6.2.12	Bandwidth optimization for low-speed links.....	831
6.2.12.1	OAM diagnostics restrictions with lower IP MTU.....	833
6.2.13	Hold up and hold down timers for IP interfaces.....	834
6.3	Configuring IES with CLI.....	835
6.4	Common configuration tasks.....	835
6.5	Configuring IES components.....	835
6.5.1	Creating an IES service.....	836
6.5.2	Configuring IES interface parameters.....	836
6.5.2.1	IES management service.....	837
6.5.2.2	IES service.....	837
6.5.2.3	IES IPv6 service.....	839
6.5.3	Configuring IES SAP parameters.....	840
6.5.3.1	IES management SAP.....	840
6.5.3.2	IES service SAP.....	842
6.5.4	Configuring IES spoke SDP parameters.....	842
6.5.5	Configuring VRRP.....	843
6.5.6	Configuring a security zone within IES.....	844
6.5.7	Configuring serial raw socket transport within IES.....	845
6.6	Service management tasks.....	846
6.6.1	Modifying IES service parameters.....	846
6.6.2	Disabling an IES service.....	846
6.6.3	Re-enabling an IES service.....	847
6.6.4	Deleting an IES service.....	847
6.7	IES command reference.....	848
6.7.1	Command hierarchies.....	848
6.7.1.1	Configuration commands.....	848
6.7.1.2	Show commands.....	854
6.7.1.3	Clear commands.....	855
6.7.1.4	Debug commands.....	855
6.7.2	Command descriptions.....	856
6.7.2.1	IES generic configuration commands.....	856
6.7.2.2	IES global configuration commands.....	858

6.7.2.3	IES management configuration commands.....	859
6.7.2.4	IES service configuration commands.....	870
6.7.2.5	Show commands.....	949
6.7.2.6	Clear commands.....	985
6.7.2.7	Debug commands.....	988
7	VPRN services.....	989
7.1	VPRN service overview.....	989
7.1.1	Routing prerequisites.....	991
7.1.2	BGP support.....	991
7.1.2.1	BGP fast reroute with prefix-independent convergence in a VPRN.....	992
7.1.2.2	BGP next-hop resolution and peer tracking.....	993
7.1.3	IPSec support.....	993
7.1.4	Security zones and VPRN.....	993
7.1.5	Static one-to-one NAT and VPRN.....	995
7.1.6	Unicast and multicast address translation.....	996
7.1.7	Route distinguishers.....	997
7.1.7.1	PE-to-CE route exchange.....	998
7.1.8	Route target constraint.....	1000
7.1.8.1	Configuring the route target address family.....	1000
7.1.8.2	Originating RTC routes.....	1000
7.1.8.3	Receiving and readvertising RTC routes.....	1001
7.1.8.4	Using RTC routes.....	1002
7.1.9	In-band management using a VPRN.....	1003
7.2	VPRN features.....	1006
7.2.1	IP interfaces.....	1007
7.2.1.1	Unnumbered interfaces.....	1007
7.2.1.2	DHCP and DHCPv6.....	1007
7.2.1.3	IPCP.....	1012
7.2.1.4	Troubleshooting and fault detection services.....	1013
7.2.1.5	VRRP on VPRN interfaces.....	1013
7.2.1.6	IP ECMP load balancing.....	1013
7.2.1.7	Proxy ARP.....	1014
7.2.1.8	Configurable ARP retry timer.....	1014
7.2.1.9	Hold up and hold down timers for IP interfaces.....	1014
7.2.2	SAPs.....	1015

7.2.2.1	Encapsulations.....	1015
7.2.2.2	QoS policies.....	1015
7.2.2.3	QinQ (VPRN).....	1016
7.2.2.4	Filter policies on a VPRN SAP.....	1016
7.2.3	PE-to-CE routing protocols.....	1016
7.2.3.1	Using OSPF or OSPFv3 in IP VPNs.....	1017
7.2.3.2	TTL security.....	1018
7.2.4	PE-to-PE tunneling mechanisms.....	1018
7.2.5	Per-VRF route limiting.....	1018
7.2.6	RIP metric propagation in VPRNs.....	1019
7.2.7	Multicast VPN (MVPN).....	1019
7.2.7.1	Multicast in IP-VPN applications.....	1020
7.2.7.2	MVPN building blocks.....	1021
7.2.7.3	Provider tunnel support.....	1028
7.2.7.4	Inter-AS and intra-AS solutions.....	1030
7.2.7.5	NG-MVPN non-segmented inter-AS solution.....	1032
7.2.7.6	Mrinfo and Mtrace.....	1036
7.2.7.7	Multicast-only fast reroute.....	1036
7.2.7.8	mLDP point-to-multipoint support.....	1036
7.2.7.9	mLDP fast upstream switchover.....	1037
7.2.7.10	Multicast Source Discovery Protocol.....	1037
7.2.8	VPRN autobinding tunnels.....	1037
7.2.9	Spoke SDPs.....	1038
7.2.10	Spoke-SDP termination to VPRN.....	1038
7.2.11	IPv6 on virtual private edge router.....	1041
7.2.12	IPv6 over IPv4 LAN-to-LAN IPsec tunnels.....	1043
7.2.13	Bandwidth optimization for low-speed links.....	1043
7.2.13.1	OAM diagnostics restrictions with lower IP MTU.....	1045
7.2.14	Support for NTP.....	1046
7.3	Configuring a VPRN service with CLI.....	1046
7.4	Basic configuration.....	1047
7.5	Common configuration tasks.....	1048
7.6	Configuring VPRN components.....	1048
7.6.1	Creating a VPRN service.....	1049
7.6.2	Configuring global VPRN parameters.....	1049
7.6.3	Configuring router interfaces.....	1050

7.6.4	Configuring static route entries for VPRN.....	1050
7.6.5	Configuring BGP for VPRN.....	1052
7.6.5.1	Configuring VPRN BGP group and neighbor parameters.....	1052
7.6.5.2	Configuring route reflection.....	1053
7.6.5.3	VPRN BGP CLI syntax.....	1053
7.6.6	Configuring IPv6 parameters for VPRN BGP.....	1054
7.6.7	Configuring VPRN IPv6 neighbor discovery parameters.....	1055
7.6.8	Configuring OSPF or OSPFv3 for VPRN.....	1055
7.6.9	Configuring RIP for VPRN.....	1056
7.6.10	Configuring IGMP for VPRN.....	1057
7.6.11	Configuring PIM for VPRN.....	1058
7.6.12	Configuring MVPN for VPRN.....	1059
7.6.13	Configuring a VPRN interface.....	1060
7.6.14	Configuring a VPRN IPv6 interface.....	1063
7.6.15	Configuring VPRN interface routed VPLS IPv6 parameters.....	1064
7.6.16	Configuring VPRN interface SAP parameters.....	1064
7.6.17	Configuring VPRN interface SAP IPv6 parameters.....	1067
7.6.18	Configuring VPRN interface spoke SDP parameters.....	1067
7.6.19	Configuring VPRN interface spoke SDP IPv6 parameters.....	1068
7.6.20	Configuring VRRP.....	1069
7.6.21	Configuring a security zone within a VPRN.....	1069
7.6.22	Configuring serial raw socket transport within a VPRN.....	1070
7.6.23	Configuring VPRN router advertisement.....	1071
7.7	Service management tasks.....	1072
7.7.1	Modifying VPRN service parameters.....	1072
7.7.2	Deleting a VPRN service.....	1073
7.7.3	Disabling a VPRN service.....	1073
7.7.4	Re-enabling a VPRN service.....	1074
7.8	VPRN services command reference.....	1075
7.8.1	Command hierarchies.....	1075
7.8.1.1	Configuration commands.....	1076
7.8.1.2	Show commands.....	1101
7.8.1.3	Clear commands.....	1104
7.8.1.4	Debug commands.....	1105
7.8.2	Command descriptions.....	1106
7.8.2.1	VPRN service configuration commands.....	1106

7.8.2.2	Show service commands.....	1394
7.8.2.3	Show router commands.....	1454
7.8.2.4	Clear service commands.....	1550
7.8.2.5	Debug MSDP commands.....	1561
7.8.2.6	Debug service commands.....	1563
8	IPSec.....	1569
8.1	IPSec overview.....	1569
8.1.1	IPSec implementation.....	1570
8.1.1.1	IPSec overview.....	1570
8.1.1.2	Hardware support.....	1571
8.1.1.3	IPSec encryption features.....	1572
8.1.1.4	SHA2 support.....	1573
8.1.1.5	IPSec security policy, IKE policy, and IPSec transform.....	1573
8.1.1.6	IKEv2 fragmentation.....	1573
8.1.1.7	Tunnel group.....	1574
8.1.1.8	Tunnel interfaces and SAPs.....	1574
8.1.1.9	IPSec tunnel configuration.....	1575
8.1.1.10	IPSec over MPLS with public-side IES.....	1575
8.1.1.11	IPSec transport tunnels with public-side VPRN.....	1575
8.1.1.12	GRE-encapsulated VLLs/VPLS over IPSec VPNs.....	1576
8.1.1.13	GRE-encapsulated VLLs/VPLS over IPSec over MPLS.....	1579
8.1.2	X.509v3 certificate overview.....	1583
8.1.2.1	X.509v3 certificate support on the 7705 SAR.....	1584
8.1.2.2	Local storage.....	1584
8.1.2.3	CA profile.....	1585
8.1.2.4	CA chain computation.....	1586
8.1.2.5	Certificate enrollment.....	1586
8.1.2.6	Certificate revocation check.....	1587
8.1.2.7	Certificate, CRL, and key cache.....	1587
8.1.3	Using certificates for IPSec tunnel authentication.....	1588
8.1.4	Trust anchor profile.....	1588
8.1.5	Certificate profile.....	1588
8.1.6	Certificate Management Protocol version 2.....	1589
8.1.7	OCSP.....	1590
8.1.8	Applications.....	1591

8.1.8.1	Metrocell deployment.....	1591
8.1.8.2	Small business deployment.....	1591
8.1.9	NAT-traversal for IKEv1/v2 and IPSec.....	1592
8.1.10	BFD over IPSec tunnel.....	1593
8.1.11	QoS for IPSec.....	1594
8.1.11.1	Network and access ingress QoS (decryption QoS).....	1594
8.1.11.2	Network and access egress QoS (encryption QoS).....	1595
8.1.12	Fragmentation and IP MTU.....	1595
8.1.12.1	Fragmentation configuration.....	1595
8.1.12.2	Reassembly.....	1596
8.1.13	Support for private VPRN service features.....	1596
8.1.14	Routing in private services.....	1596
8.1.15	IPSec on the 10-port 1GigE/1-port 10GigE X-Adapter card.....	1596
8.1.16	IPSec sequence number.....	1596
8.1.17	PBR and MFC.....	1597
8.1.17.1	PBR.....	1597
8.1.17.2	MFC.....	1597
8.1.18	OSPFv3 packet authentication with IPv6 IPSec.....	1597
8.1.19	Network security with IPv6 IPSec.....	1598
8.1.20	IPSec over r-VPLS on a public-side service.....	1599
8.1.21	Statistics.....	1599
8.1.22	Security support.....	1599
8.2	Public key infrastructure.....	1600
8.2.1	CA role in PKI.....	1600
8.2.2	Digital signature and certificates.....	1600
8.2.2.1	Certificates.....	1601
8.2.3	Certificate storage.....	1602
8.2.4	CMPv2 certificate management.....	1603
8.2.4.1	CMPv2 initial registration.....	1603
8.2.4.2	Key update.....	1604
8.2.4.3	CRL.....	1604
8.2.5	OCSP.....	1604
8.2.6	Certificate or CRL expiration warning.....	1604
8.2.7	Automatic CRL update.....	1605
8.3	IPSec best practices recommendations.....	1605
8.4	Configuration notes.....	1606

8.5	Configuring IPsec with CLI.....	1606
8.6	Basic configuration overview.....	1607
8.7	Common configuration tasks.....	1607
8.7.1	Configuring an IPsec tunnel group.....	1607
8.7.2	Configuring router interfaces for IPsec.....	1608
8.7.3	Configuring IPsec parameters.....	1608
8.7.4	Configuring IPsec and IPsec tunnels in services.....	1609
8.7.5	Configuring IPsec IPv6 parameters for a VPRN private service.....	1611
8.7.6	Configuring X.509v3 certificate parameters.....	1611
8.7.7	Configuring CMPv2.....	1613
8.7.8	Configuring OCSP.....	1614
8.8	Configuring IPsec over MPLS.....	1614
8.8.1	IPsec over BGP 3107 label routes.....	1615
8.8.1.1	Static LAN-to-LAN tunnel configuration.....	1615
8.8.1.2	Policy option configuration.....	1617
8.8.1.3	BGP configuration with BGP 3107 label route advertisement.....	1618
8.8.1.4	LDP or RSVP-TE tunnel configuration.....	1618
8.8.2	IPsec over IGP shortcut.....	1620
8.8.2.1	Static LAN-to-LAN tunnel configuration.....	1620
8.8.2.2	IGP shortcut configuration.....	1622
8.8.2.3	RSVP-TE LSP configuration.....	1622
8.9	Service management tasks.....	1623
8.9.1	Deleting an IPsec IKE policy or an IPsec transform.....	1623
8.9.2	Deleting a public-side IPsec tunnel SAP and interface.....	1624
8.9.3	Deleting a private-side IPsec tunnel SAP and interface.....	1624
8.9.4	Deleting an IPsec security policy.....	1624
8.9.5	Deleting an IPsec tunnel.....	1625
8.10	IPsec command reference.....	1626
8.10.1	Command hierarchies.....	1626
8.10.1.1	IPsec configuration commands.....	1626
8.10.1.2	PKI configuration commands.....	1629
8.10.1.3	Show commands.....	1631
8.10.1.4	Clear commands.....	1632
8.10.1.5	Debug commands.....	1632
8.10.2	Command descriptions.....	1633
8.10.2.1	IPsec configuration commands.....	1633

8.10.2.2	PKI configuration commands.....	1674
8.10.2.3	Show commands.....	1716
8.10.2.4	Clear commands.....	1732
8.10.2.5	Debug commands.....	1732
9	Network group encryption.....	1735
9.1	NGE overview.....	1735
9.1.1	NGE key groups and encryption partitions.....	1737
9.1.2	NGE domains.....	1738
9.1.2.1	Private IP/MPLS network NGE domain.....	1740
9.1.2.2	Private over intermediary network NGE domain.....	1740
9.1.3	Network Services Platform management.....	1741
9.2	Key groups.....	1742
9.2.1	Key group algorithms.....	1743
9.2.1.1	Encapsulating security payload.....	1744
9.2.2	Security associations.....	1744
9.2.2.1	Active outbound SA.....	1744
9.3	Services encryption.....	1745
9.3.1	Services encryption overview.....	1745
9.3.2	Assigning key groups to services.....	1747
9.3.3	Pseudowire switching for NGE traffic.....	1748
9.3.4	Pseudowire control word for NGE traffic.....	1749
9.3.5	VPRN Layer 3 spoke-SDP encryption and MP-BGP-based VPRN encryption interaction.....	1749
9.3.6	NGE and RFC 3107.....	1749
9.3.7	NGE for NG-MVPN.....	1750
9.4	Router interface encryption.....	1750
9.4.1	Router interface NGE domain concepts.....	1751
9.4.2	GRE-MPLS packets inside the NGE domain.....	1753
9.4.3	Router encryption exceptions using ACLs.....	1753
9.4.4	IPSec packets crossing an NGE domain.....	1754
9.4.5	Multicast packets traversing the NGE domain.....	1755
9.4.6	Assigning key groups to router interfaces.....	1757
9.4.7	Router interface NGE firewall considerations.....	1757
9.4.8	NGE and BFD support.....	1758
9.4.9	NGE and ACL interactions.....	1758

9.4.10	Router interface NGE and ICMP interactions over the NGE domain.....	1758
9.4.11	OAM considerations for router interface encryption.....	1759
9.5	Layer 2 encryption.....	1759
9.6	NGE packet overhead and MTU considerations.....	1760
9.6.1	GRE fragmentation for NGE packets.....	1762
9.7	1588v2 encryption with NGE.....	1762
9.8	QoS for NGE traffic.....	1763
9.8.1	Network ingress.....	1763
9.8.2	Network egress.....	1764
9.9	Statistics.....	1764
9.10	Remote network monitoring support.....	1764
9.11	Configuration notes.....	1765
9.12	Configuring NGE with CLI.....	1766
9.13	Basic NGE configuration overview.....	1766
9.14	Configuring NGE components.....	1766
9.14.1	Configuring the global encryption label.....	1767
9.14.2	Configuring a key group.....	1767
9.14.3	Assigning a key group to an SDP or VPRN service.....	1768
9.14.4	Assigning a key group to a router interface.....	1770
9.14.5	Assigning a key group to an Ethernet port.....	1770
9.15	NGE management tasks.....	1771
9.15.1	Modifying a key group.....	1771
9.15.2	Removing a key group.....	1773
9.15.2.1	Removing a key group from an SDP or VPRN service.....	1773
9.15.2.2	Removing a key group from a router interface.....	1774
9.15.2.3	Removing a key group from an Ethernet port.....	1774
9.15.3	Changing key groups.....	1775
9.15.3.1	Changing the key group for an SDP or VPRN service.....	1775
9.15.3.2	Changing the key group for a router interface.....	1776
9.15.3.3	Changing the key group for an Ethernet port.....	1776
9.15.4	Deleting a key group from a 7705 SAR.....	1777
9.16	NGE command reference.....	1778
9.16.1	Command hierarchies.....	1778
9.16.1.1	Configuration commands.....	1778
9.16.1.2	Show commands.....	1779
9.16.1.3	Clear commands.....	1779

9.16.2	Command descriptions.....	1780
9.16.2.1	Configuration commands.....	1780
9.16.2.2	Show commands.....	1786
9.16.2.3	Clear commands.....	1791
10	Ethernet virtual private networks.....	1793
10.1	Overview and EVPN applications.....	1793
10.1.1	EVPN for MPLS tunnels in E-LAN services.....	1793
10.1.2	EVPN for MPLS tunnels in E-Line services.....	1794
10.2	EVPN for MPLS tunnels.....	1795
10.2.1	BGP-EVPN control plane for MPLS tunnels.....	1795
10.2.1.1	EVPN route type 2 – MAC/IP advertisement route.....	1797
10.2.1.2	EVPN route type 3 – inclusive multicast Ethernet tag route.....	1798
10.2.1.3	EVPN route type 1 – Ethernet auto-discovery route.....	1799
10.2.1.4	EVPN route type 4 – Ethernet segment route.....	1800
10.2.1.5	EVPN route type 5 – IP prefix route.....	1801
10.2.1.6	RFC 5512 – BGP tunnel encapsulation extended community.....	1802
10.2.2	EVPN-VPLS for MPLS tunnels.....	1803
10.2.2.1	Overview.....	1803
10.2.2.2	EVPN and VPLS integration.....	1805
10.2.2.3	Auto-derived route distinguisher in services.....	1808
10.2.2.4	EVPN multihoming in VPLS services.....	1809
10.2.3	EVPN-VPWS for MPLS tunnels.....	1830
10.2.3.1	BGP-EVPN control plane for EVPN-VPWS.....	1830
10.2.3.2	EVPN for MPLS tunnels in Epipe services.....	1831
10.2.3.3	Using active/standby PWs and MC-LAG with EVPN-VPWS Epipes.....	1834
10.2.3.4	EVPN multihoming for EVPN-VPWS services.....	1835
10.2.4	EVPN for MPLS tunnels in r-VPLS services.....	1837
10.2.4.1	Overview.....	1837
10.2.4.2	EVPN-MPLS multihoming and passive VRRP.....	1838
10.2.5	MPLS entropy label.....	1839
10.2.6	Preference-based and non-revertive designated forwarder election.....	1840
10.2.7	IPv6 tunnel resolution for EVPN MPLS Services.....	1842
10.2.8	EVPN multi-homing support for MPLS tunnels resolved to non-system IPv4/IPv6 addresses.....	1843
10.3	General EVPN topics.....	1843

10.3.1	BGP-EVPN MAC mobility.....	1843
10.3.2	BGP-EVPN MAC duplication.....	1844
10.3.3	Conditional static MAC and protection.....	1845
10.3.4	Blackhole MAC.....	1846
10.3.5	CFM interaction with EVPN services.....	1847
10.3.6	BGP and EVPN route selection for EVPN routes.....	1848
10.3.7	Interaction of EVPN and other features.....	1848
10.3.7.1	Interaction of EVPN-MPLS with existing VPLS features.....	1848
10.3.7.2	Routing policies for BGP-EVPN IP prefixes.....	1849
10.4	Configuring an EVPN service with CLI.....	1851
10.4.1	EVPN all-active multihoming configuration example.....	1852
10.4.2	EVPN single-active multihoming configuration example.....	1854
10.4.3	EVPN-MPLS r-VPLS configuration examples.....	1855
10.4.3.1	EVPN-MPLS r-VPLS without multihoming.....	1855
10.4.3.2	EVPN-MPLS r-VPLS with all-active multihoming.....	1860
10.4.3.3	EVPN-MPLS r-VPLS with single-active multihoming.....	1869
10.5	EVPN command reference.....	1874
10.5.1	Command hierarchies.....	1874
10.5.1.1	EVPN configuration commands.....	1874
10.5.1.2	Show commands.....	1877
10.5.2	Command descriptions.....	1879
10.5.2.1	EVPN configuration commands.....	1879
10.5.2.2	Routed VPLS EVPN commands.....	1900
10.5.2.3	EVPN service system commands.....	1900
10.5.2.4	EVPN redundancy commands.....	1911
10.5.2.5	Show commands.....	1913
11	List of acronyms.....	1933
12	Supported standards and protocols.....	1960
12.1	Security standards.....	1960
12.2	Telecom standards.....	1960
12.3	Protocol support.....	1961
12.3.1	ATM.....	1961
12.3.2	BFD.....	1961
12.3.3	BGP.....	1962

12.3.4	DHCP/DHCPv6.....	1962
12.3.5	Differentiated services.....	1963
12.3.6	Digital data network management.....	1963
12.3.7	ECMP.....	1963
12.3.8	Ethernet VPN (EVPN).....	1963
12.3.9	Frame relay.....	1963
12.3.10	GRE.....	1963
12.3.11	Internet protocol (IP) – version 4.....	1964
12.3.12	Internet protocol (IP) – version 6.....	1964
12.3.13	IPSec.....	1964
12.3.14	IS-IS.....	1965
12.3.15	LDP.....	1966
12.3.16	LDP and IP FRR.....	1966
12.3.17	MPLS.....	1966
12.3.18	MPLS – OAM.....	1967
12.3.19	Multicast.....	1967
12.3.20	Network management.....	1967
12.3.21	OSPF.....	1969
12.3.22	OSPFv3.....	1969
12.3.23	PPP.....	1969
12.3.24	Pseudowires.....	1969
12.3.25	RIP.....	1970
12.3.26	RADIUS.....	1970
12.3.27	RSVP-TE and FRR.....	1970
12.3.28	Segment routing (SR).....	1971
12.3.29	SONET/SDH.....	1971
12.3.30	SSH.....	1971
12.3.31	Synchronization.....	1971
12.3.32	TACACS+.....	1972
12.3.33	TLS.....	1972
12.3.34	TWAMP.....	1972
12.3.35	VPLS.....	1973
12.3.36	VRRP.....	1973
12.4	Proprietary MIBs.....	1973

List of tables

Table 1: Configuration process.....	46
Table 2: Pseudowire service types.....	49
Table 3: Service types and SAP encapsulations.....	56
Table 4: GRE header descriptions.....	61
Table 5: GRE pseudowire payload packet descriptions.....	62
Table 6: IP pseudowire payload packet descriptions.....	65
Table 7: Spoke-SDP termination support.....	66
Table 8: 802.1ag terminology.....	72
Table 9: Y.1731 terminology.....	73
Table 10: CFM frame processing.....	77
Table 11: FC and VLAN priority mappings for Up and Down MEPs per frame type.....	88
Table 12: Special QinQ SAP identifiers.....	105
Table 13: SDP echo reply response conditions.....	149
Table 14: Ethernet ring command field descriptions.....	166
Table 15: Customer command field descriptions.....	169
Table 16: Service SDP field descriptions.....	171
Table 17: Service SDP-using field descriptions.....	175
Table 18: Service service-using field descriptions.....	178
Table 19: SAP ID configurations.....	180
Table 20: Port and encapsulation values.....	184
Table 21: SAToP serial payload size minimums and defaults	195

Table 22: T1 framing for CAS (RBS) support in a T1 ESF multiframe.....	200
Table 23: SAToP and CESoPSN support on the 7705 SAR.....	204
Table 24: Unstructured payload defaults.....	206
Table 25: Default and minimum payload sizes for CESoPSN without CAS.....	208
Table 26: Default values for the payload size for T1 and E1 CESoPSN with CAS.....	210
Table 27: Control word bit descriptions.....	218
Table 28: MEF 8 support on the 7705 SAR.....	224
Table 29: Ingress SAP tagging rules.....	229
Table 30: Egress SAP tagging rules.....	230
Table 31: VLAN tagging examples (Epipe)	230
Table 32: Security zone interfaces and endpoints per context.....	236
Table 33: Control word bit descriptions.....	240
Table 34: Control word bit descriptions.....	244
Table 35: Hardware support for interworking PWs.....	249
Table 36: Supported pseudowire switching options.....	257
Table 37: Pseudowire switching TLV field descriptions.....	261
Table 38: MTU points and descriptions.....	275
Table 39: Service MTU default and maximum values.....	277
Table 40: MTU calculator – service creation (worst case) access ports and SAPs.....	277
Table 41: MTU calculator – service creation (worst case) network ports.....	278
Table 42: Matching MTU or payload values for signaled VLL services.....	280
Table 43: Service MTU values.....	363
Table 44: SAP ID configurations.....	366

Table 45: Port and encapsulation values.....	369
Table 46: SAP ID preconfiguration for SAP aggregation groups.....	375
Table 47: Match-qinq-dot1p matching behavior.....	387
Table 48: Dot1p re-marking behavior for the qinq-mark-top-only command.....	389
Table 49: Service-ID all command field descriptions.....	422
Table 50: Service-ID base field descriptions.....	466
Table 51: Service egress label field descriptions.....	469
Table 52: Service ingress label field descriptions.....	470
Table 53: Service-ID endpoint field descriptions.....	472
Table 54: Service-ID labels field descriptions.....	473
Table 55: Service-ID MACsec field descriptions.....	474
Table 56: Service-ID network latency measurement field descriptions.....	475
Table 57: Service-ID SAP field descriptions.....	479
Table 58: Service-ID SAP aggregation group field descriptions.....	501
Table 59: Service SAP-using field descriptions.....	508
Table 60: Service SCADA bridge SAP field descriptions.....	509
Table 61: SDP field descriptions.....	511
Table 62: SDP-using field descriptions.....	517
Table 63: Service-using field descriptions.....	518
Table 64: Security zone interfaces and endpoints per context.....	549
Table 65: Ingress behavior for VPLS next-hop routing.....	554
Table 66: Egress behavior for VPLS next-hop routing.....	554
Table 67: Service service-MTU field descriptions.....	631

Table 68: Match-qinq-dot1p matching behavior.....	667
Table 69: Dot1p re-marking behavior for the qinq-mark-top-only command.....	669
Table 70: Service egress labels field descriptions.....	699
Table 71: Service FDB-info field descriptions.....	700
Table 72: Service FDB-MAC field descriptions.....	703
Table 73: Service service-ID (all) field descriptions.....	711
Table 74: Service service-ID (base) field descriptions.....	714
Table 75: Service service-ID (DHCP statistics) field descriptions.....	716
Table 76: Service service-ID (DHCP summary) field descriptions.....	718
Table 77: Service service-ID (endpoint) field descriptions.....	719
Table 78: Service service-ID (FDB) field descriptions.....	722
Table 79: Service service-ID (IGMP and MLD snooping) field descriptions.....	730
Table 80: Service service-ID IGMP and MLD snooping (base) field descriptions.....	738
Table 81: Service service-ID IGMP and MLD snooping (port-DB) field descriptions.....	742
Table 82: Service service-ID IGMP and MLD snooping (proxy-DB) field descriptions.....	746
Table 83: Service service-ID IGMP and MLD snooping (querier) field descriptions.....	747
Table 84: Service service-ID IGMP and MLD snooping (static) field descriptions.....	749
Table 85: Service service-ID IGMP and MLD snooping (statistics) field descriptions.....	752
Table 86: Service service-ID (labels) field descriptions.....	755
Table 87: Service service-ID (MAC move) field descriptions.....	756
Table 88: Service-ID MACsec field descriptions.....	758
Table 89: Service ID (SAP) field descriptions.....	763
Table 90: Service ID (SDP) field descriptions.....	772

Table 91: Service ID (split horizon group) field descriptions.....	777
Table 92: Service ID (STP) field descriptions.....	779
Table 93: Service ingress-label field descriptions.....	784
Table 94: Service VPLS PIM snooping group field descriptions.....	785
Table 95: Service VPLS PIM snooping neighbor field descriptions.....	787
Table 96: Service VPLS PIM snooping port field descriptions.....	788
Table 97: Service VPLS PIM snooping statistics field descriptions.....	789
Table 98: Service VPLS PIM snooping status field descriptions.....	791
Table 99: Service SAP-using field descriptions.....	794
Table 100: Security zone interfaces and endpoints for IES.....	826
Table 101: Port MTU requirements for OAM diagnostics (GRE tunnels).....	833
Table 102: Port MTU requirements for OAM diagnostics (LDP tunnels).....	834
Table 103: Message interval configuration ranges.....	912
Table 104: Match-qinq-dot1p matching behavior.....	923
Table 105: Dot1p re-marking behavior for the qinq-mark-top-only command.....	925
Table 106: Valid DSCP names.....	944
Table 107: Service customer field descriptions.....	950
Table 108: Service egress field descriptions.....	951
Table 109: Service ID all field descriptions.....	954
Table 110: Service ID ARP field descriptions.....	961
Table 111: Service ID base field descriptions.....	962
Table 112: Service ID DHCP statistics field descriptions.....	964
Table 113: Service ID DHCP summary field descriptions.....	965

Table 114: Service ID interface field descriptions.....	967
Table 115: Service IP transport summary field descriptions.....	968
Table 116: Service IP transport detailed field descriptions.....	970
Table 117: IP transport subservice remote host summary field descriptions.....	973
Table 118: IP transport subservice remote host detailed field descriptions.....	974
Table 119: Service-ID MACsec field descriptions.....	976
Table 120: Service ingress label field descriptions.....	979
Table 121: IP transport-using field descriptions.....	980
Table 122: Service SAP-using field descriptions.....	984
Table 123: Service service-using field descriptions.....	985
Table 124: BGP FRR scenarios	992
Table 125: Security zone interfaces and endpoints per context.....	994
Table 126: VPRN interfaces supported for static one-to-one NAT.....	996
Table 127: Route distinguisher Type-Value fields.....	997
Table 128: IPv4 and IPv6 GRT-supported management protocols.....	1004
Table 129: Recursive opaque types.....	1033
Table 130: 6VPE access control list, SAP	1042
Table 131: 6VPE access control list, SDP	1042
Table 132: 6VPE access control list, r-VPLS override	1043
Table 133: Port MTU requirements for OAM diagnostics (GRE tunnels).....	1045
Table 134: Port MTU requirements for OAM diagnostics (LDP tunnels).....	1046
Table 135: Applications and support for configurable DSCP or dot1p markings.....	1128
Table 136: DSCP-to-default forwarding class mapping.....	1131

Table 137: Default route preference.....	1140
Table 138: Route preference defaults by route type.....	1211
Table 139: Valid DSCP names.....	1280
Table 140: VPRN interface state and IP address.....	1316
Table 141: Match-qinq-dot1p matching behavior.....	1363
Table 142: Dot1p re-marking behavior for the qinq-mark-top-only command.....	1365
Table 143: Message interval configuration ranges.....	1380
Table 144: Service egress field descriptions.....	1395
Table 145: Service ID all field descriptions.....	1406
Table 146: Service ID ARP field descriptions.....	1417
Table 147: Service ID base field descriptions.....	1419
Table 148: Service ID DHCP statistics field descriptions.....	1422
Table 149: Service ID DHCP summary field descriptions.....	1423
Table 150: Service ID interface detailed field descriptions.....	1426
Table 151: Service IP transport subservice summary field descriptions.....	1429
Table 152: Service IP transport subservice detailed field descriptions.....	1430
Table 153: IP transport subservice remote host summary field descriptions.....	1433
Table 154: IP transport subservice remote host detailed field descriptions.....	1435
Table 155: Service-ID MACsec field descriptions.....	1436
Table 156: Service ID SAP detailed field descriptions.....	1441
Table 157: Service ingress label field descriptions.....	1446
Table 158: IP transport-using field descriptions.....	1447
Table 159: Service ID SDP detailed field descriptions.....	1449

Table 160: TWAMP Light field descriptions.....	1452
Table 161: Service service-using field descriptions.....	1454
Table 162: Aggregate route field descriptions.....	1456
Table 163: ARP table field descriptions.....	1457
Table 164: BGP damping field descriptions.....	1462
Table 165: BGP group field descriptions.....	1465
Table 166: BGP neighbor (standard, detailed, and dynamic) field descriptions.....	1473
Table 167: BGP neighbor (advertised-routes and received-routes) field descriptions.....	1482
Table 168: BGP neighbor (graceful restart) field descriptions.....	1484
Table 169: BGP next-hop field descriptions.....	1486
Table 170: BGP path field descriptions.....	1488
Table 171: BGP routes field descriptions.....	1508
Table 172: BGP summary field descriptions.....	1513
Table 173: DHCP or DHCPv6 server field descriptions.....	1515
Table 174: DHCP statistics field descriptions.....	1516
Table 175: DHCPv6 statistics field descriptions.....	1517
Table 176: DHCP summary field descriptions.....	1518
Table 177: DHCPv6 summary field descriptions.....	1519
Table 178: IP interface field descriptions.....	1521
Table 179: MSDP group field descriptions.....	1523
Table 180: MSDP peer field descriptions.....	1525
Table 181: MSDP source field descriptions.....	1526
Table 182: MSDP source-active field descriptions.....	1528

Table 183: MSDP source-active-rejected field descriptions.....	1529
Table 184: MSDP statistics field descriptions.....	1530
Table 185: MSDP status field descriptions.....	1532
Table 186: Standard route table field descriptions.....	1535
Table 187: Route table extensive field descriptions.....	1536
Table 188: Route table alternative field descriptions.....	1537
Table 189: Application QoS field descriptions.....	1540
Table 190: DSCP-to-FC mapping field descriptions.....	1542
Table 191: Static ARP table field descriptions.....	1543
Table 192: Static route field descriptions.....	1546
Table 193: Tunnel table field descriptions.....	1550
Table 194: IPSec IKE policy field descriptions.....	1722
Table 195: IPSec security policy field descriptions.....	1725
Table 196: IPSec transform field descriptions.....	1726
Table 197: IPSec tunnel field descriptions.....	1729
Table 198: Inside and outside NGE domains – configuration scenarios.....	1752
Table 199: NGE overhead for MPLS.....	1760
Table 200: NGE overhead for router interface and Ethernet port NGE.....	1761
Table 201: Accounting for NGE overhead SDP and service MTU – calculation examples.....	1761
Table 202: Encryption key group field descriptions.....	1788
Table 203: Group encryption summary field descriptions.....	1791
Table 204: EVPN route types and usage.....	1795
Table 205: Route type 2 fields and values	1797

Table 206: Route type 3 fields and values	1798
Table 207: Route type 1 fields and values (Ethernet AD per-ESI route)	1799
Table 208: Route type 1 fields and values (Ethernet AD per-EVI route)	1800
Table 209: Route type 4 fields and values	1801
Table 210: Route type 5 fields and values	1802
Table 211: Service EVPN-MPLS field descriptions.....	1915
Table 212: Service ID BGP field descriptions.....	1916
Table 213: Service BGP-EVPN field descriptions.....	1918
Table 214: Service ID Ethernet segment field descriptions.....	1920
Table 215: Service ID EVPN-MPLS field descriptions.....	1923
Table 216: Service system BGP-EVPN field descriptions.....	1925
Table 217: Service system BGP-EVPN Ethernet segment name field descriptions.....	1927
Table 218: Service system BGP route distinguisher field descriptions.....	1930
Table 219: Redundancy BGP-EVPN multihoming field descriptions.....	1932
Table 220: Acronyms.....	1933

List of figures

Figure 1: Service entities and the service model.....	52
Figure 2: Service access point (SAP).....	54
Figure 3: Multiple SAPs on a single port/channel.....	55
Figure 4: SDP tunnel pointing from ALU-A to ALU-B.....	59
Figure 5: GRE header.....	61
Figure 6: GRE pseudowire payload packet over Ethernet.....	62
Figure 7: IP example of pseudowire payload packet over Ethernet.....	65
Figure 8: 7705 SAR-A HSDPA offload example.....	70
Figure 9: MEPs and MAs.....	74
Figure 10: ETH-CFM frame format.....	75
Figure 11: ETH-CFM OAMPDU message.....	76
Figure 12: ICC-based MEG-ID format.....	78
Figure 13: Dot1ag loopback test.....	80
Figure 14: MEP on Ethernet access.....	84
Figure 15: Down MEP at Ethernet SAP.....	85
Figure 16: Dot1ag Down MEPs on spoke SDPs.....	86
Figure 17: Y.1731 MEP support on the 7705 SAR.....	87
Figure 18: Ethernet protection switching.....	90
Figure 19: G.8032 ring in the idle state.....	91
Figure 20: G.8032 ring in the protecting state.....	92
Figure 21: Ring example.....	93

Figure 22: G.8032 sub-ring.....	95
Figure 23: Sub-ring configuration example.....	96
Figure 24: Sub-ring homed to VPLS.....	98
Figure 25: Multi-ring hierarchy.....	99
Figure 26: QinQ frame.....	101
Figure 27: QinQ tagging example.....	103
Figure 28: QinQ using VPLS Ethernet SAPs.....	104
Figure 29: QinQ using VPLS ATM SAPs.....	105
Figure 30: IP transport service.....	109
Figure 31: TCP/UDP packet transport over IP/MPLS.....	110
Figure 32: IES/VPRN IP transport service.....	111
Figure 33: Raw socket and Cpipe support on the 7705 SAR.....	112
Figure 34: Service creation and implementation flowchart.....	114
Figure 35: ATM VLL for end-to-end ATM service.....	190
Figure 36: ATM virtual trunk over IP/ MPLS packet-switched network.....	191
Figure 37: E1 framing for CAS support in an E1 multiframe.....	199
Figure 38: SAToP MPLS encapsulation.....	202
Figure 39: CESoPSN MPLS encapsulation.....	202
Figure 40: CESoPSN packet payload format for trunk-specific n x 64 kb/s (with and without CAS transport).....	203
Figure 41: AMP hitless simultaneous transmission of TDM traffic over two paths with symmetrical TDM service.....	216
Figure 42: Control word bit structure.....	218
Figure 43: Ethernet VLL frame with MPLS encapsulation.....	221

Figure 44: Epipe service.....	221
Figure 45: Epipe network configuration with ATM SAP.....	223
Figure 46: CESoETH encapsulation.....	223
Figure 47: TDM SAP to Ethernet SAP.....	224
Figure 48: TDM SAP to spoke SDP.....	224
Figure 49: FR VLL for end-to-end FR service.....	238
Figure 50: FR frame.....	239
Figure 51: FR PW 1-to-1 MPLS PSN encapsulation.....	239
Figure 52: FR PW control word.....	240
Figure 53: HDLC VLL for end-to-end HDLC service.....	242
Figure 54: HDLC VLL frame.....	243
Figure 55: HDLC VLL frame with MPLS encapsulation.....	243
Figure 56: HDLC PW control word.....	244
Figure 57: IP pseudowires between 7705 SAR nodes.....	246
Figure 58: FR header with NLPID support for IP interworking PW.....	252
Figure 59: cHDLC header frame.....	252
Figure 60: SLARP keepalive frame.....	253
Figure 61: PPP and MLPPP to IP PWs on 2-port OC3/STM1 Channelized Adapter cards.....	253
Figure 62: Simplex to simplex pseudowire switching.....	257
Figure 63: Simplex to redundant pseudowire switching.....	258
Figure 64: Pseudowire switching network.....	259
Figure 65: Pseudowire switching TLV.....	260
Figure 66: N-to-1 cell mode encapsulation.....	267

Figure 67: VCC cell mode encapsulation with $N > 1$	269
Figure 68: MTU points on the 7705 SAR.....	275
Figure 69: Pseudowire redundancy.....	282
Figure 70: Implicit and explicit endpoint objects.....	283
Figure 71: Pseudowire redundancy with four spoke SDPs.....	285
Figure 72: Access side failure with ICB protection.....	286
Figure 73: Network side failure with ICB protection.....	287
Figure 74: Active/standby mode for redundant pseudowires.....	288
Figure 75: Standby-signaling-master enabled.....	289
Figure 76: Standby-signaling-slave enabled.....	290
Figure 77: Mixed microwave link scenario.....	309
Figure 78: Ethernet-to-ATM interworking on the 7705 SAR-M.....	310
Figure 79: SDPs – unidirectional tunnels.....	311
Figure 80: VPLS service architecture.....	526
Figure 81: Access port ingress packet format and lookup.....	527
Figure 82: Network port egress packet format and flooding.....	528
Figure 83: Access port egress packet format and lookup.....	529
Figure 84: Typical pseudowire-based mobile backhaul.....	530
Figure 85: Local VPLS on 7705 SAR in mobile backhaul.....	531
Figure 86: Spoke-SDP termination to VPLS using 7705 SAR-18 routers.....	532
Figure 87: ATM and IP DSLAM backhaul.....	534
Figure 88: PPPoE initialization and agent ID push function.....	537
Figure 89: Agent circuit ID information.....	537

Figure 90: Video over mobile backhaul.....	542
Figure 91: Metro cell multicast.....	543
Figure 92: Multiple hosts in an r-VPLS receiving the same channel.....	544
Figure 93: Multiple hosts in an r-VPLS receiving different channels.....	545
Figure 94: PIM snooping example.....	547
Figure 95: Simulating r-VPLS.....	548
Figure 96: H-VPLS with spoke SDP redundancy.....	557
Figure 97: Dual-homed MTUs in two-tier hierarchical VPLS.....	558
Figure 98: Dual-homed CE connection to VPLS.....	561
Figure 99: Example of ATM PVC access and termination on a VPLS.....	562
Figure 100: SDPs – unidirectional tunnels.....	578
Figure 101: IES for customer access to the Internet.....	822
Figure 102: SDP ID and VC label service identifiers (conceptual view of the service).....	830
Figure 103: IES spoke-SDP termination.....	830
Figure 104: Pseudowire-based backhaul (spoke-SDP termination at 7750 SR).....	831
Figure 105: Virtual private routed network.....	990
Figure 106: Firewall protection for the network core.....	995
Figure 107: Route distinguisher structure.....	997
Figure 108: Directly connected IP target.....	999
Figure 109: Multiple hops to IP target.....	999
Figure 110: IPv4 in-band management using a VPRN configured with GRT lookup.....	1005
Figure 111: IPv6 in-band management using a VPRN configured with GRT lookup.....	1006
Figure 112: RIP metric propagation in VPRNs.....	1019

Figure 113: Multicast in an IP-VPN application.....	1021
Figure 114: I-PMSI and S-PMSI.....	1023
Figure 115: P2MP FEC and MP LDP opaque value as per RFC 6388.....	1024
Figure 116: BGP MVPN address family updates.....	1025
Figure 117: I-PMSI sender-receiver, sender-only, and receiver-only: optimized I-PMSI mesh.....	1030
Figure 118: Inter-AS option B: non-segmented solution.....	1032
Figure 119: Unicast VPN option C with segmented MPLS.....	1033
Figure 120: Non-segmented mLDP PMSI establishment (option C).....	1034
Figure 121: Non-segmented mLDP C-multicast exchange (option C).....	1035
Figure 122: SDP ID and VC label service identifiers (conceptual view of the service).....	1039
Figure 123: VPRN spoke-SDP termination.....	1039
Figure 124: Pseudowire-based backhaul (spoke-SDP termination at 7750 SR).....	1040
Figure 125: VPRN in mobile backhaul application.....	1041
Figure 126: Spoke-SDP termination to VPRN.....	1041
Figure 127: Access IPv6 traffic aggregation and encryption.....	1043
Figure 128: IPSec implementation architecture.....	1571
Figure 129: VPRN public service IPSec transport tunnels.....	1576
Figure 130: Routing GRE-encapsulated packets over IPSec.....	1578
Figure 131: VLL/VPLS over IPSec over MPLS using BGP 3107 label routes.....	1580
Figure 132: VLL/VPLS over IPSec over MPLS using a loopback address.....	1582
Figure 133: Typical metrocell deployment.....	1591
Figure 134: Typical small business deployment.....	1592
Figure 135: UDP header injected by a NAT-T-enabled IPSec tunnel.....	1593

Figure 136: IPSec over BGP 3107 label route.....	1615
Figure 137: NGE network with NSP NFM-P management.....	1736
Figure 138: Key group partitioning.....	1738
Figure 139: NGE domain transit.....	1739
Figure 140: Private IP/MPLS network NGE domain.....	1740
Figure 141: Private over intermediary network NGE domain.....	1741
Figure 142: Key groups and a typical NGE packet.....	1743
Figure 143: NGE MPLS/GRE label stack.....	1746
Figure 144: NGE encryption and packet formats.....	1747
Figure 145: Inbound and outbound key group assignments.....	1748
Figure 146: Router interface encryption packet format (IPSec transport mode).....	1751
Figure 147: Inside and outside NGE domains.....	1752
Figure 148: Router interface NGE exception filter example.....	1754
Figure 149: IPSec packets transiting an NGE domain.....	1755
Figure 150: Processing multicast packets.....	1756
Figure 151: Firewall considerations.....	1757
Figure 152: Encrypted Layer 2 packet.....	1759
Figure 153: QoS for NGE traffic (network ingress).....	1764
Figure 154: EVPN for MPLS in VPLS services.....	1794
Figure 155: EVPN-MPLS route type 2 and type 3 (required routes and communities).....	1796
Figure 156: EVPN route type 1 and type 4.....	1797
Figure 157: EVPN route type 5.....	1801
Figure 158: EVPN-VPLS integration.....	1807

Figure 159: DF election.....	1810
Figure 160: Split horizon.....	1810
Figure 161: Aliasing.....	1811
Figure 162: ES discovery and DF election.....	1813
Figure 163: All-active multihoming ES failure.....	1821
Figure 164: Black hole caused by SAP or service shutdown.....	1822
Figure 165: Transient issues caused by slow MAC learning.....	1823
Figure 166: Backup PE.....	1824
Figure 167: Single-active multihoming ES failure.....	1829
Figure 168: EVPN-VPWS BGP extensions.....	1830
Figure 169: EVPN-MPLS in Epipe services.....	1832
Figure 170: Active/standby PW and MC-LAG support on EVPN-VPWS.....	1834
Figure 171: EVPN-VPWS single-active multihoming.....	1836
Figure 172: EVPN-MPLS multihoming in r-VPLS services.....	1838
Figure 173: DF election extended community attribute.....	1840
Figure 174: IP-VPN import and EVPN export BGP workflow.....	1850
Figure 175: EVPN import and IP-VPN export BGP workflow.....	1851
Figure 176: R-VPLS with EVPN tunnel, without multihoming.....	1856
Figure 177: EVPN-MPLS r-VPLS with all-active multihoming ES.....	1861
Figure 178: EVPN-MPLS r-VPLS with single-active multihoming.....	1870

1 Preface

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 25.x content and may contain some content that will be released in later maintenance loads. See the 7705 SAR 25.x.Rx Software Release Notes, part number 3HE21362000xTQZZA, for information about features supported in each load of the Release 25.x software.



Note: As of Release 23.4, software support for the following hardware has been deprecated:

- 8-port Ethernet Adapter card, version 2 (a8-ethv2) (3HE02776)
- 12-port Serial Data Interface card, version 1 (a12-sdi) (3HE03391)
- 7705 SAR-W (3HE07349)

These components are no longer recognized in the release.

If information about any of the above components is required, please see the applicable installation guides in Release 22.10.

1.1 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

1.2 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 a Nokia service agreement, follow this link to contact a Nokia support representative and to access product manuals and documentation updates:

[Product Support Portal](#)

2 7705 SAR services configuration process

The following table lists the tasks that are required 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: Configuration process

Area	Task/description	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
	Ipip 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
VPNR	Configure a Layer 3 multipoint-to-multipoint VPN service as defined in RFC 2547bis	VPNR services
IPSec	Configure IPSec service	IPSec
Network group encryption	Configure network group encryption	Network group encryption
Ethernet virtual private networks	Configure EVPN, which allows VPLS services to be operated as IP-VPNs	Ethernet virtual private networks
Reference	List of security and telecom standards, supported protocols, and proprietary MIBs	Supported standards and protocols

3 Services overview

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](#)
- [Nokia service model](#)
- [Service entities](#)
- [High-speed download packet access offload](#)
- [ETH-CFM \(802.1ag and Y.1731\)](#)
- [G.8032 Ethernet ring protection switching](#)
- [QinQ support](#)
- [Raw socket IP transport service](#)
- [Service creation overview](#)
- [Port and SAP CLI identifiers](#)
- [Configuring global service entities with CLI](#)
- [Global service command reference](#)

3.1 Introduction to services on the 7705 SAR

Topics in this section include:

- [Overview](#)
- [Service types](#)
- [Service policies](#)

3.1.1 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 [Nokia 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 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 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 (for example, a 7750 SR), specifying 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 instead of 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 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 when it is in 10-port 1GigE 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.

3.1.2 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](#).

The following table 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

PW service type (Ethertype)	Value
Ethernet raw	0x0005
HDLC	0x0006
ATM N-to-one VCC cell mode ¹	0x0009
ATM N-to-one VPC cell mode ¹	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.

3.1.3 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, Epipe and Ipipe SAPs, 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 about provisioning QoS policies, including queuing behaviors, see the 7705 SAR Quality of Service Guide. For information about configuring IP and MAC filter policies, see the 7705 SAR Router Configuration Guide.

3.2 Nokia 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.

3.3 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](#)
- [Service destination points](#)

The following figure 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 the figure, 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 about SAP encapsulation types, see [SAP encapsulation types and identifiers](#). For information about SDP encapsulation types, see [SDP encapsulation types](#).

Figure 1: Service entities and the service model



19686

3.3.1 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.

3.3.2 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.

3.3.2.1 Service names

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

3.3.3 Service access points

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: Service access point \(SAP\)](#) 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. See the 7705 SAR 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)



19479

3.3.3.1 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](#)
- [Service types and SAP encapsulations – summary](#)

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.

For SAPs configured on the 16-port T1/E1 ASAP Adapter card, 32-port T1/E1 ASAP Adapter card, or 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 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:

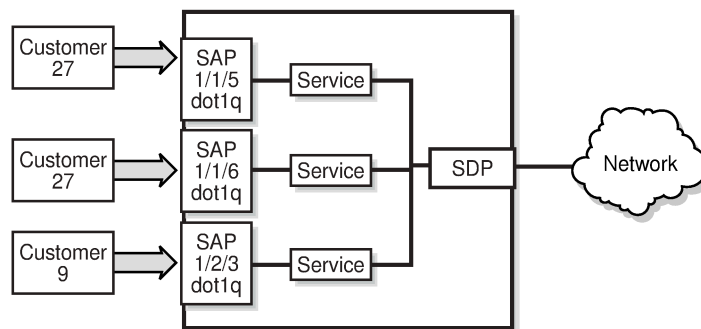
- Throughout this guide, the term "channel group" is often simplified to "channel".
- Do not confuse the term "encapsulation ID" (described here) with the term "Encapsulation ID", which is used with the SNMP and MIBs for the 7705 SAR.

3.3.3.1.1 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: Multiple SAPs on a single port/channel](#)). An example of dot1q use may 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 about QinQ, see [QinQ support](#).

Figure 3: Multiple SAPs on a single port/channel



19480

3.3.3.1.1.1 Default SAP on a dot1q or 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 functions 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 functions 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.

3.3.3.1.2 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

3.3.3.1.3 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 DS1/E1 interfaces or on unstructured (unframed) E1 interfaces. Unframed E1 can be used for lpipe support. Channelized interfaces are 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)

3.3.3.1.4 Service types and SAP encapsulations – summary

The following table 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 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

Service (port) type	Encapsulation option
Ethernet	dot1q
Ethernet	qinq
SONET/SDH	atm
TDM	cem
TDM	atm
TDM	ipcp
TDM	frame-relay
TDM	cisco-hdlc
TDM	hdlc

3.3.3.2 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 specific 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 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))

3.3.4 Service destination points

A service destination point (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 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. After 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](#).

3.3.4.1 SDP binding

To configure a distributed service pointing from ALU-A to ALU-B, the SDP ID on the ALU-A side (see the following figure) 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 devices 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



3.3.4.2 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.

3.3.4.3 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). This restriction is relaxed for some combinations of the transport methods when the mixed-LSP mode option is enabled on the SDP. See [Mixed-LSP SDPs](#) for more information.

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 about BGP route tunnels, see the 7705 SAR Routing Protocols Guide, "BGP route tunnels".

3.3.4.4 SDP encapsulation types

The Nokia 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.

3.3.4.4.1 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 allows 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)

3.3.4.4.2 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. Because 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.

3.3.4.4.2.1 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 the following figure and described in the table.

Figure 5: GRE header



19874

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

Field	Description
	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.
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 the following figure and described in the table.

Figure 6: GRE pseudowire payload packet over Ethernet



19873

Table 5: GRE pseudowire payload packet descriptions

Field	Description
Eth	The Layer 2 transport header The only Layer 2 protocol supported is Ethernet

Field	Description
	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 about 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 about ARP, see the 7705 SAR Router Configuration Guide.

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, the far-end interface address, or the loopback address. 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 may 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).

3.3.4.4.2.2 GRE fragmentation

IP fragmentation can be enabled for GRE tunnels. Services for which fragmentation is typically not available can make use of IP fragmentation performed at the IP layer of the GRE tunnel. The IP fragmentation feature can be enabled at the GRE tunnel ingress by enabling the **allow-fragmentation** command on the SDP. The IP fragmentation size limits are derived from the MTU of the network port used by the GRE tunnel.

At the GRE tunnel egress, IP reassembly can be performed as specified by a reassembly profile assigned to the network interfaces on which the GRE packets are expected to arrive. The IP reassembly function is performed on the IP fragments received at the GRE tunnel egress before any underlying service label is processed. A reassembly profile is used to specify the amount of buffer space allocated for the IP reassembly function and to configure a reassembly timeout. These parameters are configured for each forwarding class to isolate different types of GRE traffic.



Note: The GRE reassembly function is not supported for fragments of a single GRE packet that arrive on different MDAs. All fragments of a GRE packet must arrive on the same MDA.

When **allow-fragmentation** is enabled on an SDP, the current MTU algorithm is overwritten with the configured path MTU. The administrative MTU and operational MTU both show the specified MTU value. If the path MTU is not configured or available, the operational MTU is set to 2000 bytes, and the administrative MTU displays a value of 0. When **allow-fragmentation** is disabled, the operational MTU reverts to the previous value.

Fragmentation is supported on the following types of GRE SDPs:

- VPLS
- Layer 3 spoke SDP
- Epipe

The GRE SDPs can be NGE-encrypted; however, the NGE interface must be Ethernet. Fragmentation is not supported on NGE PPP/MLPPP interfaces. See [GRE fragmentation for NGE packets](#) for more information.

IP packets that are transported over a Layer 3 spoke SDP using a fragmentation-enabled GRE tunnel are handled differently depending on the DF bit setting and the size of the packet:

- If the packet DF bit setting is 0 (Fragment), the GRE fragment size is determined by the network port MTU value.
- If the packet DF bit setting is 1 (Do not fragment) and the packet size is smaller or equal to the smaller value of either the operational MTU of the spoke SDP or the Layer 3 spoke SDP interface MTU (if configured), the packet is sent through the GRE tunnel.
- If the packet DF bit setting is 1 (Do not fragment) and the packet size is larger than the smaller value of either the operational MTU of the spoke SDP or the Layer 3 spoke SDP interface MTU (if configured), the packet is discarded and an ICMP message "Fragmentation Needed and Don't Fragment was Set" is sent back to the source IP address.

IP reassembly profiles are required to ensure that all packet fragments are received within an expected time frame for each forwarding class. When the reassembly profile timers expire, all fragments of the corresponding incomplete frame are dropped and a "Fragment Reassembly Time Exceeded" ICMP error message is sent to the source node.



Note: The system checks the reassembly queues every 64 ms in a constant loop, which may cause a maximum of 63 ms variation between the user-configured value and the actual detection time. For example, using the default configuration of 2000 ms, the system may check the reassembly queue timer at 1999 ms, in which case the timeout would not occur during that cycle and would instead take place during the next cycle at 2063 ms.

Traffic ingressing a GRE tunnel can use different forwarding classes and different queues. If multiple queues are transmitting fragments, a higher-priority queue could interrupt the transmission of fragments of a frame in a lower-priority queue by interleaving fragments of another frame. If the fragments from the different frames have similar IP identifiers, they could be reassembled incorrectly into one frame at the tunnel egress. To prevent this incorrect reassembly of frames, the 7705 SAR that is performing the IP fragmentation uses 4 bits of the 16-bit IP identifier to indicate the transmitting queue at the tunnel ingress. The IP identifier is part of the IP reassembly tuple, which also contains the protocol, source address, and destination address. Using the IP reassembly tuple, fragments of frames from different queues are always differentiated. However, reserving 4 bits in the IP identifier field leaves only 12 bits to act as a sequence number, which causes a shorter identifier wraparound. The time required for the IP identifier to wrap around is a function of the traffic rate on a specific queue at the GRE tunnel ingress.

If fragments are dropped along the GRE tunnel due to congestion or bit errors, the 7705 SAR that is performing the IP reassembly at the tunnel egress normally drops partially reassembled packets due to expiration of the reassembly timeout interval. If fragment loss occurs in the network along with an IP identifier wraparound due to a high packet rate, the IP reassembly block may incorrectly insert fragments of a new frame into a frame of older fragments that are waiting for timeout. When configuring the timeout interval, it is therefore important to factor in the pre-fragmentation frame rate for forwarding classes on a GRE tunnel. As a guideline, higher-priority packets should have shorter timeout intervals than other packets because their queue interruption at the ingress GRE tunnel is minimal, and the timeout intervals should be shorter than the transmission time of 4096 packets for that forwarding class.

The 7705 SAR does not support double fragmentation.

3.3.4.4.3 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.

The 7705 SAR supports the transport of pseudowires over IP tunnels. [Figure 7: IP example of pseudowire payload packet over Ethernet](#) shows an example of an application using Apipes over IP over Ethernet.

3.3.4.4.3.1 Payload packet

A typical payload encapsulation format for pseudowires over IP is shown in the following figure and described in the table.

Figure 7: IP example of pseudowire payload packet over Ethernet



20695

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

Field	Description
	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
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 about 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 about ARP, see the 7705 SAR Router Configuration Guide.

3.3.4.5 Spoke-SDP terminations

The 7705 SAR supports spoke SDP as termination points for IES and VPRN services. The following table 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. See [Spoke-SDP termination to IES](#) and [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	✓	✓	✓	✓	✓	✓

	Epipespoke SDP	Epipespoke SDP redundancy (standby-signaling master enabled)	IES interface	VPRN interface	VPLS spoke SDP	VPLS spoke SDP redundancy (suppress-standby-signaling disabled)
VPLS spoke SDP redundancy (suppress-standby-signaling disabled)	✓	✓	✓	✓	✓	✓

3.3.4.6 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. See the 7705 SAR OAM and Diagnostics Guide, "SDP ping", for more information.

3.3.4.7 SDP keepalives

The SDP keepalive application allows a system operator to actively monitor the SDP operational state using periodic Nokia 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 specific 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, see [Configuring SDPs](#).

3.3.4.8 Mixed-LSP SDPs

If mixed-LSP SDP mode is enabled on an SDP, a maximum of two LSP types can be configured on the SDP: a primary LSP and a secondary (backup) LSP. Two combinations are possible:

- an RSVP-TE primary LSP backed up by an LDP LSP
- an LDP primary LSP backed up by a BGP LSP

The **config>service>sdp mpls>mixed-lsp-mode** command is used to configure a mixed-LSP SDP.



Note: Mixed-LSP SDPs do not support static LSPs on either the primary or backup.

Mixed-LSP mode of operation

The service manager programs only one type of LSP in the line card, which activates it to forward service packets. The LSPs are programmed in the following priority order:

1. RSVP-TE LSP type

This is the highest-priority LSP type. Up to eight RSVP-TE LSPs can be entered by the user and programmed by the service manager in the ingress line card to load-balance service packets.

2. LDP LSP type

One LDP FEC is programmed by the service manager, but the ingress line card can use up to eight LDP ECMP paths for the FEC to load-balance service packets when ECMP is enabled on the node.

3. BGP LSP type

One RFC 3107-labeled BGP prefix is programmed by the service manager. The ingress line card can use more than one next hop for the prefix.

For an RSVP-TE/LDP SDP, the service manager programs the NHLFEs for the active LSP type, preferring the RSVP-TE LSP type over the LDP LSP type. If no RSVP-TE LSP is configured, or if all of the configured RSVP-TE LSPs go down, the service manager reprograms the line card with the LDP LSP, if available. If no LDP LSP is available, the SDP goes operationally down.

For LDP/BGP SDPs, the service manager prefers the LDP LSP type over the BGP LSP type. If no LDP LSP is configured or all configured LDP LSPs go down, the service manager reprograms the line card with the BGP LSP if it is available; otherwise, the SDP goes operationally down.

An LDP/BGP SDP differs from an RSVP/LDP SDP in the number of routes available. For any given /32 prefix, only a single route exists in the routing table: the IGP route or the BGP route. Therefore, only the LDP FEC or the BGP label route is active at any given time. The impact of this is that the tunnel table must be reprogrammed each time a route is deactivated and the other route is activated. In this scenario, the SDP **revert-time** command cannot be used because there is no situation where both LSP types are active for the same /32 prefix.

When a higher-priority LSP type becomes available, the service manager resets the SDP configuration to this LSP type when the revert timer expires or when the current active LSP fails, whichever occurs first. The length of time that the service manager must wait can be configured with the **config>service>sdp mpls>mixed-lsp-mode>revert-time** command. After the SDP has reverted to the higher-priority LSP, the service manager reprograms the line card accordingly. If the revert timer is configured with the **infinite** parameter, the service manager never resets the SDP to the highest-priority LSP type unless the current active LSP fails.

If the value of the revert time timer is changed, it takes effect when the timer is next activated. Any timer that is currently active when the value is changed is restarted with the new value.



Note: LDP uses a tunnel-down-damp timer that is set to 3 seconds by default. If the LDP LSP fails, the SDP reverts to the RSVP-TE LSP type after the expiry of this timer. For an immediate switchover, this timer must be set to 0 with the **config>router>ldp>tunnel-down-damp-time** command.

3.3.4.9 Multiple load-balancing LSPs under a single SDP

Configuring multiple LSPs under a single SDP allows load distribution among multiple LSPs to the same destination. This load distribution is handled by the node without the need for any operator intervention. LSP additions or deletions result in automatic rehashing of services onto remaining LSPs, making it transparent to the operator. No new hashing algorithms are required; existing hashing algorithms are extended to select an LSP from multiple LSPs under an SDP.

Up to eight RSVP-TE or SR-TE LSPs can be configured under a single SDP. However, a mix of RSVP-TE and SR-TE LSPs is not supported. When the first LSP is configured under the SDP, all other LSPs configured under that SDP must be of the same type.

Multiple LSPs are only supported for SDPs configured for MPLS encapsulation.

3.4 High-speed download packet access offload

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.

Nokia 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.

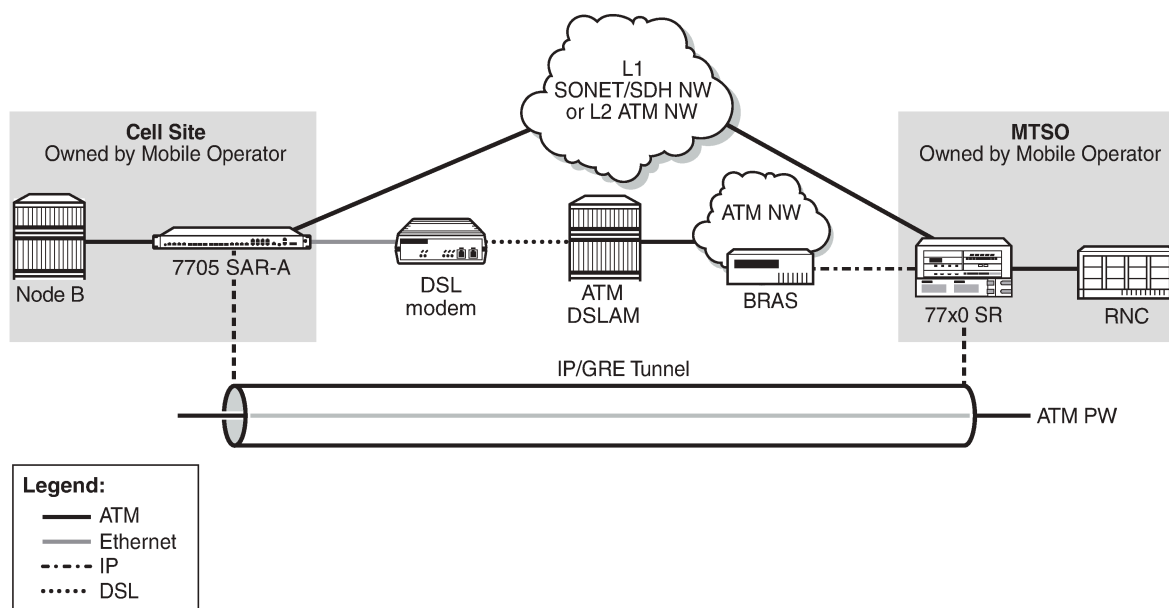
3.4.1 HSDPA offload

The Nokia 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-A is an ideal candidate for this scenario. The 7705 SAR-A is a small-scale version of the 7705 SAR product family, optimized for use in standalone small or mid-sized sites where traffic aggregation from multiple cell sites is not needed. For more information about the 7705 SAR-A, see the 7705 SAR-A Chassis Installation Guide.

The following figure shows an example of HSDPA offload with the 7705 SAR-A.

Figure 8: 7705 SAR-A HSDPA offload example



19872

A 3G Node B is connected to a 7705 SAR-A over an ATM/IMA access port (SAP endpoint). An ATM SAP-to-SAP connection is set up in the 7705 SAR-A 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.

3.4.1.1 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 NSP NFM-P to the ATM ports used for backhauling the traffic. The mission-critical traffic is still serviced before the best-effort HSDPA traffic.

For information about OAM and SAA tools, see the 7705 SAR OAM and Diagnostics Guide. For information about BFD, see the 7705 SAR Router Configuration Guide.

3.5 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 (Epipe and VPLS), Ethernet spoke SDPs, and VPLS spoke/mesh SDPs, 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, see the "ITU-T Y.1731 performance monitoring (PM)" section in the 7705 SAR OAM and Diagnostics Guide.

For information about running Ethernet OAM tests, see the "ETH-CFM (802.1ag and Y.1731)" section in the 7705 SAR OAM and Diagnostics Guide.

The information in this section is specific to Ethernet SAPs and spoke and mesh SDPs, although most of it also applies to Ethernet network interfaces. For information about ETH-CFM support specific to network interfaces, see the 7705 SAR 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 about 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 SAPs (**config>service>epipe>sap>eth-cfm**)
- Epipe spoke SDPs (**config>service>epipe>spoke-sdp>eth-cfm**)
- VPLS SAPs (**config>service>vpls>sap>eth-cfm**)
- VPLS spoke SDPs (**config>service>vpls>spoke-sdp>eth-cfm**)

- VPLS mesh SDPs (**config>service>vpls>mesh-sdp>eth-cfm**)
- network interface (**config>router>if>eth-cfm**)
- show (**show>eth-cfm**)
- oam (**oam>eth-cfm**)

3.5.1 802.1ag and Y.1731 terminology

The following table defines 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 service
MA-ID	Maintenance association identifier	A unique combination of MD index (<i>md-index</i>), MD level (<i>level</i>), and MA index (<i>ma-index</i>), where <i>md-index</i> , <i>level</i> , and <i>ma-index</i> 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, spoke SDP, or mesh SDP (VPLS only)
MEP	Maintenance association endpoint	An (edge) endpoint that can terminate, respond to, or initiate the OAM messages for a configured MD-MA combination
MEP-ID	Maintenance association endpoint identifier	A MEP is identified by its MEP-ID, which is a unique combination of MD index (<i>md-index</i>) and MA index (<i>ma-index</i>), where <i>md-index</i> and <i>ma-index</i> 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

The following table illustrates the similarities and differences between Y.1731 and 802.1ag terms.

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 Although MEG level and MD level are equivalent terms, there is no Y.1731 equivalent to an MD
MEP	Maintenance association endpoint	Same as MEP for 802.1ag

3.5.1.1 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. MD levels are used to set up a messaging hierarchy for the CFM architecture.


An MA consists of up to eight MEPs (one local and up to seven remote) for Up MEPs on Epipe and VPLS services and two MEPs (one local and one remote) for Down MEPs on Epipe and VPLS services. The MA and the service are associated by configuring the MA bridge identifier parameter to have the same value as the service ID of the 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, spoke SDP, or mesh SDP (VPLS only). 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.

The following figure shows 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 about MEP support, see [Ethernet OAM](#).

Figure 9: MEPs and MAs



20497

3.5.1.2 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).

3.5.2 ETH-CFM frame format

ETH-CFM OAMPDU messages for 802.1ag and Y.1731 use a standard Ethernet frame (see the following figure and descriptions).

Figure 10: ETH-CFM frame format

Destination MAC
Source MAC
T = 8100
Vlan / Dot1p
T = 8902
ETH-CFM OAMPDU
FCS

20472

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-C2-00-00-3x; 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 (loopbacks).

Ethertype (T)

If dot1q or qinq encapsulation is not configured, 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

The Vlan/Dot1p tag is the VLAN/dot1p identifier. If null encapsulation is configured (for Ethernet SAPs or spoke or mesh 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 information about the dot1ag or Y.1731 OAMPDU, see [ETH-CFM OAMPDU](#).

FCS


The FCS is the frame check sequence field.

3.5.2.1 ETH-CFM OAMPDU

As shown in [Figure 11: ETH-CFM OAMPDU message](#), 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 11: ETH-CFM OAMPDU message



20847

3.5.2.2 CFM frame processing

The following table shows whether a CFM frame received by various MEP types is processed. 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
	Up	Raw	Any	Processed	Not processed
		VLAN	Any	Not processed	Processed
Mesh SDP	Down	Raw	Any	Processed	Not processed
		VLAN	Any	Not processed	Processed
	Up	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 or VPLS pseudowire configured to handle untagged frames. The SAP identifier 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.

3.5.2.2.1 Processing SAP 0 and SAP 0.* OAM packets

The following points describe the processing of OAM packets on SAP 0 (dot1q) and SAP 0.* (qinq). 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
- 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.*

3.5.2.3 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:

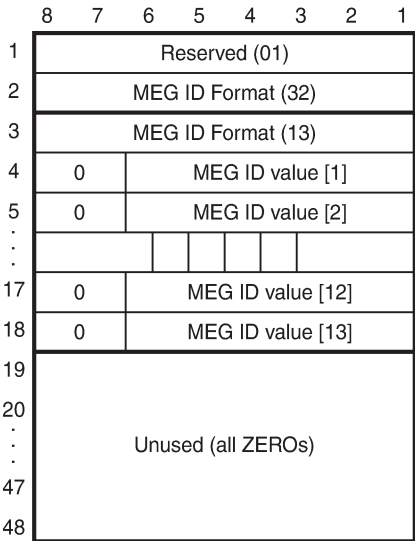
- each MEG-ID must be globally unique
- if the MEG may be required for path setup across interoperator boundaries, the MEG-ID must be available to other network operators
- the MEG-ID should not change while the MEG remains in existence
- the MEG-ID should be able to identify the network operator that 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. The following figure 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 12: ICC-based MEG-ID format



20848

3.5.3 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, spoke SDP, or mesh SDP (VPLS only)
- 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
 - virtual private networks identifier (RFC 2685)
 - 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, spoke SDP, or mesh 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)

3.5.3.1 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



Note: The 7705 SAR also supports Ethernet bandwidth notification (ETH-BN) on the client side. For information, see the 7705 SAR OAM and Diagnostics Guide, "ITU-T Y.1731 Ethernet bandwidth notification (ETH-BN)".

3.5.3.1.1 Loopback


The loopback (LB) 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 the following figure, 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 is processed. After 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 13: Dot1ag loopback test



20480

3.5.3.1.2 Linktrace

The linktrace (LT) 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 a MEP
- MEP – supports generating LTMs and responding with LTR messages
- displays linktrace test results on the originating MEP

3.5.3.1.3 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. Therefore, 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. The MAC address cannot be a multicast MAC address; it must be the MEP MAC destination address (DA).

Datapath processing of LBMs is supported with dot1ag and Y.1731 for the following MEPs:

- SAP Up and Down MEPs
- spoke SDP Up and Down MEPs
- mesh SDP Up and Down MEPs (VPLS only)

For spoke or mesh SDP Up and 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), the fastpath processing of LBMs is terminated and LBM processing continues via the CSM.

3.5.3.1.4 Continuity check

The continuity check (CC) 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. 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

3.5.3.1.5 Remote defect indication

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

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.

3.5.3.1.6 Alarm indication signal

The Ethernet alarm indication signal (ETH-AIS) function 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 SAP 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. After the AIS condition is cleared, a MEP resumes alarm generation upon detecting defect conditions.

The following specific configuration information is required by a SAP 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

3.5.3.1.7 Ethernet test

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



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 SAP or facility 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 toward a targeted peer MEP. The receiving MEP detects the frames with ETH-Test information and performs the requested measurements.

3.5.4 MEP support (802.1ag and Y.1731)

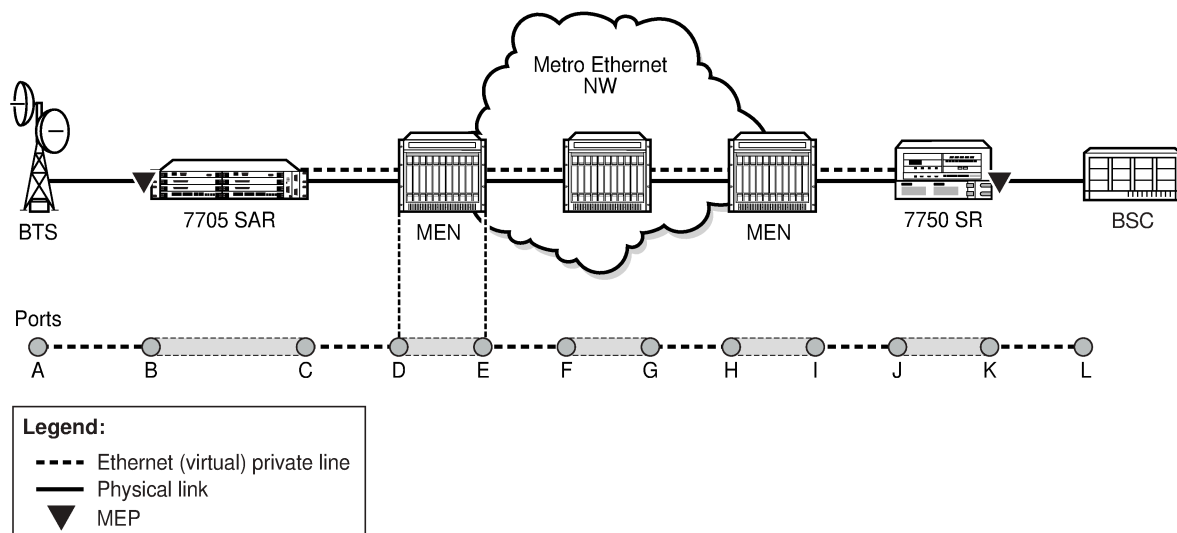
The 7705 SAR supports the following MEPs for 802.1ag and Y.1731:

- Up and Down MEPs on Ethernet (Epipe and VPLS) SAPs
- Up and Down MEPs on Ethernet (Epipe and VPLS) spoke SDPs
- Up and Down MEPs on mesh SDPs (VPLS only)
- facility MEPs on network interfaces

3.5.4.1 802.1ag MEP support on Ethernet SAPs

The 7705 SAR supports Up and Down MEPs on Ethernet SAPs. The following figure 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 toward port A.

Figure 14: MEP on Ethernet access




20474

The following figure shows how a Down MEP on an Ethernet SAP may 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). Because CFM offers OAM capabilities on a per-service basis, which in this case means per SAP (or endpoint), each service can run CFM. 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 the figure, separate dot1ag instances initiated on the Wavence 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 15: Down MEP at Ethernet SAP



28769

3.5.4.2 802.1ag and Y.1731 MEP support on Ethernet spoke SDPs and mesh SDPs

The 7705 SAR supports Up and Down MEPs with 802.1ag and Y.1731 on Ethernet spoke SDP endpoints (Epipe and VPLS) and mesh SDP endpoints (VPLS only). [Figure 16: Dot1ag Down MEPs on spoke SDPs](#) illustrates a dot1ag Down MEP on an Ethernet spoke SDP.

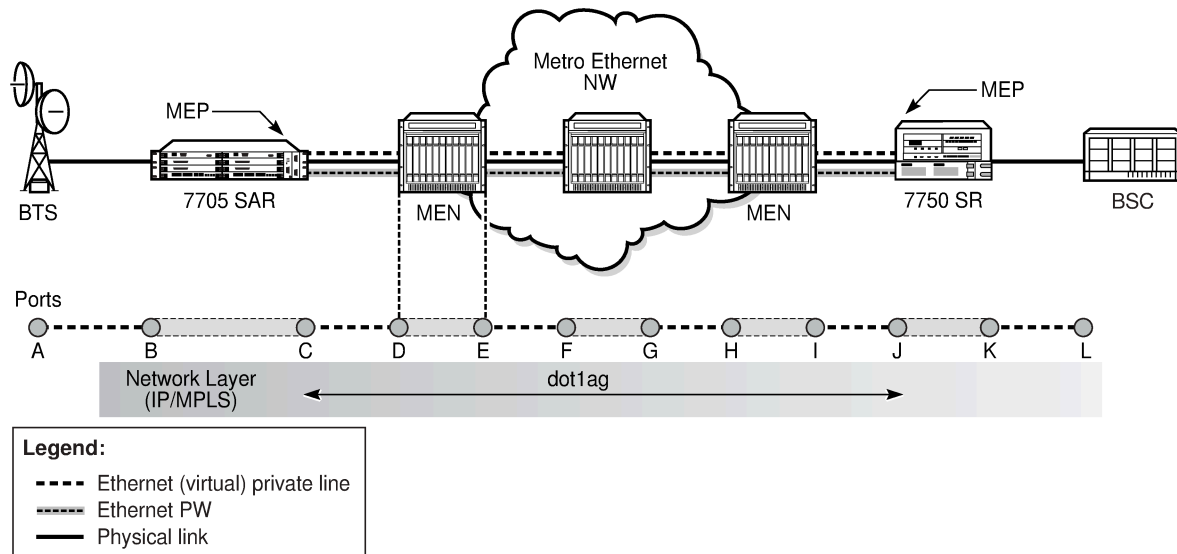
CFM messages can be generated and switched across an Ethernet PW. CFM messages that are received and have an MD that matches the value configured on the 7705 SAR are extracted and processed. Any received CFM messages with an MD level that does not match the configured value are not terminated and are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

Up and Down MEPs on Ethernet spoke and mesh 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.

Figure 16: Dot1ag Down MEPs on spoke SDPs



20476


In the figure, assuming that the MEP is enabled both on the 7750 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 7750 SR node. The Ethernet PW running between the 7705 SAR and the 7750 SR generates a pipe-like connectivity; therefore, 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 7750 SR and 7705 SAR endpoints.

3.5.4.3 Y.1731 MEP support on Ethernet SAPs

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

The figure 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 or mesh 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 or mesh SDP is on the network side. If a spoke or mesh SDP Down MEP is used, packets are terminated and extracted on the network side without being switched through the switching fabric.

Figure 17: Y.1731 MEP support on the 7705 SAR



20849

3.5.5 Priority mapping (802.1ag and Y.1731)

Operators often run OAM tests over a single, specific forwarding class (FC). For example, an operator may 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** 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. Therefore, 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 because the OAM packet follows the same path as the data packets.

3.5.5.1 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.

The following table 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 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

3.5.5.2 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).

3.6 G.8032 Ethernet ring protection switching

This section contains the following topics:

- [Overview of G.8032 operation](#)
- [Ethernet ring sub-rings](#)


The 7705 SAR supports Ethernet ring protection switching in accordance with ITU-T G.8032 to achieve resiliency for Ethernet Layer 2 networks. A G.8032 Ethernet ring is built on Ethernet OAM and is also referred to as Ring Automatic Protection Switching (RAPS).

Ethernet rings are supported on VPLS SAPs. VPLS services supporting Ethernet rings can connect to other rings and Ethernet services using VPLS and r-VPLS SAPs. Ethernet rings provide rings for core network or access network resiliency. A single point of interconnection to other services is supported. The Ethernet-ring service is a VLAN service providing protection for ring topologies and the ability to interact with other protection mechanisms for overall service protection. This combined service protection ensures that higher layers are isolated from failures because there will only be a RAPS switchover when the lower layer cannot recover.

Rings are preferred in data networks where the native connectivity is laid out in a ring or where there is a requirement for simple resilient LAN services. Due to the symmetry and the simple topology, rings are considered a good solution for access and core networks where resilient LANS are required. The 7705 SAR implementation can be used for interconnecting access rings and to provide traffic engineered backbone rings.

Even though 7705 SAR nodes are often connected via IP/MPLS links to each other or to higher hierarchies, the first level of connected aggregation nodes can significantly benefit from G.8032 protection switching. The following figure shows a common example where standalone microwave nodes are deployed in a ring that are connected to a 7705 SAR node acting as the head-end of the ring. Providing G.8032 Ethernet ring protection switching would provide significantly better reconvergence times in the ring and ensure minimal service disruption in case of failure.

Figure 18: Ethernet protection switching



28787

Ethernet rings use one VLAN ID per control per ring instance and use one or more VIDs for data instances per control instance. A dedicated control VLAN (ERP VLAN) is used to run the protocol on the control VLAN ID. G.8032 controls the active state for the data VLANs (ring data instances) associated with a control instance. Multiple control instances allow logically separate rings on the same topology.

The 7705 SAR supports dot1q and qinq encapsulation for data ring instances. The control channel supports dot1q and qinq encapsulation. The control channel can support dot1q while the data channels use qinq if the global **new-qinq-untagged- sap** command is enabled.

3.6.1 Overview of G.8032 operation


RAPS messages that carry the G.8032 protocol are sent on a dedicated protocol VLAN called the Ethernet ring protection (ERP) instance. Revertive and non-revertive behaviors are supported. In revertive mode, the G.8032 protocol ensures that one ring protection link (RPL) owner blocks the RPL. RAPS messages are periodically sent around the ring to inform other nodes in the ring about the blocked port in the RPL owner node. In non-revertive mode, any link may be the RPL.

Y.1731 Ethernet OAM CC is the basis of the RAPS messages. Nodes in the ring typically use Y.1731 CC messages to monitor the health of each link in the ring in both directions. CC messages are not mandatory. Other link layer mechanisms could be used, for example, loss of signal (LOS) for instances when the nodes are directly connected.

Initially each ring node blocks one of its links and notifies other nodes in the ring about the blocked link. After a ring node in the ring learns that another link is blocked, the node unblocks its blocked link possibly causing an FDB flush of all links in the ring for the affected service VLANs controlled by the ring control

instance. This results in unblocking all links except one so that the ring is in the normal, or idle, state. In revertive mode, the link that is blocked when all other links are operable after the revert time has expired becomes the RPL. In non-revertive mode the RPL is no different than other ring links. Revertive mode offers predictability, especially when there are multiple ring instances and the operator can control which links are blocked on each instance. When there is a topology change that affects reachability, the nodes may flush the FDB and MAC learning occurs for the affected service VLANs, which allows packet forwarding to continue. The following figure depicts this operational state.

Figure 19: G.8032 ring in the idle state




28779

When a ring failure occurs, any node that detects the failure (by using Y.1731 OAM CC monitoring) sends RAPS message in both directions. After receiving a failure notification, the nodes at both ends of the failed link block forwarding to the failed link to prevent it from becoming active.

When a ring failure occurs in revertive mode, the RPL owner unblocks the previously blocked RPL and triggers an FDB flush for all nodes for the affected service instances. The ring is now in the protecting state and full ring connectivity is restored. MAC learning occurs, which allows Layer 2 packet forwarding on the ring. The following figure depicts a G.8032 ring in the protecting state.

Figure 20: G.8032 ring in the protecting state



28780

After the failed link recovers, the nodes that blocked the link send RAPS messages indicating no failure. This triggers the RPL owner to block the RPL link and indicate the blocked link in its own RAPS message. When the nodes at the recovered link receive the RPL RAPS message, they unblock the specified link, restore connectivity, and all nodes in the ring perform an FBD flush. MAC learning takes place and the ring returns to the normal, or idle, state.

Each path uses Y.1731 maintenance entity groups (MEGs) and Maintenance Endpoints (MEPs) to exchange RAPS-specific information to co-ordinate switchovers. As well Y.1731 MEGs and MEPs optionally use fast Continuity Check Messages (CCM), which provide an inherent fault detection mechanism as part of the protocol. When a ring path failure is detected, the protection links are activated. During a failure, reconvergence times depend on the failure detection mechanisms being used. For Y.1731, the CCM transmit interval determines the response time. The router supports message timers as low as 10 ms, providing restoration times comparable to SONET/SDH. Alternatively, 802.3ah (Ethernet in the First Mile) or simple LOS can act as a trigger for a protection switch where appropriate. Where nodes are directly connected there is no need to use Ethernet CC messaging for liveness detection.

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

Ethernet RAPS can be configured on any port configured for access mode by using dot1q or qinq encapsulation, enabling support for Ethernet RAPS protected services on the service edge toward the customer site or within the Ethernet backbone.

The Ethernet ring is built from a VPLS service on each node with VPLS SAPs that provide ring paths with SAPs. As a result, most of the VPLS SAP features are also available on Ethernet rings resulting in a feature-rich ring service.

The control tag defined under each Ethernet ring is used for encapsulating and forwarding the CCMs and the G.8032 messages used for the protection function. If a failure of a link or node affects an active Ethernet ring segment, the services will fail to receive the CCMs exchanged on that segment or will receive a fault indication from the link layer OAM module. CCMs are optional but MEPs are always configured to provide G.8032 control. The forwarding of CCMs and G.8032 RAPS messages continues in the control VPLS even if the service or its SAPs are administratively shut down. The Ethernet ring instance can be shut down to stop the operation of the ring on a node.


For fault detection using CCMs, three CC messages and a configurable hold-off timer must be missed for a fault to be declared on the associated path. The hold-off timer is required in order to support an additional 50 ms resiliency mechanism in the optical layer. After receiving the fault indication, the protection module declares the associated ring link down and the G.8032 state machine sends the appropriate messages to open the RPL and flush the learned addresses.

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

The following figure illustrates a resilient ring service. In the example, a ring (solid line) that uses dot1q VLAN ID 100 carries service VID 500. The RPL for the ring is between A and B, where B is the RPL owner. Also illustrated is a QinQ service on the (dotted line) ring that uses dot1q VLAN ID 600 for the ring to connect service VLAN 100.50. The two rings have RPLs on different nodes, which allows a form of load balancing.

The figure illustrates that service encapsulations and ring encapsulation can be mixed in various combinations. Neither of the rings is a closed loop. When any one node or link fails, a ring can restore connectivity to all remaining nodes within the 50 ms transfer time (the signaling time after detection).

Figure 21: Ring example



28781

Example:

```
configure
eth-ring 1
description "Ethernet Ring on Node B"
```

```


    revert-time 100
    guard-time 5
    ccm-hold-time down 100 up 200
    rpl-node owner
    path a 1/6/1 raps-tag 100
    # Control Channel Tag 100
    description "To A ring link"
    rpl-end
    eth-cfm
        mep 1 domain 1 association 1
        control-mep
        no shutdown
    exit
    exit
    no shutdown
exit
path b 1/6/2 raps-tag 100 # Control Channel Tag 100
description "to D Ring Link"
eth-cfm
    mep 1 domain 1 association 1
    control-mep
    no shutdown
    exit
    exit
    no shutdown
exit
    no shutdown
exit
service
    vpls 10 customer 1 create # Ring APS SAPs
    description "Ring Control VLAN ID 100"
    sap 1/6/1:100 eth-ring 1 create
    # TAG for the Control Path a
    no shutdown
    exit
    sap 1/6/2:100 eth-ring 1 create
    # TAG for the Control Path b
    no shutdown
    exit
no shutdown
vpls 40 customer 1 create
description "Data service over Ethernet Ring 1 VLAN ID
500"
    sap 1/1/1:100 create
    # Traffic from uW node
    no shutdown
    exit
    sap 1/6/1:500 eth-ring 1 create
    # TAG for the Data Channel Path a
    no shutdown
    exit
    sap 1/6/2:500 eth-ring 1 create
    # TAG for the Data Channel Path b
    no shutdown
    exit
no shutdown
exit

```

3.6.2 Ethernet ring sub-rings

Ethernet sub-rings offer a dual redundancy with interconnected rings. The router supports sub-rings connected to major rings and a sub-ring connected to an LDP-based VPLS for access ring support in VPLS networks. [Figure 22: G.8032 sub-ring](#) and [Figure 23: Sub-ring configuration example](#) illustrate a major ring and sub-ring scenario. In this scenario, any link can fail in either ring (ERP1 or ERP2) and each ring is protected. The sub-ring (ERP2) relies on the major ring (ERP1) as part of its protection for traffic from nodes C and D, which are configured as interconnection nodes.

Figure 22: G.8032 sub-ring




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

3.6.2.1 Virtual and non-virtual channels

The 7705 SAR supports sub-ring control communication using either a virtual channel or non-virtual channel. In virtual channel mode, a dedicated VLAN ID, other than the major ring RAPS control channel, is configured as a data instance on the major ring. This allows the sub-ring control messages and state machine logic to function similar to a major ring. In non-virtual channel mode, the sub-ring is only connected by the RAPS control channels on the sub-ring itself. This mode offers slightly less redundancy in the RAPS messaging than the virtual channel mode because sub-ring RAPS messages are not propagated across the major ring. When a non-virtual link is configured, the protocol allows RPL messages over the sub-ring blocked link. The following figure shows a sub-ring configuration using virtual and non-virtual channels.

Figure 23: Sub-ring configuration example



28784

Sub-ring configuration is similar to major ring configuration and consists of three parts:

- Ethernet ring instance configuration
- control VPLS configuration
- data VPLS configuration (data instance or data channel)

The Ethernet ring configuration of a sub-ring is tied to a major ring and only one path is allowed. A split horizon group is mandatory to ensure that sub-ring control messages from the major ring are only passed to the sub-ring control.

As with a major ring, CCMs and RAPS messages continue to be forwarded in the control VPLS even if the service or its SAPs are administratively shut down. The Ethernet ring instance can be shut down to stop the operation of the ring on a node.

The data VPLS can be configured on the major ring and shares the same VLAN ID (SAP encapsulation) on both the major ring and the sub-ring to keep data on the same VLAN ID everywhere. See the configuration example shown below. Like other services in the router, the encapsulation VLAN ID is controlled by SAP configuration and the association to the controlling ring is by the Ethernet ring ID.

The following CLI output is an example of a sub-ring configuration on Node C shown in the figure:

```
eth-ring 2
  description "Ethernet Sub Ring on Ring 1"
```

```

sub-ring virtual-link // Using a virtual link
interconnect ring-id 1 // Link to Major Ring 1
propagate-topology-change
exit
exit
path a 1/1/3 raps-tag 100 // Ring control uses VLAN ID 100
eth-cfm
    mep 9 domain 1 association 4
        ccm-enable
        control-mep
        no shutdown
    exit
exit
no shutdown
exit
no shutdown
exit

```

If the sub-ring had been configured as a non-virtual-link, the sub-ring configuration above and on all the other sub-ring nodes for this sub-ring would instead be:

```

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

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

# Data configuration for the Sub-Ring

vpls 11 customer 1 create
description "Data on VLAN ID 11 for Ring 1"
stp shutdown
sap 1/1/1:11 eth-ring 1 create // VLAN ID 11 used for ring
stp shutdown
exit
sap 1/1/4:11 eth-ring 1 create
stp shutdown
exit
sap 1/1/3:11 eth-ring 2 create // Sub-ring data
stp shutdown
exit
sap 3/2/1:1 create
description "Local Data SAP"
stp shutdown
no shutdown
exit

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

vpls 100 customer 1 create
description "Control VLAN ID 100 for Ring 2 Interconnection"

```


```

split-horizon-group "s1" create //Ring Split horizon Group
exit
stp shutdown
sap 1/1/1:100 split-horizon-group "s1" eth-ring 1 create
stp shutdown
exit
sap 1/1/4:100 split-horizon-group "s1" eth-ring 1 create
stp shutdown
exit
sap 1/1/3:100 eth-ring 2 create
stp shutdown
exit
no shutdown
exit

```


The 7705 SAR supports a special configuration of the non-virtual link sub-ring that can be homed to a VPLS service as illustrated in the following figure. This is an economical way to have a redundant ring connection to a VPLS service. This configuration is currently supported only for dot1q and qinq sub-rings and only on LDP-based VPLS. The primary application for this is access rings that require resiliency. This example in the figure shows a sub-ring at an interconnection node without a virtual channel and interconnected to a VPLS. A VPLS service 1 is used to terminate the ring control. The Ethernet ring data SAP appears in the associated LDP-based VPLS service 5.

Figure 24: Sub-ring homed to VPLS



Ethernet rings and sub-rings offer a way to build a scalable and resilient Ethernet transport network. The following figure illustrates a hierarchical ring network where dual-homed services are connected to an Ethernet ring network.

Figure 25: Multi-ring hierarchy



28786

Major rings are connected by sub-rings to the top level major ring. These sub-rings require a virtual channel and will not work with non-virtual channels. Ring flushing is contained to major rings or, in the case of a sub-ring link failure or node failure, to the sub-ring and any directly attached major rings.

The following CLI output is an example of a sub-ring configuration for VPLS PE1 in the figure:

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

# Configuration for the ring control interconnection termination:

vpls 1 customer 1 create
  description "Ring 1 Control termination"
  stp shutdown
  sap 1/1/3:1.1 eth-ring 1 create //path a control
  stp shutdown
exit
```

```

    no shutdown
    exit

# Configuration for the ring data into the LDP based VPLS Service

vpls 5 customer 1 create
    description "VPLS Service at PE1"
    stp
        no shutdown
    exit
    sap 1/1/3:2.2 eth-ring 1 create
        stp shutdown
    exit
    sap 1/1/5:1 create
    exit
    mesh-sdp 5001:5 create //sample LDP MPLS LSPs
    exit
    mesh-sdp 5005:5 create
    exit
    mesh-sdp 5006:5 create
    exit
    no shutdown
exit

```

3.6.2.1.1 LAG support

Ethernet rings support LAG on Ethernet ring SAPs. However, using LAG impacts the resiliency response time. In many cases, operators may achieve better resiliency response time and QoS by using multiple ring instances, each on a single link, instead of LAG on Ethernet rings. If a response time of less than 100 ms is not required, LAG is an acceptable option for Ethernet rings.

3.6.2.2 OAM considerations

Ethernet CFM is enabled by configuring MEPs on each individual path under an Ethernet ring. Only Down MEPs can be configured on each path. CCM sessions can also be enabled to monitor the liveness of the path using an interval of 10 ms or 100 ms. Different CCM intervals can be supported on path A and path B in an Ethernet ring. CFM is optional if the node supports LOS, which is controlled by configuring **no-ccm-enable**.

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

When an Ethernet ring is configured on two ports located on different cards, the SAP queues and virtual schedulers are created with the actual parameters on each card.

Ethernet ring CC messages that are transmitted over the SAP queues using the default egress QoS policy use network class (NC) as a forwarding class. If user traffic is assigned to the NC forwarding class, it will compete for bandwidth resources with the Ethernet CCMs and cause congested queues. This congestion can result in CCM loss that could lead to unnecessary switching of the Ethernet ring. To avoid potential congestion, configure different QoS policies that will control the amount of traffic assigned into the corresponding queue.

3.6.2.3 Support service and solution combinations

Ethernet rings are supported on VPLS and r-VPLS instances. The following considerations apply:

- Only ports in access or hybrid mode can be configured as Ethernet ring paths. The ring ports can be located on the same or different media adapter cards.
- Dot1q and qinq ports are supported as Ethernet ring path members.
- A mix of regular and multiple Ethernet ring SAPs and pseudowires can be configured in the same services.

3.7 QinQ support

This section contains the following topics:

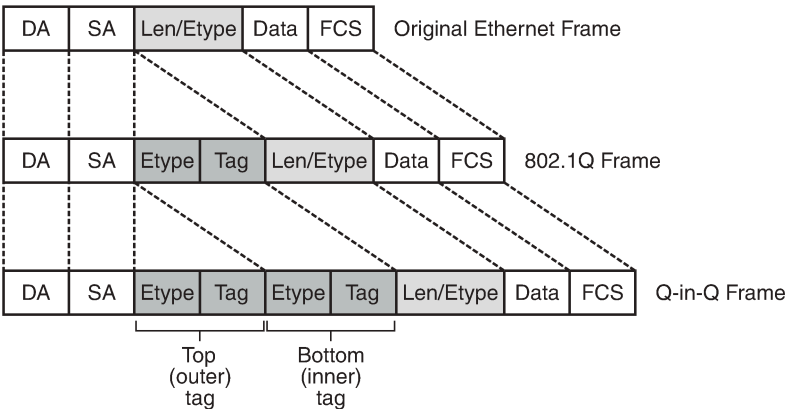
- [Overview of QinQ](#)
- [QinQ support with forced c-tag forwarding \(VPLS only\)](#)
- [QinQ support on Ethernet ports](#)
- [QinQ configuration overview](#)

3.7.1 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 the following figure. 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 26: QinQ frame



23834

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 lpipe)
- VPLS
- VPRN
- IES


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

- 6-port Ethernet 10Gbps Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card (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-Ax
- 7705 SAR-Wx (Ethernet ports)
- 7705 SAR-H
- 7705 SAR-Hc
- 7705 SAR-X

The following figure 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 27: QinQ tagging example



24021

3.7.2 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.

3.7.2.1 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.

3.7.2.2 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 may occur in the scenario shown in the following figure, where VPLS QinQ SAPs with **force-c-vlan-forwarding** enabled connect Customer 1 and Customer 2 and preserve the ingress VLAN tag.

Figure 28: QinQ using VPLS Ethernet SAPs



24017

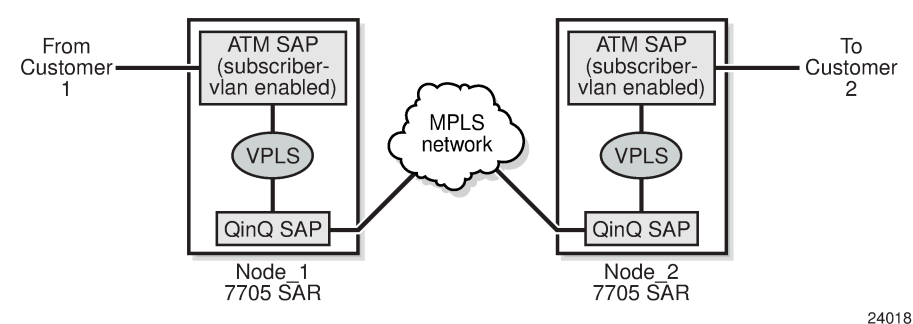
In the figure, the following events occur:

- 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)

3.7.2.3 Example 3: QinQ using VPLS with ATM SAPs

A second example of triple-tag behavior is shown in the following figure, where customers are connected using ATM SAPs with **subscriber-vlan** enabled.

Figure 29: QinQ using VPLS ATM SAPs



24018

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.

3.7.3 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.

3.7.3.1 Special QinQ SAP identifiers

The following table describes the special SAP identifiers used on the 7705 SAR. In the table, 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)	<i>port-id:qtag1.qtag2</i> ¹	1/1/1:20.2000	<i>qtag1</i> and <i>qtag2</i> range: 0 to 4094, and *
Null	<i>port-id</i> <i>port-id:0</i>	1/1/1 1/1/1:0	No suffix means no VLANs 0 suffix means SAP 0

QinQ SAP type	SAP identifier	Example	Notes
	<i>port-id:0.*</i> <i>port-id:qtag1.0</i> ²	1/1/1:0.*	0.* suffix means SAP 0 and any unused VLAN on the port
Default	<i>port-id:*</i> <i>port-id:.*</i> <i>port-id:qtag1.*</i>	1/1/1:.* 1/1/1:.*.* 1/1/1:40.*	*.* means any unused VLAN on the port ³

Notes:

1. *qtag1* is the top (outer) tag; *qtag2* is the bottom (inner) tag.
2. This configuration becomes available when the **new-qinq-untagged-sap** command is enabled.
3. 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 the table 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 or any frames with a VLAN tag of 0
 - example: SAP 1/1/1:0.* allows any untagged frames or frames with an outer tag of "0" only
If the outer tag is 10, the frame is dropped. If the outer tag is 0 and the inner tag is any valid VLAN ID, 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*)

3.7.3.2 QinQ dot1p match behavior

Because a qinq-encapsulated packet has top and bottom Q-tags, the user 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.

3.7.3.3 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.

3.7.3.4 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.

3.7.4 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).

3.8 Raw socket IP transport service

Serial data transport using raw sockets over IP transport services is a method of transporting serial data, in character form, over an IP network using Layer 3-based services. This feature can help transport Supervisory control and data acquisition (SCADA) data from remote terminal units (RTUs) to front-end processors (FEPs), or SCADA masters.

The functionality provided by the IP transport service feature for serial raw sockets is summarized as follows:

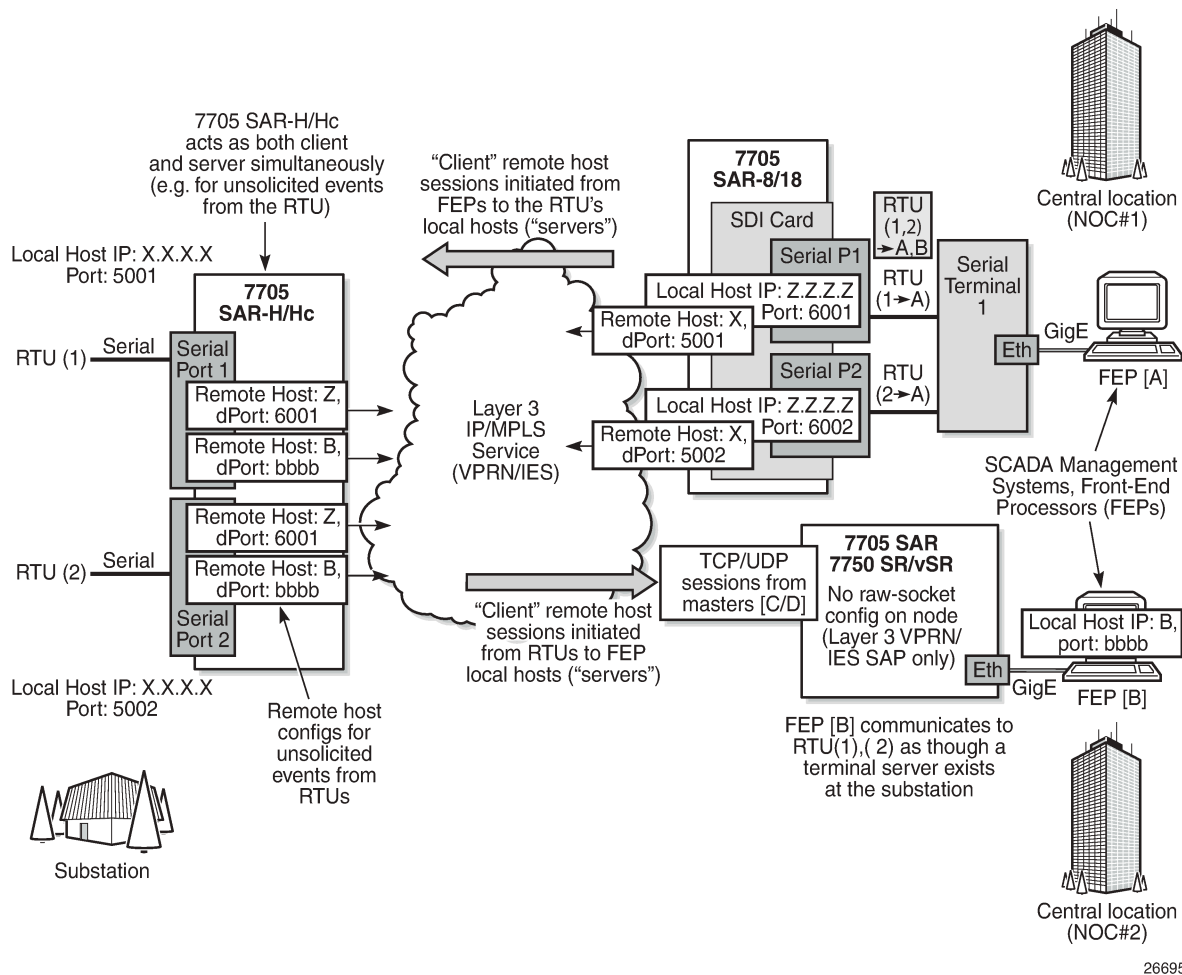
- IP transport local host (server) function, to listen to and open raw socket sessions from remote hosts
- IP transport remote host (client) function, to initiate and open new raw socket sessions to remote hosts
- both local host and remote host functions support for either TCP or UDP IP transport services
- IP transport over an IES or VPRN service
- enhanced QoS and queuing of sessions to ensure that collisions between sessions do not cause serial data to impact RTUs and end-user equipment

Figure 30: IP transport service illustrates a detailed view of the local host (server) and remote host (client) functionality that enables multiple communication streams to and from a serial port using raw socket IP transport.

The figure shows a three-node network: a 7705 SAR-H or 7705 SAR-Hc (left), a 7705 SAR-8 Shelf V2 or 7705 SAR-18 (top right) and a 7705 SAR-H/7705 SAR-Hc node, 7705 SAR-8 Shelf V2/7705 SAR-18 node, or 7750 SR/VSR node (bottom right). There are two devices, RTU (1) and RTU (2) connected to the serial ports on the 7705 SAR-H/7705 SAR-Hc. FEP server [A] can reach the RTUs via the socket sessions that originate from the 12-port Serial Data Interface card on the 7705 SAR-8 Shelf V2/7705 SAR-18 node. The bottom-right 7705 SAR or 7750 SR/VSR node is connected to FEP server [B] directly using Ethernet. This FEP server reaches the RTUs via a Layer 3 IP/MPLS service where raw socket sessions are processed directly on the FEP servers.

Through local host and remote host configurations on the 7705 SAR-H/7705 SAR-Hc or 7705 SAR-8 Shelf V2/7705 SAR-18, serial raw socket IP transport sessions are established to carry serial data over a wireless IP/MPLS network. The source and destination IP addresses and port numbers for these sessions are derived directly from the local/remote host configurations associated with each serial port or master head-end server.

Figure 30: IP transport service



The 7705 SAR-H/7705 SAR-Hc supports the ability to configure a raw socket IP transport interface for each serial port. This allows the raw socket IP transport to receive TCP or UDP session packets from multiple remote hosts when operating as a local host (server), or to create new multiple sessions to remote hosts to send and receive serial data when operating as a client.

There are two main configurations required for a serial raw socket IP transport service to be operational and to support the sending and receiving of serial data:

- **port-level configuration**

This includes configuring rudimentary serial link parameters such as baud rate, start/stop values, and bits. Socket-level configuration is also required, such as configuring end-of-packet checking parameters (idle-time, length, special character) and the inter-sessions delay for transmitting sessions data out the serial link. For information about the required port-level configuration, see the 7705 SAR Interface Configuration Guide, Configuration Command Reference chapter, “Serial Commands”.

- **IP transport service-level configuration**

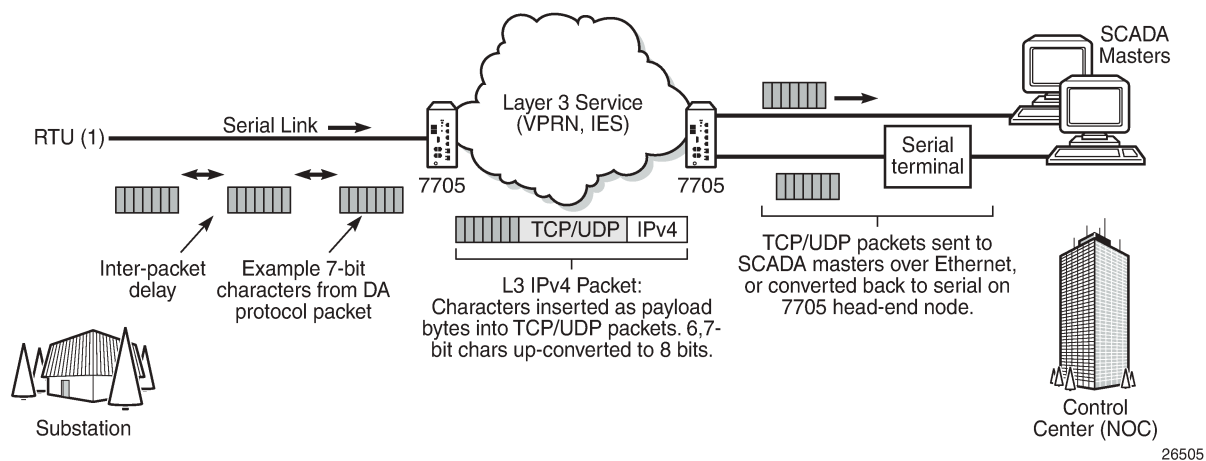
This includes creating an IP transport subservice to associate the serial port within a Layer 3 IES/VPRN service, so that TCP/UDP encapsulated serial data can be routed within the corresponding Layer 3

service. The IP transport subservice ID is modeled and created in the same way that the SAP IDs are created under the same service types. IP transport configuration includes configuring IP transport local host items and remote host items, such as setting TCP timers and sessions controls. See [IES command reference](#) and [VPRN services command reference](#) for the required commands.

The 7705 SAR-H/7705 SAR-Hc supports the configuration of a raw socket IP transport service for each serial port. This allows each serial port's local host to listen to and open raw socket sessions from remote hosts that need to communicate over the serial port, and for each serial port's local host to initiate and open raw socket sessions to remote hosts when serial data needs to be sent to those remote hosts. The local and remote host functions support TCP or UDP sessions (but not both concurrently) over the IES/VPRN service.

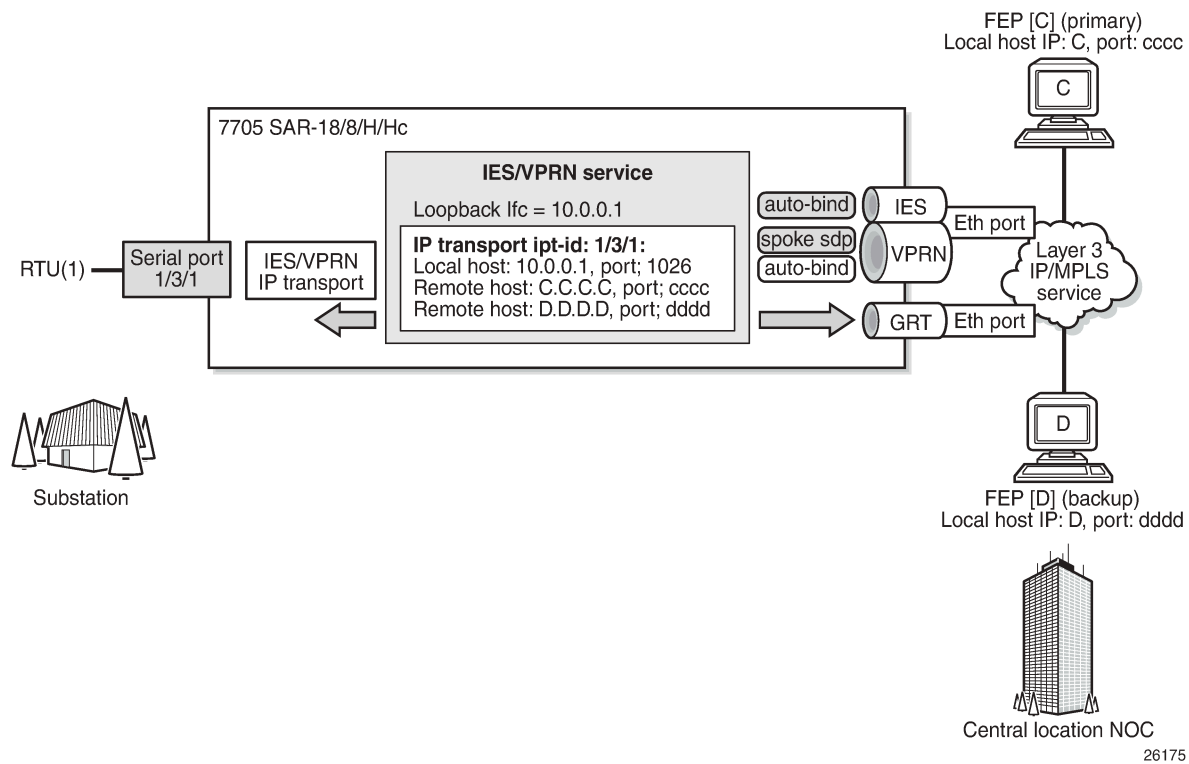
The serial data is received as characters that represent bytes in a packet. These bytes are packetized into Layer 3 TCP/UDP packets that are then transported or forwarded across the IP/MPLS network using the node's Layer 3 IES/VPRN service constructs for routing. The following figure illustrates how serial data is encapsulated into TCP/UDP packets and transported over IP/MPLS.

Figure 31: TCP/UDP packet transport over IP/MPLS



For raw socket packets to be routed within an IES/VPRN service, an IP transport subservice must be configured within an IES/VPRN context. The IP transport subservice context is where users configure local host and remote host information, such as IP addresses and ports for establishing TCP/UDP sessions, and other per-session parameters. TCP/UDP encapsulated serial data is routed within the corresponding Layer 3 IES/VPRN service. The following figure illustrates this concept.

Figure 32: IES/VRN IP transport service



To create an IP transport subservice, the **ip-transport** command is used with the corresponding serial port as the *ipt-id* to bind the serial port SAP to the IP transport subservice. After the IP transport service is created, local host and remote host configurations can proceed. A local host must be configured before remote hosts can be configured.

Each local host uses a local address (from a loopback or local interface configured under the IES/VRN service context) as the local host IP address of the IP transport subservice associated with the serial port. The local host IP address is the source IP address in the raw socket packets leaving the node within the IES/VRN service. The local host is used to terminate TCP/UDP sessions from remote hosts. The local host can select either the TCP or UDP protocol for raw socket sessions but not both concurrently.

Multiple remote hosts can be configured under the IP transport subservice associated with the serial port so that each remote host receives the serial data that is received on the serial port. Each remote host has its own remote destination IP address and port value for establishing sessions. The configured remote hosts use the TCP or UDP protocol configured for the IP transport subservice.



Note: It is not necessary to configure remote hosts when the IP transport service is not originating sessions. If sessions are only established toward the IP transport local host (for example, remote servers polling the local host), the remote host configuration is not necessary. Remote host configurations may still be desirable when using the **filter-unknown-host** command.

IP transport processing of TCP/UDP packets occurs on the CSM of the 7705 SAR-H/7705 SAR-Hc. Filters configured for protecting the CSM must take into account the raw socket IP transport packets and ensure that the filter is not blocking associated IP transport sessions. For example, operators must ensure that

interface IP addresses and ports configured on the node are not blocked and that remote host IP/port combinations are not blocked.

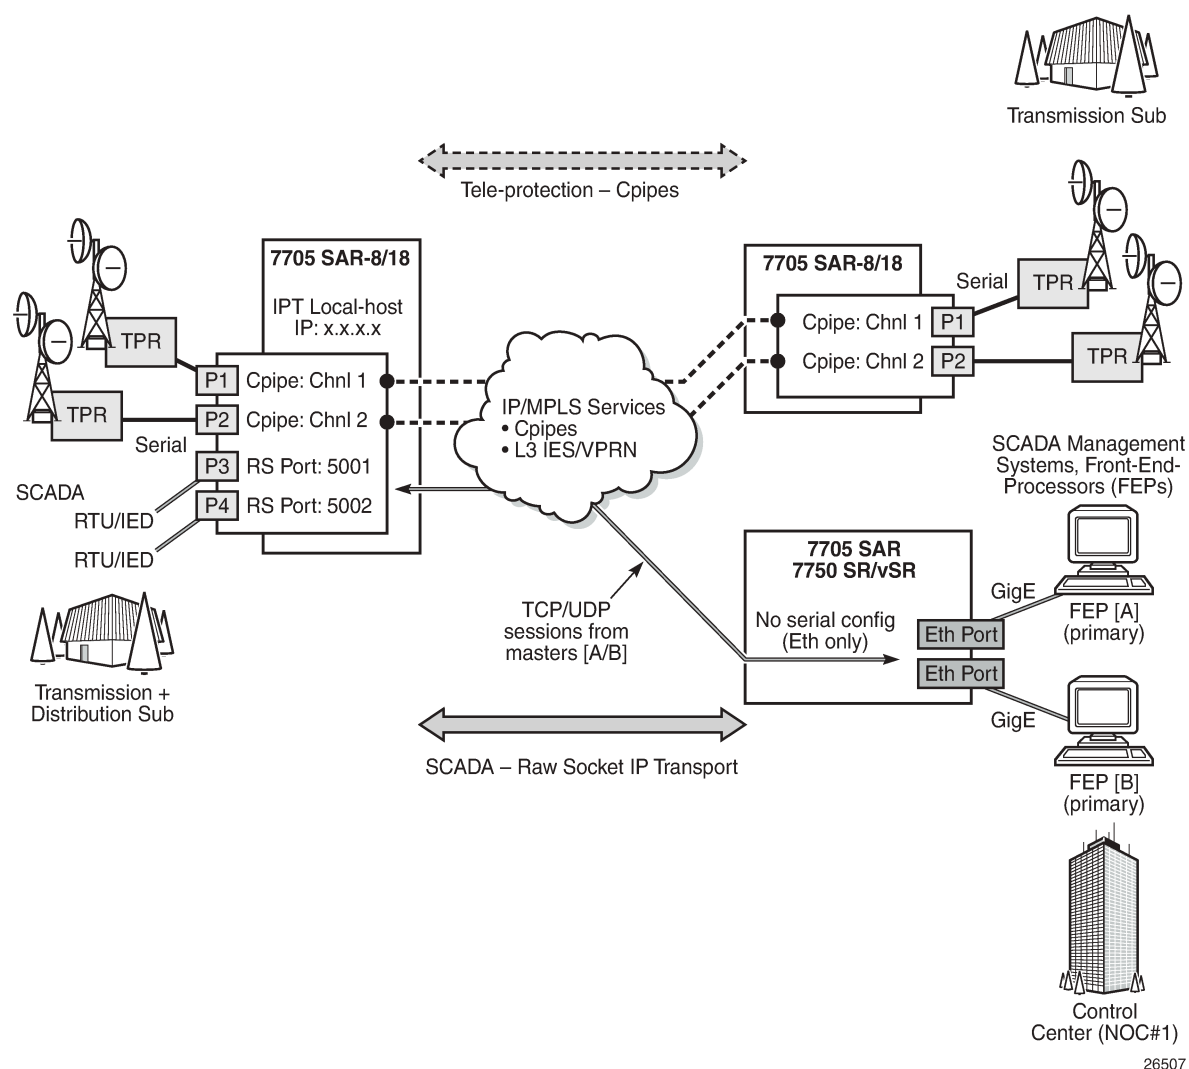
For IES/VRN IP transport services, all tunnel types supported by the IES/VRN service are also supported for the IP transport service. This includes all types of MPLS tunnels (such as RSVP-TE, LDP, autobind, static LSP) and GRE tunnels.



Note: IP transport-to-IP transport raw socket data on the same node is not supported. If serial-to-serial communication is needed on the same node, customers must use Cpipes.

The 7705 SAR supports the concurrent operation of raw sockets and Cpipes on the 12-port Serial Data Interface card, version 2 and version 3, as shown in the following figure.

Figure 33: Raw socket and Cpipe support on the 7705 SAR



3.8.1 Remote host manual TCP connection check

A manual TCP connection check can be performed for each remote host configured for a raw socket IP transport subservice. When executed by an operator, the TCP connection check attempts to establish a TCP session toward the configured remote host. Only one TCP connection check attempt is made, with a fixed timeout of 5 s. If the attempt is successful, the session is torn down immediately without data being sent.

The TCP connection check is initiated in the CLI using the **tools>perform>service>id>ip-transport>remote-host>check-tcp** command. The result is displayed in the CLI using the **tools>dump>service>id>ip-transport>remote-host>check-tcp** command. Equivalent management is available via SNMP.

If a TCP connection to a remote host already exists because of serial traffic being transmitted, the check returns "successful" without impacting the existing TCP connection.

3.8.2 QoS requirements for IP transport

Serial raw socket data that is transported using an IP transport service can be DSCP marked and FC classified at the source node. This allows the source node (local host) of the traffic to mark packets correctly so that downstream nodes prioritize them as needed, and to queue local traffic in the right egress queue based on the classification assigned to the IP transport service.

Additionally, the DSCP setting is assigned per IP transport subservice for all traffic from the local host and all traffic destined for each remote host. The DSCP setting is not set per remote host.

See the **dscp** and **fc** commands under [IES raw socket IP transport configuration commands](#) and [VPRN raw socket IP transport configuration commands](#) for more information about configuring the QoS settings for an IES or VPRN IP transport subservice.

3.9 Service creation overview

[Figure 34: Service creation and implementation flowchart](#) 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 before subscriber services tasks.

Before starting the process shown in the figure, ensure that the 7705 SAR system has been configured with an IP address and (for the 7705 SAR-8 Shelf V2 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 Lpipe), IES, VPLS, or VPRN services

- configure SAPs
- bind SDPs
- create exclusive QoS policies
- (optionally) assign IP filter policies to Epipe SAPs, Lpipe 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 34: Service creation and implementation flowchart



21829

3.10 Port and SAP CLI identifiers

When entering text in the CLI, *port-id* is often displayed to indicate that a port identifier may need to be entered in the command line. Similarly, to identify a SAP, the *port-id* is used, but more 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/VC1 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 a 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.

3.11 Configuring global service entities with CLI

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

Topics in this section include:

- [Service model entities](#)
- [Basic configuration](#)
- [Common configuration tasks](#)
- [ETH-CFM \(802.1ag and Y.1731\) tasks](#)
- [Service management tasks](#)

3.12 Service model entities

The Nokia 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 Lpipe SAP parameters](#)
- [Configuring a VPLS SAP](#)
- [Configuring IES SAP parameters](#)
- [Configuring VPRN interface SAP parameters](#)

3.13 Basic configuration

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

- 7705 SAR Quality of Service Guide
- 7705 SAR Router Configuration Guide
- 7705 SAR System Management Guide
- 7705 SAR 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
...
...
    pipe 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#

```

3.14 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](#)

3.14.1 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 "Nokia 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 "Nokia Customer"
    phone "650 555-5100"
exit
```

3.14.2 Configuring SDPs

When configuring an SDP, consider the following points:

- SDPs can be configured for MPLS, GRE, or IP encapsulation.
- If you do not specify an encapsulation type, the default is MPLS.
- 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.
- By default, SDPs are not associated with services. After an SDP is created, services can be associated with that SDP.
- A distributed service must be bound to an SDP.
- 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.
- If an SDP configuration does not include the IP address of the associated far-end router, VLL and VPLS services to the far-end router cannot be provided.
- Up to eight RSVP-TE or SR-TE LSPs can be configured under a single MPLS-encapsulated SDP. However, a mix of RSVP-TE and SR-TE LSPs is not supported. RSVP-TE LSPs are configured using the **config>service>sdp>lsp** command and SR-TE LSPs are configured using the **config>service>sdp>sr-te-lsp** command.
- For MPLS-encapsulated SDPs, LSPs must be configured before the LSP-to-SDP associations can be assigned. The LSP-to-SDP associations must be created explicitly.
- LSPs are configured in the **config>router>mpls** context. See the 7705 SAR MPLS Guide for configuration and command information.
- The destination address of the LSPs must match the far-end IP address of the SDP.

- The far-end SDP IP address can be the system IP address of a 7705 SAR or an SR-series router, or loopback or network interface of the far-end router
- When configuring MPLS SDP parameters, you can either specify up to eight RSVP-TE LSPs using the **config>service>sdp>lsp** command or enable LDP using the **config>service>sdp>ldp** command. There cannot be two methods of transport in a single SDP unless the **mixed-lsp-mode** option is selected (the **lsp** and **ldp** commands are mutually exclusive).

If **mixed-lsp-mode** is enabled and an LSP is specified, RSVP-TE is used for dynamic signaling in the LSP.

- 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, ingress and egress vc-labels for the services using that SDP must be configured manually.

For a basic SDP configuration, perform the following steps:

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

The following examples show the CLI syntax for a basic MPLS SDP configuration. The first two show the syntax for configuring the SDP without **mixed-lsp-mode** enabled; one shows an RSVP-TE LSP configuration and the other shows an LDP configuration. The third example shows the syntax for configuring the SDP with **mixed-lsp-mode** enabled, with both an RSVP-TE LSP and LDP configured.

CLI syntax:

```
config>service>sdp sdp-id mpls create
    far-end ip-addr
    lsp lsp-name
    keep-alive
        shutdown
    no shutdown
```

Example:

```
config>service# sdp 100 mpls create
config>service>sdp# far-end 10.10.10.10
config>service>sdp# lsp "LSP1"
config>service>sdp# no shutdown
config>service>sdp# exit
```

The following example displays a basic SDP LSP configuration output.

```
*A: Sar18 Dut-B>config>service# info
-----
.....
    sdp 100 create
        no shutdown
        far-end 10.10.10.10
        lsp "LSP1"
        keep-alive
            shutdown
        exit
    exit
.....
```

CLI syntax:

```
config>service>sdp sdp-id mpls create
    far-end ip-addr
    ldp
    keep-alive
        shutdown
    no shutdown
```

Example:

```
config>service# sdp 300 mpls create
config>service>sdp# far-end 10.10.10.10
config>service>sdp# ldp
config>service>sdp# no shutdown
config>service>sdp# exit
```

The following example displays a basic SDP LDP configuration output.

```
*A: Sar18 Dut-B>config>service# info
-----
.....
    sdp 300 create
        far-end 10.10.10.10
        ldp
        keep-alive
            shutdown
        exit
        no shutdown
    exit
.....
-----
```

CLI syntax:

```
config>service>sdp sdp-id mpls create
    description description-string
    far-end ip-addr
    mixed-lsp-mode
    exit
    ldp
    lsp lsp-name
    keep-alive
        shutdown
    no shutdown
```

Example:

```
config>service# sdp 1 mpls create
config>service>sdp# description "SDI4"
config>service>sdp# far-end 10.10.10.10
config>service>sdp# mixed-lsp-mode
config>service>sdp# exit
config>service>sdp# ldp
config>service>sdp# lsp "LSP1"
config>service>sdp# no shutdown
config>service>sdp# exit
```

The following example displays a basic SDP **mixed-lsp-mode** configuration output.

```
A:Sar18 Dut-B>config>service# info
```

```
-----
.....
      sdp 1 create
        description "SDI4"
        far-end 10.10.10.10
        mixed-lsp-mode
        exit
        ldp
        lsp "LSP1"
        keep-alive
          shutdown
        exit
        no shutdown
      exit
.....
-----
```

The following examples show the CLI syntax for a basic GRE SDP configuration.

CLI syntax:

```
config>service>sdp sdp-id gre create
  description description-string
  far-end ip-addr
  keep-alive
    shutdown
  no shutdown
```

Example:

```
config>service# sdp 2 gre create
config>service>sdp# description "GRE-10.10.10.10"
config>service>sdp# far-end 10.10.10.10
config>service>sdp# no shutdown
config>service>sdp# exit
```

The following example displays a basic GRE SDP 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
      exit
.....
-----
A:ALU-12>config>service#
```

3.14.2.1 Enabling IP fragmentation for GRE SDPs

Use the following CLI syntax to enable fragmentation of IP packets for GRE SDPs.

CLI syntax:

```
config>service>sdp sdp-id gre [create]
allow-fragmentation
```

Example:

```
config>service# sdp 2 gre create
config>service>sdp# allow-fragmentation
```



Note: Fragmented IP packets require a reassembly profile to ensure that packets that cannot be reassembled are disposed of in a timely manner. See the 7705 SAR Router Configuration Guide for configuration and command information.

3.14.3 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#
```

3.15 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](#)

3.15.1 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 (global or Epipe or VPLS service). For information about Ethernet OAM commands for 802.1ag and Y.1731 OAM, see the "Ethernet OAM capabilities" section in the 7705 SAR 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 or VPLS 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 2 customer 1 create
*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
-----

*A:ALU-1>config>service>vpls 2 customer 1 create
*A:ALU-1>config>service>vpls# info
-----
shutdown
sap 1/2/1:0 create
  eth-cfm
    mep 1 domain 1 association 1 direction down
```

```

        shutdown
    exit
exit
mesh-sdp 7:2 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 or VPLS 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
...
-----

*A:ALU-1>config>service>vpls# info
-----
    shutdown
    sap 1/2/1:0 create
        eth-cfm
            mep 1 domain 1 association 1 direction up
                ais-enable
                    interval 60
                    priority 2

```



```

        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.

3.15.2 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
- VPLS SAP
- VPLS spoke or mesh SDP
- OAM tests (loopback, linktrace, Ethernet test, delay measurement, loss measurement, synthetic loss measurement)

The MAC address for a MEP on an Epipe SAP or on an Epipe or VPLS SDP cannot be changed. For a MEP on an Epipe SAP, the MAC address is the port MAC address. For a MEP on an Epipe or VPLS SDP, the MAC address is the system MAC address. The MAC address for a MEP on a VPLS SAP can be changed; the default is the port MAC address.

The 7705 SAR supports the following MEPs for both 802.1ag and Y.1731:

- SAP Up and Down MEPs
- spoke SDP Up and Down MEPs
- mesh SDP Up and Down MEPs (VPLS only)

The following two syntax examples are for an Epipe service. Configuration for VPLS is the same except that the **hold-mep-up-on-failure** and **dual-ended-loss-test-enable** parameters are not supported on VPLS SAPs.

The third syntax shows the OAM tests that can be applied to MEPs.

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}]
        ais-enable
        ccm-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
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
oam
    eth-cfm eth-test mac-address mep mep-id domain md-index
association ma-index [priority priority] [data-length data-length]
    eth-cfm linktrace mac-address mep mep-id domain md-index
association ma-index [ttl tll-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 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 interval] [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]

```

3.16 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](#)

3.16.1 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"
```

3.16.2 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
```

3.16.3 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: After 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
```

3.16.4 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
```

3.16.5 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
```

3.17 Global service command reference

3.17.1 Command hierarchies

- [Global service configuration commands](#)
 - [Customer commands](#)
 - [SDP commands](#)
 - [SAP commands](#)
 - [Ethernet ring commands](#)
 - [ETH-CFM commands](#)
- [Show commands](#)
- [Monitor commands](#)

3.17.1.1 Global service configuration commands

3.17.1.1.1 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
```

3.17.1.1.2 SDP commands

```
config
- service
- sdp sdp-id [gre | mpls | ip] [create]
- no sdp sdp-id
- [no] adv-mtu-override
- [no] allow-fragmentation
- [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 | ipv6-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
- [no] mixed-lsp-mode
  - revert-time {revert-time | infinite}
  - no revert-time
- path-mtu bytes
- no path-mtu
- [no] shutdown
- signaling {off | tldp}
- [no] sr-isis
- [no] sr-ospf
- [no] sr-te-lsp lsp-name
- vlan-vc-etype 0x0600..0xffff
- no vlan-vc-etype [x0600.0xffff]
- [no] weighted-ecmp

```

3.17.1.1.3 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
- system

```

- ethernet
 - [no] new-qinq-untagged-sap

3.17.1.1.4 Ethernet ring commands

```

config
- [no] eth-ring ring-index
- ccm-hold-time [down down-timeout] [up up-timeout]
- no ccm-hold-time
- compatible-version version
- no compatible-version
- description description-string
- no description
- guard-time time
- no guard-time
- node-id xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
- no node-id
- path {a | b} [{port-id | lag-id | mw-link-id} raps-tag qtag1[.qtag2]]
- no path {a | b}
  - description description-string
  - no description
  - eth-cfm
    - [no] mep mep-id domain md-index association ma-index
      - [no] ccm-enable
      - ccm-ltm-priority priority
      - no ccm-ltm-priority
      - [no] control-mep
      - description description-string
      - no description
      - no eth-test-enable
        - bit-error-threshold bit-errors
        - test-pattern {all-zeros | all-ones} [crc-enable]
        - no test-pattern
      - low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon |
xcon | noXcon}
        - one-way-delay-threshold seconds
        - [no] shutdown
    - [no] rpl-end
    - [no] shutdown
  - revert-time time
  - no revert-time
  - rpl-node {owner | nbr}
  - no rpl-node
  - [no] shutdown
  - [no] sub-ring {virtual-link | non-virtual-link}
    - interconnect ring-id ring-index
    - interconnect vpls
    - [no] interconnect
      - [no] propagate-topology-change

```

3.17.1.1.5 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 about configuring ETH-CFM commands, see the 7705 SAR OAM and Diagnostics Guide.

3.17.1.2 Show commands

```

show
- eth-ring [status]
- eth-ring [ring-index] hierarchy
- eth-ring ring-index [path {a | b}]
- service
  - customer customer-id
  - sdp sdp-id keep-alive-history
  - sdp [sdp-id] [detail]
  - sdp far-end ip-addr keep-alive-history
  - sdp far-end ip-addr [detail]
  - sdp-using sdp-id[:vc-id] | far-end ip-address]
  - service-using [epipe] [fpipe] [hpipe] [ies] [vprn] [mirror] [vpls] [apipe] [cpipe]
[sdp sdp-id] [customer customer-id]

```

3.17.1.3 Monitor commands

```

monitor
- service
- id service-id
  - sap sap-id [interval seconds] [repeat repeat] [absolute | rate]
  - sap-aggregation-group group-id [interval seconds] [repeat repeat] [absolute | rate]

```


3.17.2 Command descriptions

- [Global service configuration commands](#)
- [Show commands](#)
- [Monitor commands](#)

3.17.2.1 Global service configuration commands

- [Generic commands](#)
- [Customer commands](#)
- [SDP commands](#)
- [SDP keepalive commands](#)
- [Ethernet ring commands](#)

3.17.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>service>customer

config>service>sdp

config>eth-ring

config>eth-ring>path

config>eth-ring>path>eth-cfm>mep

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 (such as #, \$, or spaces), 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.

new-qinq-untagged-sap

Syntax

[no] new-qinq-untagged-sap

Context

config>system>ethernet

Description

This command controls the behavior of qinq SAP *y.0* (for example, 1/1/1:3000.0). If this command is enabled, the *y.0* SAP maps all ingress frames tagged with outer tag VLAN ID of *y* (qinq Ethertype) and no inner tag, or with an inner tag VLAN ID of zero (0). This behavior applies to all existing and future *y.0* SAPs.

The **no** form of this command disables qinq untagged SAP, and the *y.0* SAP will work like a *y.** SAP (for example, 1/1/1:3000.*); all frames tagged with outer VLAN *y* and no inner VLANs, or inner VLAN *x*, where inner VLAN *x* is not specified in a SAP *y.x* configured on the same port (for example, 1/1/1:3000.10).

Default

no new-qinq-untagged-sap. This setting ensures that there will be no disruption for existing usage of this SAP type.

3.17.2.1.2 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 (such as #, \$, or spaces), 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 (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

3.17.2.1.3 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, the 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. Segment routing (SR) is another MPLS tunnel type and is used to allow service binding to an SR tunnel programmed in the tunnel table manager (TTM) by OSPF or IS-IS. An SDP of type **sr-isis** or **sr-ospf** can be used with the far-end option.

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**). After 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 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 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

allow-fragmentation

Syntax

[no] **allow-fragmentation**

Context

config>service>sdp

Description

This command enables GRE-encapsulated packets transmitted from the SDP to be fragmented if their size exceeds the configured network port MTU.

When **allow-fragmentation** is enabled, GRE-encapsulated packets that are larger than the network port MTU are fragmented and the DF bit is set to 0 by the router. Packets that are smaller than the port MTU are left unfragmented, the DF bit is set to 1 by the router, and the identification number of the packet is set to 0.

Fragmentation is supported on an SDP that has NGE (**encryption-keygroup**) enabled. To determine if an encrypted packet needs to be fragmented, the system compares the total packet size after NGE encryption to the network port MTU. If the encrypted packet size is larger than the MTU, the packet is fragmented. NGE decryption is performed after the packet is fully reassembled.

The **no** form of the command disables fragmentation of oversize GRE-encapsulated packets transmitted from the SDP.

Default

no allow-fragmentation

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, BGP-tunnel, SR-ISIS, SR-OSPF, or SR-TE-LSP. This restriction is relaxed for some combinations of the transport methods when the mixed-LSP mode option is enabled on the SDP.

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*|*ipv6-address*]

no far-end

Context

config>service>sdp

Description

This command configures the system IP address of the far-end destination 7705 SAR, 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, 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 may 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* and *ipv6-address* parameters are 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 IPv4 address of the far-end destination router for the SDP

ipv6-address

the IPv6 address of the far-end destination router for the SDP

ldp

Syntax

[no] ldp

Context

config>service>sdp

Description

This command enables LDP-signaled LSPs on MPLS-encapsulated SDPs.

In MPLS SDP configurations, up to eight RSVP-TE LSPs can be specified or LDP can be enabled. The SDP **ldp** and **lsp** commands are mutually exclusive unless mixed-LSP SDP is enabled with the [mixed-lsp-mode](#) command.

If mixed-LSP SDP is not enabled and an LSP is specified on an MPLS SDP, 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.

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-encapsulated SDPs. Up to eight RSVP-TE LSPs can be associated with a single SDP. The LSPs must have already been created in the **config>router>mpls** context with a valid far-end IP address. See the 7705 SAR MPLS Guide for CLI syntax and command usage.

In MPLS SDP configurations, LSPs can be specified or LDP can be enabled. The SDP **ldp** and **lsp** commands are mutually exclusive unless mixed-LSP SDP is enabled with the [mixed-lsp-mode](#) command.

If mixed-LSP SDP is not enabled and LDP is already enabled on an MPLS SDP, an LSP cannot be specified on the SDP. To specify an LSP on the SDP, LDP must be disabled.

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* can 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

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 tiebreaker by tunnel table manager users to select the route with the lower value.

Parameters

metric

specifies the SDP metric

Values 1 to 17407

mixed-lsp-mode

Syntax

[no] mixed-lsp-mode

Context

config>service>sdp

Description

This command enables the mixed-LSP mode of operation on an SDP, which allows a primary LSP type and a backup LSP type in the same SDP configuration. For example, RSVP-TE LSPs and LDP LSPs can be configured on an SDP with the **lsp** and **ldp** commands. If mixed-LSP mode is not enabled, these commands are mutually exclusive.

Two combinations of mixed LSPs are possible:

- an RSVP-TE primary LSP backed up by an LDP LSP
- an LDP primary LSP backed up by a BGP LSP



Note: Mixed-LSP SDPs do not support static LSPs on either the primary or backup.

The **no** form of this command disables mixed-LSP; however you must remove one of the LSP types from the SDP configuration first or the command will fail.

revert-time

Syntax

revert-time {*revert-time* | **infinite**}

no revert-time

Context

config>service>sdp>mixed-lsp-mode

Description

This command configures the length of time that the service manager must wait before it resets the SDP configuration to a higher-priority LSP type when one becomes available.

The **no** form of the command resets the timer to the default value of 0. This means that the service manager resets the SDP configuration immediately to a higher-priority LSP type when one becomes available.

Default

0

Parameters

seconds

the time, in seconds, to wait before reverting to a higher-priority LSP type when one becomes available.

Values 0 to 600

infinite

the service manager never resets the SDP configuration to the highest-priority LSP type unless the currently active LSP fails

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

sr-isis

Syntax

[**no**] **sr-isis**

Context

config>service>sdp

Description

This command configures an MPLS SDP of LSP type IS-IS segment routing. The SDP of LSP type **sr-isis** can be used with the **far-end** option. The **signaling** protocol for the service labels for an SDP using an SR tunnel can be configured to static (**off**) or T-LDP (**tldp**).

Default

no sr-isis

sr-ospf

Syntax

[no] sr-ospf

Context

config>service>sdp

Description

This command configures an MPLS SDP of LSP type OSPF segment routing. The SDP of LSP type **sr-ospf** can be used with the **far-end** option. The **signaling** protocol for the service labels for an SDP using an SR tunnel can be configured to static (**off**) or T-LDP (**tldp**).

Default

no sr-ospf

sr-te-lsp

Syntax

[no] sr-te-lsp *lsp-name*

Context

config>service>sdp

Description

This command configures an MPLS SDP of LSP type SR-TE. Up to eight SR-TE LSPs can be configured under an SDP.

The **mixed-lsp-mode** option does not support the **sr-te** tunnel type.

The **signaling** protocol for the service labels for an SDP using an SR-TE LSP can be configured to static (**off**) or T-LDP (**tldp**).

Default

none

Parameters

lsp-name

the name of an LSP that has already been created

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

weighted-ecmp

Syntax

[no] **weighted-ecmp**

Context

config>service>sdp

Description

This command enables weighted ECMP for IES or VPRN Layer 3 spoke SDP interfaces. This command is applicable when the SDP has RSVP-TE LSPs configured using the **config>service>sdp>lsp** command or SR-TE LSPs configured using the **config>service>sdp>sr-te-lsp** command.

When weighted ECMP is enabled on an SDP, a path is selected based on the configured hash. Paths are then load-balanced across the LSPs used by the SDP according to the normalized LSP load-balancing weight configured using the **load-balancing-weight** command described in the 7705 SAR MPLS Guide, "MPLS Commands". This means that consecutive packets of a particular service use the same LSP, but the overall load handled by LSPs to the SDP far end is balanced according to the load-balancing weight if all services using the SDP send the same bandwidth and there are more services using the SDP than there are LSPs for the SDP.

If an LSP in the ECMP set has no load-balancing weight configured, then ECMP is applied to packets based on the output of the hash for the service ID.

The **no** form of the command disables weighted ECMP for the SDP.

Default

no weighted-ecmp

3.17.2.1.4 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. After 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.

The following table describes the 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 <i>orig-sdp-id</i> ¹	Orig-SDP Non-Existent	Down
Keepalive request not sent due to administratively down <i>orig-sdp-id</i>	Orig-SDP Admin-Down	Down
Keepalive reply received, invalid origination-id	Far End: Originator-ID Invalid	Down
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 resets 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 resets 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 resets 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 resets 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

3.17.2.1.5 Ethernet ring commands

eth-ring

Syntax

[no] eth-ring *ring-index*

Context

config

Description

This command configures a G.8032 Ethernet ring. Ethernet rings may be configured as major rings with two paths (A and B), as sub-rings with two paths, or in the case of an interconnection node, a single path.

The **no** form of this command deletes the specified Ethernet ring.

Default

no eth-ring

Parameters

ring-index

specifies the Ethernet ring ID

Values 1 to 128

ccm-hold-time

Syntax

ccm-hold-time [*down down-timeout*] [*up up-timeout*]
no ccm-hold-time

Context

config>eth-ring

Description

This command configures Ethernet ring dampening timers.
The **no** form of the command sets the up and down timers to the default values.

Default

down 0
up 20

Parameters

down-timeout
specifies the down timeout, in centiseconds

Values 0 to 5000
Default 0

up-timeout
specifies the hold-time for reporting the recovery, in deciseconds

Values 0 to 5000
Default 20

compatible-version

Syntax

compatible-version *version*
no compatible-version

Context

config>eth-ring

Description

This command configures Ethernet ring compatibility version for the G.8032 state machine and messages. The default is version 2 and all router switches use version 2. The version can be changed if there is a need to interwork with third party devices that only support version 1.

The **no** form of this command sets the compatibility version to 2.

Default

2

Parameters

version
specifies the version of the G.8032 state machine

Values 1 or 2

guard-time

Syntax

guard-time *time*
no guard-time

Context

config>eth-ring

Description

This command configures the guard time for an Ethernet ring.

The **no** form of this command restores the default guard time.

Default

5

Parameters

value
specifies the guard time, in deciseconds

Values 1 to 20

Default 5

node-id

Syntax

node-id *xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx*

no node-id

Context

config>eth-ring

Description

This optional command configures the MAC address of the RPL link. The default is to use the chassis MAC address for the ring control. This command allows the chassis MAC address to be overridden with another MAC address.

The **no** form of the command removes the RPL link.

Default

no node-id

Parameters

xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

specifies the MAC address

path

Syntax

path {a | b} [{port-id | lag-id| mw-link-id} **raps-tag** qtag1[.qtag2]]

no path {a | b}

Context

config>eth-ring

Description

This command assigns the Ethernet ring (major or sub-ring) path to a port and defines the Ethernet ring APS tag. Rings typically have two paths, A or B.

The **no** form of this command removes the configured path.

Default

no path

Parameters

- port-id*
the port ID in the *slot/mda/port* format
- lag-id*
the LAG identifier in the form **lag-lag-id**
Values 1 to 32
- mw-link-id*
the microwave link ID number in the form **mw-link-num**
Values 1 to 24
- raps-tag**
the Ethernet ring member's encapsulation
- qtag1*
specifies the top or outer VLAN ID
Values 1 to 4094
- qtag2*
specifies the bottom or inner VLAN ID
Values 1 to 4094

eth-cfm

Syntax

eth-cfm

Context

config>eth-ring>path

Description

This command enables the context to configure Ethernet CFM parameters.

mep

Syntax

[no] **mep** *mep-id* **domain** *md-index* **association** *ma-index*

Context

config>eth-ring>path>eth-cfm

Description

This command provisions an 802.1ag maintenance endpoint (MEP).
The **no** form of the command deletes the MEP.

Parameters

- mep-id*
specifies the MEP identifier
Values 1 to 81921
- md-index*
specifies the MD index value
Values 1 to 4294967295
- ma-index*
specifies the MA index value
Values 1 to 4294967295

ccm-enable

Syntax

[no] ccm-enable

Context

config>eth-ring>path>eth-cfm>mep

Description

This command enables the generation of CCM messages.
The **no** form of the command disables the generation of CCM messages.

Default

no ccm-enable

ccm-ltm-priority

Syntax

ccm-ltm-priority *priority*
no ccm-ltm-priority

Context

config>eth-ring>path>eth-cfm>mep

Description

This command specifies the priority value for CCMs and LTMs transmitted by the MEP.

The **no** form of the command removes the priority value from the configuration.

Default

The highest priority on the bridge port.

Parameters

priority

specifies the priority of CCM and LTM messages

Values 0 to 7

control-mep

Syntax

[no] control-mep

Context

config>eth-ring>path>eth-cfm>mep

Description

This command enables Ethernet ring control on the MEP. Configuration of this command is mandatory for a ring. MEP detection of failure using CCM can be enabled or disabled independently of the control MEP.

The **no** form of this command disables Ethernet ring control.

Default

no control-mep

eth-test-enable

Syntax

[no] eth-test-enable

Context

config>eth-ring>path>eth-cfm>mep

Description

This command enables Ethernet test functionality on the MEP. For the test to work, operators need to configure Ethernet test parameters on both the sender and receiver nodes. The Ethernet test can be performed using the **oam>eth-cfm>eth-test** command.

A check is done for both the provisioning and the test to ensure the MEP is a Y.1731 MEP, provisioned with domain format **none** and associated with format **icc-based**. If they are not, the command fails and an error message is generated in the CLI and SNMP indicating the problem.

Default

no eth-test-enable

bit-error-threshold

Syntax

bit-error-threshold *bit-errors*

Context

config>eth-ring>path>eth-cfm>mep>eth-test-enable

Description

This command specifies the lowest priority defect that is allowed to generate a fault alarm.

Default

1

Parameters

bit-errors

specifies the lowest priority defect

Values 0 to 11840

test-pattern

Syntax

test-pattern {**all-zeros** | **all-ones**} [**crc-enable**]

no test-pattern

Context

config>eth-ring>path>eth-cfm>mep>eth-test-enable

Description

This command configures the test pattern for Ethernet test frames.

The **no** form of the command removes the test pattern from the Ethernet test configuration.

Default

all-zeros

Parameters

all-zeros

configures the test pattern with all zeros

all-ones

configures the test pattern with all ones

crc-enable

generates a CRC checksum

low-priority-defect

Syntax

low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon}

Context

config>eth-ring>path>eth-cfm>mep

Description

This command specifies the lowest priority defect that will generate a fault alarm.

Default

remErrXcon

Parameters

allDef

all defects (DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM) generate a fault alarm

macRemErrXcon

DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM generate a fault alarm

remErrXcon

DefRemoteCCM, DefErrorCCM, and DefXconCCM generate a fault alarm

errXcon

DefErrorCCM and DefXconCCM generate a fault alarm

xcon

DefXconCCM generates a fault alarm

noXcon

no defects DefXcon or lower generate a fault alarm

one-way-delay-threshold

Syntax

one-way-delay-threshold *seconds*

Context

config>eth-ring>path>eth-cfm>mep

Description

This command configures a one-way delay threshold time limit.

Default

3

Parameters

seconds

the threshold value, in seconds

Values 0 to 600

rpl-end

Syntax

[no] **rpl-end**

Context

config>eth-ring>path

Description

This command configures the G.8032 path as a ring protection link (RPL) end. The ring must be declared as either an RPL owner or an RPL neighbor in order for this command to be valid. Only a path configured as A or B can be configured as an RPL end.

The **no** form of this command reverts to the default.

Default

no rpl-end

revert-time

Syntax

revert-time *time*

no revert-time**Context**

config>eth-ring

Description

This command configures the revert time for an Ethernet ring. It ranges from 60 s to 720 s.

The **no** form of this command means that the Ethernet ring is in non-revertive mode and the revert time is essentially 0. The revert timers are not set.

Default

300

Parameters

time

specifies the revert time, in seconds

Values 60 to 720

rpl-node**Syntax**

rpl-node {owner | nbr}

no rpl-node

Context

config>eth-ring

Description

This command configures the G.8032 ring protection link type as either owner or neighbor. The **no** form of the command means this node is not connected to an RPL link. When an RPL owner or neighbor is specified, either the A or B path must be configured with the [rpl-end](#) command. An owner is responsible for operation of the RPL link. The link can be left with no RPL type configured, but if this command is used the **nbr** parameter is mandatory.

On a sub-ring without a virtual channel, it is mandatory to configure a sub-ring non-virtual-link on all nodes on the sub-ring to propagate the RAPS messages around the sub-ring.

The **no** form of this command removes the RPL link.

Default

no rpl-node

sub-ring

Syntax

[no] sub-ring {virtual-link | non-virtual-link}

Context

config>eth-ring

Description

This command defines an Ethernet ring as a sub-ring as defined in G.8032. A sub-ring can have only one valid path connecting it to a major ring or to a VPLS instance. The **virtual-link** parameter indicates that a sub-ring is connected to another ring and that control messages can be sent over the attached ring to the other side of the sub-ring. The **non-virtual-link** channel parameter indicates that a sub-ring may be connected to a another ring or to a VPLS instance but that control messages from the sub-ring can not use the attached ring or VPLS instance. The non-virtual channel behavior is standard G.8032 capability.

The **no** form of this command deletes the sub-ring and its virtual channel associations.

Default

no sub-ring

Parameters

virtual-link

specifies that the interconnection is to a ring and a virtual link will be used

non-virtual-link

specifies that the interconnection is to a ring or a VPLS instance and a virtual link will not be used

interconnect

Syntax

interconnect ring-id *ring-index*

interconnect vpls

no interconnect

Context

config>eth-ring>sub-ring

Description

This command links the G.8032 sub-ring to a ring instance or to a VPLS instance. The ring instance must be a complete ring with two paths but may itself be a sub-ring or a major ring, as declared by its configuration on another node. When the interconnection is to another node, the sub-ring may have a virtual link or a non-virtual-link. When the sub-ring has been configured with a non-virtual link, the sub-ring

may be alternatively be connected to a VPLS service. This command is only valid on the interconnection node where a single sub-ring port connects to a major ring or terminates on a VPLS service.

The **no** form of this command removes the interconnect node.

Default

no interconnect

Parameters

ring-index

specifies the identifier for the ring instance of the connection ring for this sub-ring on this node

Values 0 to 128

vpls

specifies that the sub-ring is connected to the VPLS instance that contains the sub-ring SAP

propagate-topology-change

Syntax

[no] propagate-topology-change

Context

config>eth-ring>sub-ring>interconnect

Description

This command configures the G.8032 sub-ring to propagate topology changes from the sub-ring to the major ring as specified in the G.8032 interconnection flush logic. This command is only valid on a sub-ring and the interconnection node. A virtual link or non-virtual link must be specified for this command to be valid. The command is blocked on major rings (when both path A and B are specified on a ring).

The **no** form of this command disables topology propagation.

Default

no propagate-topology-change

3.17.2.2 Show commands



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

eth-ring

Syntax

```
eth-ring [status]
eth-ring [ring-index] hierarchy
eth-ring ring-index [path {a | b}]
```

Context

show

Description

This command displays Ethernet ring information.

Parameters

- status**
displays an Ethernet ring status summary
- ring-index**
specifies an Ethernet ring index
Values 1 to 32
- hierarchy**
displays Ethernet ring hierarchical relationships
- path**
displays information for a specific path

Output

The following output is an example of Ethernet ring information and [Table 14: Ethernet ring command field descriptions](#) describes the fields.

Output example

```
*A:7705:Dut-A# show eth-ring status
=====
Ethernet Ring (Status information)
=====
Ring  Admin  Oper    Path Information      MEP Information
ID    State   State   Path      Tag      State   Ctrl-MEP CC-Intvl Defects
-----
1      Up      Up      a - lag-32    1    Up      Yes      100ms  -----
          b - 1/1/1      1    Up      Yes      100ms  -----
2      Up      Up      a - lag-32    3.3  Up      Yes      100ms  -----
          b - 1/1/1      3.3  Up      Yes      100ms  -----
3      Up      Up      a - 1/1/3     5.5  Up      Yes      100ms  -----
          b - N/A        -    -      -      -      -
=====
Ethernet Tunnel MEP Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM
```

```

*A:7705:Dut-A# show eth-ring 1 hierarchy
=====
Ethernet Ring 1 (hierarchy)
=====
Ring Int  Admin Oper          Paths Summary          Path States
ID  ID  State State                      a      b
-----
1   -   Up   Up   a - lag-32          1      b - 1/1/1          1      U      U
3   1   Up   Up   a - 1/1/3          5.5    b - Not configured  B      -
=====
Ethernet Ring Summary Legend:  B - Blocked    U - Unblocked

*A:7705:Dut-A# show eth-ring 1 path a
=====
Ethernet Ring 1 Path Information
=====
Description      : Ethernet Ring : 1 Path : pathA
Port             : lag-32          Raps-Tag           : 1.0
Admin State      : Up              Oper State         : Up
Path Type        : normal          Fwd State          : unblocked
                                      Fwd State Change   : 03/19/2019 18:51:35

Last Switch Command: noCmd
APS Rx PDU        : Request State: 0x0
                  Sub-Code       : 0x0
                  Status          : 0xE0 ( RB DNF BPR )
                  Node ID        : 94:e9:8c:0c:db:d9
=====

*A:7705:Dut-A# show eth-ring 1 path b
=====
Ethernet Ring 1 Path Information
=====
Description      : Ethernet Ring : 1 Path : pathB
Port             : 1/1/1          Raps-Tag           : 1.0
Admin State      : Up              Oper State         : Up
Path Type        : normal          Fwd State          : unblocked
                                      Fwd State Change   : 03/19/2019 18:51:39

Last Switch Command: noCmd
APS Rx PDU        : Request State: 0x0
                  Sub-Code       : 0x0
                  Status          : 0xE0 ( RB DNF BPR )
                  Node ID        : 94:e9:8c:0c:db:d9
=====

*A:7705:Dut-A#

```

Table 14: Ethernet ring command field descriptions

Label	Description
Ring ID	Displays the Ethernet ring ID number
Admin State	Up – the Ethernet ring or path is administratively enabled Down – the Ethernet ring or path is administratively disabled
Oper State	Up – the Ethernet ring or path is operationally enabled Down – the Ethernet ring or path is operationally disabled
Path	Displays the configured Ethernet ring path

Label	Description
Tag	Displays the APS tag
Ctrl-MEP	Indicates whether the MEP is configured as a control MEP
CC-Intvl	Displays the configured CCM interval value
Defects	Displays any Ethernet tunnel MEP defects
Int ID	Displays the interface ID
Path States	Displays the state of Ethernet ring paths B – Blocked U – Unblocked
Description	Displays the Ethernet ring path description
Fwd State	Displays the forwarding state of the path
Fwd State Change	Displays the date and time of the last change to the path's forwarding state

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 15: Customer command field descriptions](#) describes the fields.

Output example

```
*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 15: Customer command field descriptions

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* **keep-alive-history**

sdp [*sdp-id*] [**detail**]

sdp far-end *ip-addr* **keep-alive-history**

sdp far-end *ip-addr* [**detail**]

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 16: Service SDP field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B>show>service# sdp
=====
Services: Service Destination Points
=====
SdpId  AdmMTU  OprMTU  Far End          Adm  Opr          Del    LSP    Sig
-----
 8      0       0       10.10.10.10     Down Down        MPLS           TLDP
11      0       0              Down Down        MPLS           TLDP
30      0       0              Down Down        MPLS           TLDP
900     0       0       10.10.10.10     Down Down        MPLS           TLDP
1000    0      1546     1.1.1.1         Up   Up          MPLS    L      TLDP
-----
Number of SDPs : 5
-----
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
        I = SR-ISIS, O = SR-OSPF, T = SR-TE, F = FPE
=====
*A:Sar18 Dut-B>show>service#
```

```
*A:Sar18 Dut-B>show>service# sdp 8
=====
Service Destination Point (Sdp Id : 8)
=====
SdpId  AdmMTU  OprMTU  Far End          Adm  Opr          Del    LSP    Sig
-----
 8      0       0       10.10.10.10     Down Down        MPLS           TLDP
=====
*A:Sar18 Dut-B>show>service#
```

```
*A:Sar18 Dut-B>show>service# sdp 1000 detail
=====
Service Destination Point (Sdp Id : 1000) Details
=====
-----
Sdp Id 1000 -1.1.1.1
-----
Description          : (Not Specified)
SDP Id               : 1000          SDP Source           : manual
Admin Path MTU       : 0            Oper Path MTU        : 1546
Delivery              : MPLS
Far End              : 1.1.1.1
Tunnel Far End       : 1.1.1.1      LSP Types            : LDP

Admin State          : Up            Oper State           : Up
Signaling            : TLDP          Metric               : 0
Last Status Change   : 05/23/2018 18:41:32 Adv. MTU Over.       : No
Last Mgmt Change     : 05/23/2018 18:41:05 VLAN VC Etype       : 0x8100
Flags                : None

Mixed LSP Mode Information :
Mixed LSP Mode          : Enabled    Active LSP Type       : LDP
```

```

Revert Time      : 0                      Revert Count Down : n/a

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

-----
BGP Tunnel Information :
-----
BGP LSP Id       : 0

-----
LDP Information :
-----
LDP LSP Id       : 65537

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

-----
Segment Routing
-----
ISIS              : disabled
OSPF               : disabled
TE-LSP            : disabled

-----
Group Encryption
-----
Inbound Keygroup Id : N/A
Outbound Keygroup Id : N/A
=====
*A: Sar18 Dut-B>show>service# *

*A: ALU-12>show>service# sdp 11 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 16: Service SDP field descriptions

Label	Description
SDP Id	Identifies the SDP

Label	Description
Description	Identifies the SDP by the text description stored in its configuration file
SDP Source	Specifies the SDP source type
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
Tunnel Far End	On the 7705 SAR, tunnel far end is the same as the SDP far end and is not configurable
LSP Types	Specifies the supported LSP types: R = RSVP, L = LDP, B = BGP, I = SR-ISIS, O = SR-OSPF, T = SR-TE, and n/a = Not Applicable
Adm (Admin State)	Specifies the desired state of the SDP
Opr (Oper State)	Specifies the operating state of the SDP
Del (Delivery)	Specifies the type of delivery used by the SDP: MPLS, GRE, or IP
Sig (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 tiebreaker 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
Flags	Specifies all the conditions that affect the operating status of this SDP
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified

Label	Description
Mixed LSP Mode Information:	
Mixed LSP Mode	Indicates whether mixed-LSP mode is configured on the SDP
Active LSP Type	Indicates the active LSP type on the mixed-LSP SDP: RSVP or LDP
Revert Time	The number of seconds the service manager must wait before it resets the SDP configuration to a higher priority LSP type when one becomes available
Keepalive Information:	
Admin State	Specifies the desired keepalive state
Oper State	Specifies the operating keepalive state
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
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
Keepalive History:	
Time of Probe Report	Indicates the date and time of the report
RTT (ms)	Indicates round-trip time (RTT), in milliseconds.

Label	Description
Size	Indicates the size of the packet, in bytes
Status	Indicates the status of the response
LDP Information:	
LDP LSP Id	Indicates the LDP LSP identifier
RSVP/Static LSPs:	
Associated LSP List	Lists the associated LSPs 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
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
Segment Routing	
ISIS	Indicates the state of IS-IS segment routing: enabled or disabled
OSPF	Indicates the state of OSPF segment routing: enabled or disabled
TE-LSP	Indicates the state of TE-LSP segment routing: enabled or disabled
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 17: Service SDP-using field descriptions](#) describes the fields.

Output example

```
*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
=====
```

Table 17: Service SDP-using field descriptions

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

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] [mirror] [apipe] [ipipe] [cpipe] [sdp *sdp-id*]
[customer *customer-id*]

Context

show>service

Description

This command displays the services matching specific 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

mirror

displays matching mirror 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 18: Service service-using field descriptions](#) describes the fields.

Output example – 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
```

```
*A:7705:Dut-A# show service service-using
```

```
=====
```

```
Services
```

```
=====
```

ServiceId	Type	Adm	Opr	CustomerId	Service Name
100	Mirror	Up	Up	1	XYZ Mirror 100
1000	Epipe	Up	Up	1	XYZ Epipe 1000
2000	VPLS	Up	Up	1	
2100	VPLS	Up	Up	1	rVpls2100
2200	Ipipe	Up	Up	1	

```
-----
```

```

3000      IES      Up      Up      1      XYZ Ies 3000
3500      VPRN     Up      Up      1      XYZ Vprn 3500
4050      IES      Up      Up      1      XYZ Ies 4050
-----

```

Matching Services : 8

Output example – 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#

```

Output example – 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 18: Service service-using field descriptions

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

Label	Description
Service Name	The service name

3.17.2.3 Monitor commands

service

Syntax

service

Context

monitor

Description

This command enables the context to configure criteria to monitor specific service SAP criteria.

id

Syntax

id service-id

Context

monitor>service

Description

This command displays statistics for a specific service, specified by the *service-id*, at the configured interval until the configured count is reached.

The first screen displays the current statistics related to the *service-id*. The subsequent statistical information listed for each interval is displayed as a delta to the previous screen output.

When the keyword **rate** is specified, the rate per second for each statistic is displayed instead of the delta.

Monitor commands are similar to **show** commands, but only statistical information is displayed. Monitor commands display the selected statistics according to the configured number of times at the interval specified.

Parameters

service-id

identifies the service in the service domain

Values 1 to 2147483690 or *service-name*

sap

Syntax

sap *sap-id* [*interval seconds*] [*repeat repeat*] [*absolute* | *rate*]

Context

monitor>service>id

Description

This command displays statistics for a SAP associated with this service.

This command displays statistics for a specific SAP, identified by the port ID and encapsulation value, at the configured interval until the configured count is reached.

The first screen displays the current statistics related to the SAP. The subsequent statistical information listed for each interval is displayed as a delta to the previous screen output.

When the keyword **rate** is specified, the rate per second for each statistic is displayed instead of the delta.

Monitor commands are similar to **show** commands, but only statistical information is displayed. Monitor commands display the selected statistics according to the configured number of times at the interval specified.

Parameters

sap-id

identifies the SAP for the service

The *sap-id* can be configured in one of the formats described in the following table. The range of values for the parameters follow the table.

Table 19: SAP ID configurations

Type	Syntax	Example
port-id	<i>slot/mda/port[.channel]</i>	1/1/5
bridge	<i>slot/mda/<bridge-id.branch-id></i>	1/5/16.10
null	<i>[port-id bundle-id lag-id aps-id mw-link-id]</i>	<i>port-id</i> : 1/1/3 <i>bundle-id</i> : bundle-ppp-1/1.1 <i>lag-id</i> : lag-1 <i>aps-id</i> : aps-1 <i>mw-link-id</i> : mw-link-1
dot1q	<i>[port-id lag-id aps-id mw-link-id]:qtag1</i>	<i>port-id</i> :qtag1: 1/1/3:100 <i>lag-id</i> : lag-1:10 <i>aps-id</i> : aps-1 <i>mw-link-id</i> : mw-link-1

Type	Syntax	Example
qinq	[<i>port-id</i> <i>lag-id</i>]:qtag1.qtag2	<i>port-id</i> :qtag1.qtag2: 1/1/3:100.30 <i>lag-id</i> : lag-1:10.10
atm	[<i>port-id</i> <i>aps-id</i>][:vpi/vci vpi vpi1.vpi2] ¹	<i>port-id</i> : 1/1/1 or 1/1/1.1 (for T1/E1 channelized ports) <i>aps-id</i> : aps-1 vpi/vci: 16/26 vpi: 16 vpi1.vpi2: 16.22
lag	<i>lag-id</i>	lag-2
frame	[<i>port-id</i> <i>aps-id</i>]: <i>dlci</i>	1/1/1 <i>aps-id</i> : aps-1 <i>dlci</i> : 16
frame relay	[<i>port-id</i>]: <i>dlci</i>	1/1/1 <i>dlci</i> : 16
cisco-hdlc	<i>slot/mda/port.channel</i>	1/1/1.3
cem	<i>slot/mda/port.channel</i>	1/1/1.3
ima-grp	<i>bundle-id</i> [:vpi/vci vpi vpi1.vpi2]	1/1/3.1
ipcp	<i>slot/mda/port.channel</i>	1/2/2.4
hdlc	<i>slot/mda/port.channel</i>	1/1/3.1
lag-id	<i>lag-id</i>	lag-1
mw-link-id	<i>mw-link-id</i>	mw-link-1
aps-id	<i>aps-group-id</i> [.channel]	aps-1
bundle-id	bundle-[<i>ima</i> <i>ppp</i>]- <i>slot/mda.bundle-num</i>	bundle-ima-1/1.1
tunnel-id	tunnel- <i><id></i> .[private public]: <i><tag></i>	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	<i>[port-id lag-id]:qtag1.qtag2</i>
atm	<i>[port-id aps-id][:vpi/vci vpi vpi1.vpi2]</i>
frame	<i>[port-id aps-id]:dlci</i>
cisco-hdlc	<i>slot/mda/port.channel</i>
cem	<i>slot/mda/port.channel</i>
ipcp	<i>slot/mda/port.channel</i>
ima-grp	<i>bundle-id[:vpi/vci vpi vpi1.vpi2]</i>
hdlc	<i>slot/mda/port.channel</i>
port-id	<i>slot/mda/port[.channel]</i>
bridge	<i>slot/mda/bridge-id.branch-id</i> <i>bridge-id 1 to 16</i> <i>branch-id 1 to 32</i>
bundle-id	<i>bundle-type-slot/mda.bundle-num</i> bundle keyword <i>type ima, ppp</i> <i>bundle-num 1 to 32</i>
aps-id	<i>aps-group-id[.channel]</i> aps keyword <i>group-id 1 to 24</i>
mw-link-id	<i>mw-link-id</i> <i>id 1 to 24</i>
lag-id	<i>lag-id</i> lag keyword <i>id 1 to 32</i>
qtag1	<i>*, 0 to 4094</i>
qtag2	<i>*, 0 to 4094</i>
vpi	<i>NNI 0 to 4095</i> <i>UNI 0 to 255</i>
vci	<i>1, 2, 5 to 65535</i>
dlci	<i>16 to 1022</i>
tunnel-id	<i>tunnel-id.[private public]:tag</i> tunnel keyword <i>id 1 to 16 (1 is the only valid value)</i>

tag 0 to 4094

- port-id*

specifies the physical port ID in the *slot/mda/port* 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.
- bridge-id*

specifies an existing bridge that has been configured on an Integrated Services card in the *slot/mda/<bridge-id.branch-id>* format

bridge-id value range: 1 to 16
- branch-id*

specifies an existing branch that has been configured on an Integrated Services card in the *slot/mda/<bridge-id.branch-id>* format

branch-id value range: 1 to 32
- bundle-id*

specifies the multilink (PPP or IMA) 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-xz5/1.1
*A:ALU-12>config>port# multilink-bundle
```

- qtag1, qtag2*

specifies the encapsulation value used to identify the SAP on the port or subport. 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. The following table describes the allowed values for the port and encapsulation types.

Table 20: 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. A 0 qtag1 value also accepts untagged packets on the dot1q port, and 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. A 0 qtag1 or qtag2 value also accepts untagged packets on the qinq port, and 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.

seconds

configures the interval for each display in seconds

Values 11 to 60

Default 11

repeat

configures how many times the command is repeated

Values 1 to 999

Default 10

absolute

displays the absolute rate-per-second value for each statistic

rate
displays the rate per second for each statistic instead of the delta

sap-aggregation-group

Syntax

sap-aggregation-group *group-id* [*interval seconds*] [*repeat repeat*] [*absolute | rate*]

Context

monitor>service>id

Description

This command displays the statistics for the specified SAP aggregation group that is associated with the service.

Parameters

group-id
specifies the identifier for the SAP aggregation group
Values 1 to 32 characters

seconds
configures the interval for each display in seconds
Values 11 to 60
Default 11

repeat
configures how many times the command is repeated
Values 1 to 999
Default 10

absolute
displays the absolute rate-per-second value for each statistic

rate
displays the rate per second for each statistic instead of the delta

Output

The following output is an example of statistics for a SAP aggregation group.

Output example

```
*A:SYS28# monitor service id 1570 sap-aggregation-group SAG repeat 2
=====
```

Monitor statistics for Service 1570 SAP Aggregation Group SAG

At time t = 0 sec (Base Statistics)

Sap Aggregation Group Statistics

Last Cleared Time : N/A

Dropped Egress Cells (unconfigured vpi/vci): 14

	Packets	Octets
Forwarding Engine Stats (Ingress)		
Dropped	: 0	n/a
Off. HiPrio	: 205557	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	: 205557	68605598

Forwarding Engine Stats (Egress)		
Dropped	: 0	n/a

Queueing Stats(Egress QoS Policy 1)		
Dro. InProf	: 0	n/a
Dro. OutProf	: n/a	n/a
For. InProf	: 202446	63083956
For. OutProf	: n/a	n/a

Sap Aggregation Group per Queue Stats

	Packets	Octets
Ingress Queue 1 (Priority)		
Off. HiPrio	: 205557	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	: 205557	68605598

Egress Queue 1		
For. InProf	: 202446	63083956
For. OutProf	: n/a	n/a
Dro. InProf	: 0	n/a
Dro. OutProf	: n/a	n/a

At time t = 11 sec (Mode: Delta)

Sap Aggregation Group Statistics

Last Cleared Time : N/A

Dropped Egress Cells (unconfigured vpi/vci): 14

	Packets	Octets
Forwarding Engine Stats (Ingress)		
Dropped	: 0	n/a

```

Off. HiPrio      : 233      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     : 233      77822

```

```

Forwarding Engine Stats (Egress)
Dropped          : 0        n/a

```

```

Queueing Stats(Egress QoS Policy 1)
Dro. InProf      : 0        n/a
Dro. OutProf     : n/a      n/a
For. InProf      : 232      72384
For. OutProf     : n/a      n/a

```

Sap Aggregation Group per Queue Stats

	Packets	Octets
Ingress Queue 1 (Priority)		
Off. HiPrio	: 233	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	: 233	77822
Egress Queue 1		
For. InProf	: 232	72384
For. OutProf	: n/a	n/a
Dro. InProf	: 0	n/a
Dro. OutProf	: n/a	n/a

At time t = 22 sec (Mode: Delta)

Sap Aggregation Group Statistics

Last Cleared Time : N/A

Dropped Egress Cells (unconfigured vpi/vci): 14

	Packets	Octets
Forwarding Engine Stats (Ingress)		
Dropped	: 0	n/a
Off. HiPrio	: 232	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	: 232	77488
Forwarding Engine Stats (Egress)		
Dropped	: 0	n/a
Queueing Stats(Egress QoS Policy 1)		
Dro. InProf	: 0	n/a

Dro. OutProf	: n/a	n/a
For. InProf	: 233	72696
For. OutProf	: n/a	n/a

Sap Aggregation Group per Queue Stats		

	Packets	Octets
Ingress Queue 1 (Priority)		
Off. HiPrio	: 232	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	: 232	77488
Egress Queue 1		
For. InProf	: 233	72696
For. OutProf	: n/a	n/a
Dro. InProf	: 0	n/a
Dro. OutProf	: n/a	n/a

4 VLL services

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](#)

4.1 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 about the adapter cards and chassis that support ATM VLL services.

4.1.1 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 35: ATM VLL for end-to-end ATM service](#)). 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.

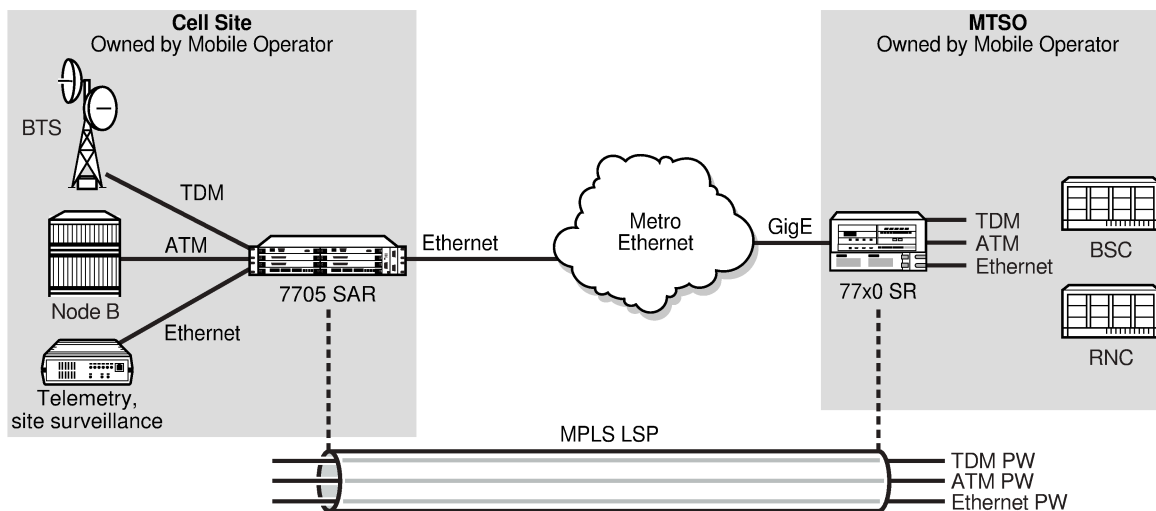
The ATM pseudowire (PW) is initiated using targeted LDP signaling as specified in RFC 4447, *Pseudowire Setup and Maintenance using LDP*; alternatively, it can be configured manually. The 7705 SAR supports MPLS, GRE, and IP as the tunneling technologies for transporting ATM PWs.

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 35: ATM VLL for end-to-end ATM service



19482

4.1.2 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 36: ATM virtual trunk over IP/ MPLS packet-switched network](#) shows two switches labeled ATM 2 and ATM 4 that appear 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 because the entire VT can be managed as a single entity from the NSP NFM-P (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-ietf-pwe3-control-protocol-xx, Pseudowire Setup and Maintenance using LDP*). In this application, there is no ATM signaling on the 7705 SAR gateway nodes because 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 functions 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 36: ATM virtual trunk over IP/ MPLS packet-switched network



22322

4.1.3 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 $N > 1$.

4.1.4 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](#)

4.1.4.1 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.

4.1.4.2 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 about ATM access egress queuing and scheduling, see the 7705 SAR Quality of Service Guide.

4.1.5 Control word

ATM VLLs support an optional control word. The control word contains protocol control information and can be enabled to guarantee ordered cell delivery for ATM (Apipe) pseudowire service. See [Pseudowire control word](#) for more information.

4.2 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](#)
- [Transparent SDH/SONET over packet \(TSoP\)](#)
- [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 about the adapter cards and chassis that support circuit emulation VLL services.

4.2.1 Cpipe service overview

Cpipe service is the Nokia 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 ($n \times 64$ 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).

4.2.2 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).

4.2.3 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 $n \times 64$ 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 \(SAToP\)](#)
- [SAToP serial](#)
- [SAToP teleprotection interface](#)
- [Structured mode \(CESoPSN\)](#)

4.2.3.1 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.

SAToP also supports asymmetric delay control (ADC). See [Asymmetric delay control](#) for information.

4.2.3.2 SAToP serial

A **satop-serial** virtual channel (vc) type can be configured on the 12-port Serial Data Interface card, version 3, to encapsulate serial traffic (subrate or super-rate) directly in the Cpipe without using high

capacity multiplexing (HCM). This capability allows the transport of serial control leads directly in the pseudowire instead of in HCM and allows the signaling to be transported with any line speed, not just subrate. It also extends support for TDM rates up to 16 Mb/s for the RS-530 interface. For subrate speeds, it can also use less bandwidth than the current HCM implementation, which occupies a full 64 kb/s timeslot.

The SAToP serial capability is supported on RS-530 and RS-232 interfaces.

4.2.3.2.1 SAToP serial payload size

The payload size for SAToP serial is configurable using the **payload-size** command as an integer number of octets and a multiple of 2 (instead of 32 for normal SAToP). This size affects the packet efficiency (that is, the payload-to-header ratio), packetization delay, and packets/s rate. The range is 2 to 1496 octets (instead of 1514 as for other Cpipes). See [Cpipe service configuration commands](#) for more information.

[Table 21: SAToP serial payload size minimums and defaults](#) shows the payload size minimum values and the defaults. Serial rates of 4800 b/s and lower only support a payload size of 2 bytes.

The minimum payload size is set so that the lowest bit rates can achieve the lowest delays possible and all rates do not exceed 4000 packets/s. The maximum payload size is set so that the packetization delay does not exceed 64 ms.

For rates of 1200 b/s and lower, the system oversamples to 2400 b/s. This results in the same packet size and packets/s as for 2400 b/s. Therefore, a 2-byte payload size is equivalent to 8 bits at 1200 b/s and 4 bits at 600 b/s of "real" serial data. Mismatched rate configurations between the two endpoints may not be identified when baud rates are 2400 b/s or lower.

Table 21: SAToP serial payload size minimums and defaults

Serial rate (kb/s)	Minimum payload size (bytes)	Minimum packetization delay	Minimum packets/s	Default payload size (bytes)	Default packetization delay	Default packets/s	Default jitter buffer (ms)
0.6	2 ¹	—	150	2 ¹	—	150	32
1.2	2 ¹	—	150	2 ¹	—	150	32
2.4	2	6.67	150	2	6.67	150	32
4.8	2	3.33	300	2	3.33	300	32
8	2	2.00	500	2	2.00	500	32
9.6	2	1.7	600	8	6.67	150	32
14.4	2	1.11	900	8	4.44	225	32
16	2	1.0	1000	8	4.00	250	32
19.2	2	0.83	1200	8	3.33	300	32
24	2	0.67	1500	8	2.67	375	32

Serial rate (kb/s)	Minimum payload size (bytes)	Minimum packetization delay	Minimum packets/s	Default payload size (bytes)	Default packetization delay	Default packets/s	Default jitter buffer (ms)
32	2	0.50	2000	8	2.00	500	32
38.4	2	0.42	2400	8	1.67	600	32
56	2	0.29	3500	8	1.14	875	32
57.6	2	0.28	3600	8	1.11	900	32
64	2	0.25	4000	8	1.00	1000	32
115.2	8	0.56	1500	64	4.44	225	16
128	8	0.50	2000	64	4.00	250	16
192	8	0.33	3000	64	2.67	375	16
256	8	0.25	4000	64	2.00	500	16
288	8	0.22	4500	128	3.56	281	8
336	8	0.19	5200	128	3.05	328	8
384	32	0.67	1500	128	2.67	375	8
512	32	0.50	2000	128	2.00	500	8
640	32	0.40	2500	128	1.60	625	8
704	32	0.36	2750	128	1.45	688	8
768	32	0.33	3000	128	1.33	750	8
896	32	0.29	3500	128	1.14	875	8
1024	32	0.25	4000	128	1.00	1000	5
1152	64	0.44	2250	256	1.78	563	5
1280	64	0.40	2500	256	1.60	625	5
1344	64	0.38	2625	256	1.52	656	5
1408	64	0.36	2750	256	1.45	688	5
1536	64	0.33	3000	256	1.33	750	5
1664	64	0.31	3250	256	1.23	813	5
1792	64	0.29	3500	256	1.14	875	5
1920	64	0.27	3750	256	1.07	938	5

Serial rate (kb/s)	Minimum payload size (bytes)	Minimum packetization delay	Minimum packets/s	Default payload size (bytes)	Default packetization delay	Default packets/s	Default jitter buffer (ms)
2048	64	0.25	4000	256	1.00	1000	5
3072	128	0.33	3000	256	0.67	1500	5
4096	128	0.25	4000	256	0.50	2000	5
5120	256	0.40	2500	1024	1.60	625	5
6144	256	0.33	3000	1024	1.33	750	5
7168	256	0.29	3500	1024	1.14	875	5
8192	256	0.25	4000	1024	1.00	1000	5
9216	512	0.44	2250	1024	0.89	1125	5
10240	512	0.40	2500	1024	0.80	1250	5
11264	512	0.36	2750	1024	0.73	1375	5
12288	512	0.33	3000	1024	0.67	1500	5
13312	512	0.31	3250	1024	0.62	1625	5
14336	512	0.29	3500	1024	0.57	1750	5
14360	512	0.27	3750	1024	0.53	1875	5
16384	512	0.25	4000	1024	0.50	2000	5

Note:

1. 600 and 1200 b/s are oversampled to 2400 b/s.

4.2.3.2.2 SAToP serial jitter buffer

For each circuit, the maximum receive jitter buffer is configurable using the **jitter-buffer** command. See [Cpipe service configuration commands](#) for more information. Playout from this buffer must start when the buffer is 50% full to give an operational packet delay variance (PDV) equal to half the maximum buffer size. The supported range is 12 to 250 ms in increments of 1 ms. The buffer size must be set to at least three times the packetization delay and not more than 32 times the packetization delay.

The default jitter buffer values are shown in [Table 21: SAToP serial payload size minimums and defaults](#).

Jitter buffer overrun and underrun counters are available via statistics and may be optionally alarmed while the circuit is up. On overruns, excess packets are discarded and counted. On underruns, all ones are sent for unstructured circuits.

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

4.2.3.3 SAToP teleprotection interface

The 8-port C37.94 Teleprotection card supports the SAToP teleprotection interface (TPIF) VC type, making it possible to transport the entire C37.94 signal, which is 2.048 Mb/s. Because this transport rate is the same rate as an E1 circuit, the encapsulation is as per RFC 4553. The SAToP TPIF VC type extends the current capability to transport $n \times 64$ kb/s channels within the C37.94 frame using CESoPSN.

The 8-port C37.94 Teleprotection card supports a maximum of four clear channel C37.94 ports for SAToP service. The card consists of four pairs of ports (ports 1/2, 3/4, 5/6, and 7/8). Only the odd-numbered port in each pair can be configured for SAToP, unframed C37.94. If a port is configured for SAToP, the even-numbered port cannot be used for framed or unframed service.

The SAToP TPIF VC has a configurable payload size and a configurable jitter buffer size. For information about configuring the payload size and jitter buffer, see [Packet payload size](#) and [Jitter buffer](#).

The SAToP TPIF VC also supports asymmetric delay control (ADC). See [Asymmetric delay control](#) for information.



Note: SAToP TPIF VCs are not supported on the 8-port Voice & Teleprotection card.

4.2.3.4 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 $n \times 64$ 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 channel associated signaling (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.

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).

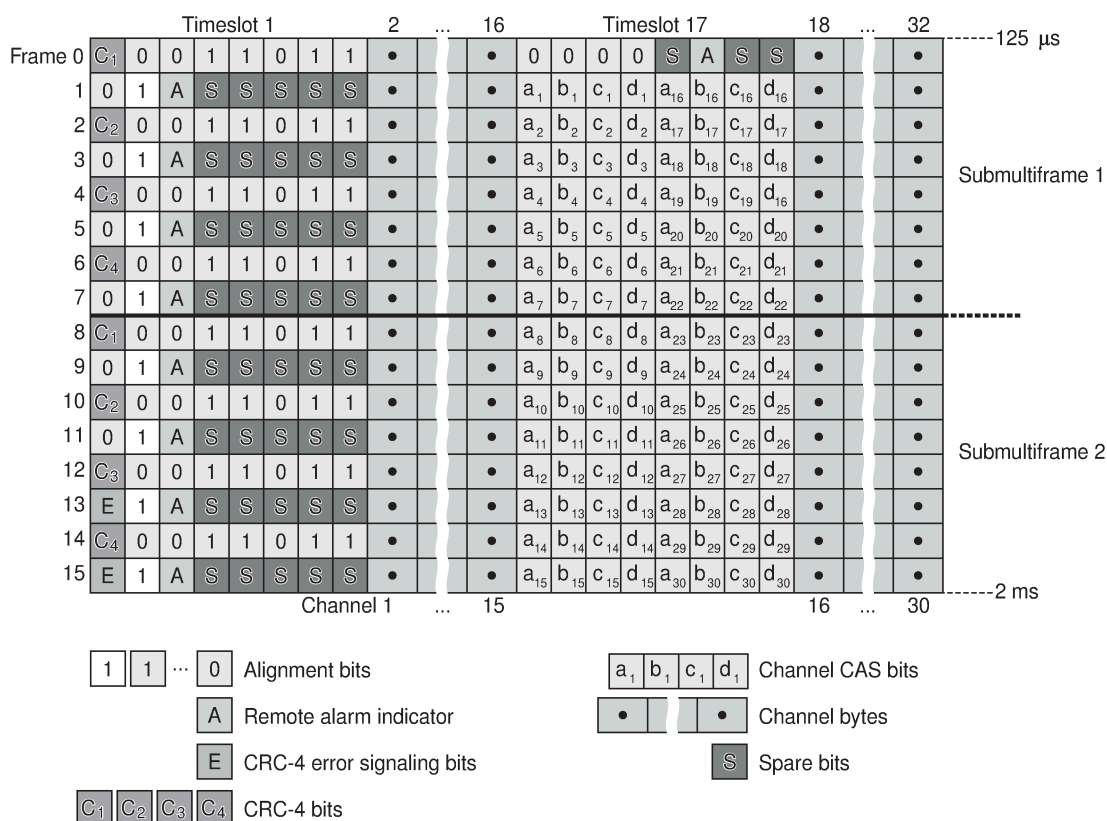
The mechanism for E1 CAS is described in ITU-T G.732. When configured for E1 CAS, timeslot 17 carries the signaling information for the timeslots used for voice trunking. Each channel requires four signaling bits, so grouping 16 E1 frames into a multiframe allows the signaling bits for all 30 channels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

As shown in the following figure, 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.

Figure 37: E1 framing for CAS support in an E1 multiframe



19966

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 extended superframe format (ESF) or superframe format (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.

In previous releases, when a Cpipe was configured for CESoPSN with CAS, the T1 ports at each end of the Cpipe had to be configured for the same framing format: either ESF or SF. If the framing formats did not match, a SapParamMismatch alarm was raised.

The 7705 SAR now supports mixed framing formats for the T1 ports on a Cpipe configured for CESoPSN with CAS; that is, one port can be configured for ESF framing and the other port for SF framing.

If the ingress port is an ESF-framed T1 port, when the packets arrive at the egress port, the ABCD bits from the ingress ESF SAP are sent out as AB bits in two consecutive superframes on the egress SF SAP. The CD bits from the ingress ESF SAP are mapped as AB bits in the second SF frame.

If the ingress port is an SF-framed T1 port, when the packets arrive at the egress port, the AB bits from every second SF frame from the ingress SF SAP are repeated twice as the ABCD bits of an egress ESF frame. The AB bits from the interlacing SF frames are dropped.

ESF and SF framing interoperability is supported on DS1 (T1) ports or channels on the following hardware:

- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- 2-port OC3/STM1 Channelized Adapter card
- 4-port DS3/E3 Adapter card
- 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card
- 4-port T1/E1 and RS-232 Combination module (supported on the 7705 SAR-H)
- 7705 SAR-A (variant with T1/E1 ports)
- 7705 SAR-M (variants with T1/E1 ports)
- 7705 SAR-X

The following table 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 22: 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	—	—

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	
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	—	—
22	4054	—	—	e6	1-8	—	—
23	4247	—	m	—	1-8	—	—
24	4440	1	—	—	1-7	8	D

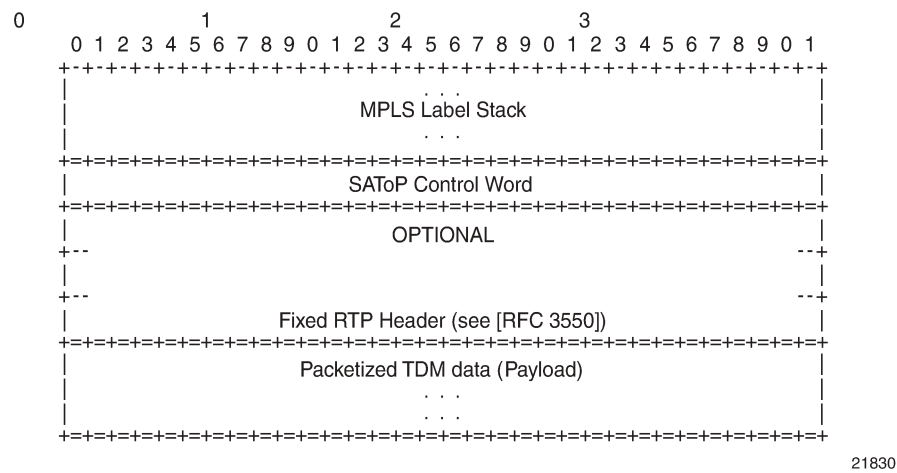
Notes:

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

4.2.4 TDM PW encapsulation

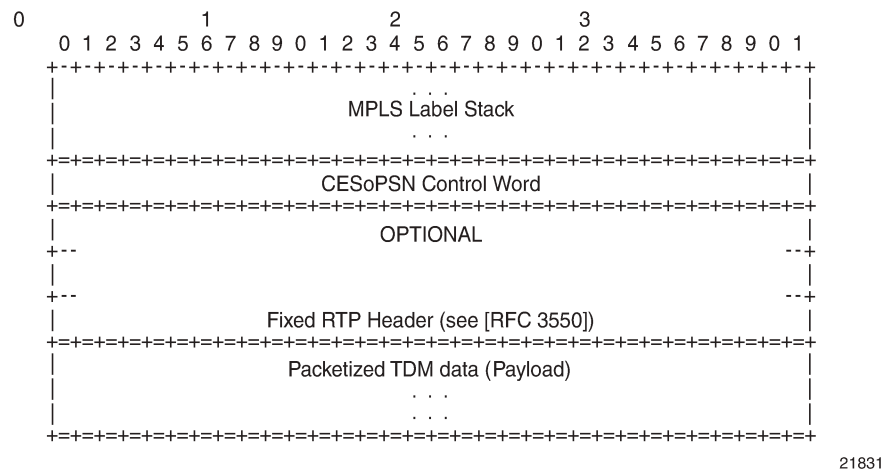
TDM circuits are MPLS-encapsulated as per RFC 4553 (SAToP) and RFC 5086 (CESoPSN). (see the following figures).

Figure 38: SAToP MPLS encapsulation



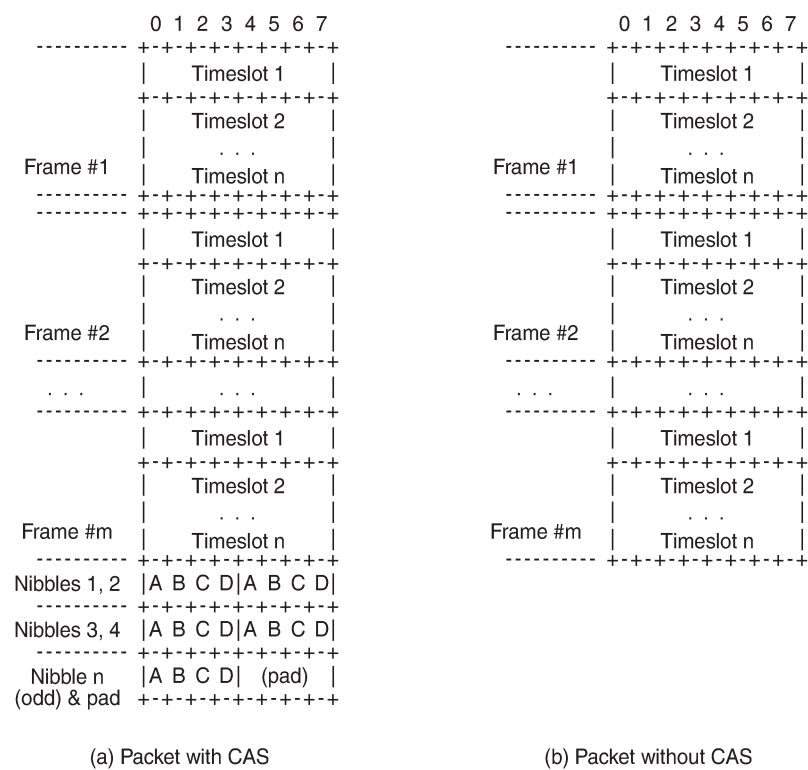
For GRE tunnels, the same encapsulations shown in the following figure are used, but GRE tunnel headers are used instead of MPLS tunnel headers.

Figure 39: CESoPSN MPLS encapsulation



The following figure shows the format of the CESoPSN TDM payload (with and without CAS) for packets carrying trunk-specific n x 64 kb/s service.

Figure 40: CESoPSN packet payload format for trunk-specific $n \times 64$ kb/s (with and without CAS transport)



21832

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 \times 24 frames/multiframe \times n multiframe, where $n = 1$ to 8. For an E1 port, the payload size is equal to the number of CAS and/or non-CAS timeslots \times 16 frames/multiframe \times n multiframe, 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.

4.2.5 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 $n \times 64$ kb/s timeslots in DS1 and E1 circuits as per RFC 5086.

The following table lists the adapter cards, modules, and chassis that support SAToP and CESoPSN.

Table 23: 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-chds1v2		
	On DS0/64k		✓
	On DS1/E1	✓	
32-port T1/E1 ASAP Adapter Card	a32-chds1v2		
	On DS0/64k		✓
	On DS1/E1	✓	
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, version 1 and version 2	a4-chds3, a4-chds3v2		
	On $n \times$ DS0		✓
	On DS1 ²	✓	
	On DS3/E3	✓	
12-port Serial Data Interface Card, version 2	a12-sdiv2		
	On V.35 and X.21 ports		✓
	On RS-232 ports		✓

Card/module/chassis		SAToP	CESoPSN
Name	CLI identifier (includes mode and channelization)		
12-port Serial Data Interface Card, version 3	a12-sdiv3		
	On V.35 and X.21 ports		✓
	On RS-232 ports	✓ (SAToP serial)	✓
	On RS-530 ports	✓ (SAToP serial)	✓
6-port E&M Adapter Card	a6-em		✓
8-port Voice & Teleprotection Card	a8-vt		✓
8-port C37.94 Teleprotection Card	a8-c3794	✓ (SAToP TPIF)	✓
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	✓	✓
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 more information about 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](#)
- [Cpipe network latency measurement](#)
- [RTP header](#)
- [Control word](#)

4.2.5.1 Unstructured

Unstructured CES is configured by choosing **satop-t1**, **satop-e1**, **satop-t3**, **satop-e3**, or **satop-tpif** as the **vc-type** when creating a Cpipe service. For DS1, E1, and TPIF unstructured circuit emulation, the framing parameter of the port must be set to **ds1-unframed**, **e1-unframed**, and **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 a multiple of 32. The payload packet size affects the packet efficiency and packetization delay. The following table 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, the facility data link (FDL) channel is transported over the Cpipe as part of the SAToP service. Regardless of the underlying framing, the framing parameter of the port must be set to unframed.

Table 24: 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

4.2.5.2 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 $n \times 64$ 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 25: Default and minimum payload sizes for CESoPSN without CAS](#) 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 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 = $2 \times I \times (\text{number of timeslots})$, where $I = 1$ to 20
- for DCR with a timestamp frequency of 19.44 MHz, the payload size = $8 \times I \times (\text{number of timeslots})$, 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.

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:

$$S = N \times F$$

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:

$$\begin{aligned} S &= 4 \text{ octets/frame} \times 16 \text{ frames} \\ &= 64 \text{ octets} \end{aligned}$$

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 μ s.

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

$$\begin{aligned} D &= 125 \mu\text{s/frame} \times 16 \text{ frames} \\ &= 2.000 \text{ ms} \end{aligned}$$

The following table 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 $2 \leq N \leq 4$, 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 25: Default and minimum payload sizes for CESoPSN without CAS

Number of timeslots (N)	Default values			Minimum values		
	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
23	8	184	1.000	2	46	0.250

	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)
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

4.2.5.3 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. See the 7705 SAR Interface Configuration Guide for information about 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 required 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-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.

The following table 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. See [Figure 40: CESoPSN packet payload format for trunk-specific \$n \times 64\$ kb/s \(with and without CAS transport\)](#) 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), the Cpipe will not become operational. This must be considered if you change the **service-mtu** from its default value.

Table 26: 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
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

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)
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
28	NA	NA	NA	16	448	2.00
29	NA	NA	NA	16	464	2.00
30	NA	NA	NA	16	480	2.00

4.2.5.4 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).

4.2.5.5 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.

4.2.5.5.1 Configuration and 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 $2 \leq N \leq 4$, the default is 16 ms
- for $5 \leq N \leq 15$, the default is 8 ms
- for $N \geq 16$, 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.

4.2.5.6 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.

Asymmetric delay control (ADC) is used to control the asymmetric latency contributed by the jitter buffer. When the **asym-delay-control** command is enabled, a special startup sequence is triggered when the TDM PW is initially started or is restarted after being brought down (caused by 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 8-port C37.94 Teleprotection card) or the configured **idle-payload-fill** value (for the other port types). See the 7705 SAR Interface Configuration Guide for information about 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:

- ADC 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 using RSVP-TE or SR-TE
 - do not use MPLS FRR or loop-free alternate paths (LFA, R-LFA, or TI-LFA) anywhere along the path 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 ADC, 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, between the two SAPs across the network. If different paths are used, latencies may be different, causing asymmetry. To prevent this situation, the 7705 SAR supports path redundancy for ADC. See [ADC for redundant paths](#).

Optionally, the ADC analysis can be set to repeat at configured time intervals (**min-repeat**) 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 the **threshold-repeat** value, octets are added or dropped as necessary.



Note: The **min-repeat** option is not configurable for the **asym-delay-control** command if active multipath (AMP) mode is enabled. For more information about AMP, see [Active multipath](#).

On-demand ADC allows users to initiate a one-time ADC analysis and correction on a live service using the **tools>perform>service>id>sap** command. Similar to the ADC repeat function, ADC uses the **threshold-repeat** value to determine if octets need to be added or dropped.

If ADC 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 ADC 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.

ADC can only be enabled for Cpipes configured as CESoPSN without CAS, SAToP TPIF, or SAToP (applies only to E1 circuits on the 16-port T1/E1 ASAP Adapter card or the 32-port T1/E1 ASAP Adapter card). If ADC is enabled, ACR, DCR, and RTP cannot be enabled on the port.

The following adapter cards, modules, and platforms support ADC:

- 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) on the 7705 SAR-8 Shelf V2 and 7705 SAR-18
- 8-port C37.94 Teleprotection card (C37.94 (TPIF) channels) on the 7705 SAR-8 Shelf V2 and 7705 SAR-18
- 12-port Serial Data Interface cards (RS-232, X.21, and V.35 on both 12-port Serial Data Interface card versions and RS-530 on 12-port Serial Data Interface card, version 3) on the 7705 SAR-8 Shelf V2 and 7705 SAR-18
- 7705 SAR-Hc (RS-232 channels)
- T1/E1 ports on the 16-port T1/E1 ASAP Adapter card and 32-port T1/E1 ASAP Adapter card on the 7705 SAR-8 Shelf V2 and 7705 SAR-18

4.2.5.6.1 ADC for redundant paths

When two paths are created between Cpipe endpoint routers, there is no guarantee that the latency of the two paths is exactly the same. Each path may be a different distance and have different numbers or types of switches or routers, and path failures may occur in a single direction. Automatic path switchover in these cases will result in asymmetry of traffic latency. This is problematic for networks that require high availability, such as power industry networks that use teleprotection. To overcome this problem, the 7705 SAR supports ADC over redundant network paths.

To enable ADC over redundant network paths in a Cpipe service, each router in the service must be configured with one SAP and two SDPs, where:

- one router is configured as the standby-signaling master and the other is configured as the standby-signaling slave
- the two SDPs on each router provide two different paths between the routers. In order to keep the service symmetric, both the master endpoint router and the slave endpoint router must use the same SDP and therefore the same path at any one time.
- each path is made up of two unidirectional LSPs with strict hop-by-hop routing over the two routers

If the active path becomes unavailable, as detected through LOS, BFD failure, LSP down, or spoke SDP down, the standby-signaling master and the standby-signaling slave routers both switch over to the available path.

After each path switchover, ADC automatically executes its analysis and resets the jitter buffer latency to the engineered value. This occurs because the switchover process may leave the path in a state that is susceptible to asymmetry.

In addition, TDM PWs enabled with ADC receive data only from the active path. Normally, incoming traffic is accepted from both active and inactive paths. However, because in-transit traffic may cause symmetry issues after a path switchover, only traffic on the active path is accepted.

4.2.5.6.2 Active multipath

Active multipath (AMP) mode allows TDM traffic to be transmitted simultaneously over up to four paths from the near end of a Cpipe to the far end. This allows the traffic to switch paths in a hitless fashion while experiencing minimal packet loss.

AMP is supported on the following:

- 8-port Voice & Teleprotection card
- 12-port Serial Data Interface card, version 2 and 3
- 8-port C37.94 Teleprotection card
- 16-port T1/E1 ASAP Adapter card
- 32-port T1/E1 ASAP Adapter card
- 6-port E&M Adapter card
- 4-port T1/E1 and RS-232 Combination module on the 7705 SAR-H
- RS-232 ports on the 7705 SAR-Hc
- E1 ports on the 7705 SAR-A

At startup or restart of the Cpipe service, at least one of the configured bidirectional paths must be up to support this functionality.

A timer with a value between 1 and 60 s is configured at each end of the Cpipe using the **active-multipath-timeout** command. If ADC is enabled, each end signals to the other end the availability of incoming paths (if ADC is not enabled, there is no signaling done between the two ends). At system startup, the system determines if locally available paths are present. After a local path becomes available, the system starts the timer. If all the remaining configured local paths become available before the timer expires, the timer stops and the available paths are immediately signaled to the far end. Otherwise, when the timer expires, the system uses only the available paths.

If the Cpipe service is ADC-enabled, the startup process for the Cpipe can only continue if there is at least one common bidirectional path that is up. If the Cpipe service is not ADC-enabled, the Cpipe can start up with any available paths.

If the Cpipe service is ADC-enabled, only one common path is selected for ADC analysis and jitter buffer adjustment. If more than one common path is available, the path with the lowest virtual circuit identifier (VCI) is selected for initial ADC analysis and jitter buffer adjustment.



Note: The path with the lowest VCI may not be the fastest path. This does not affect correct ADC startup.

During the ADC process, the selected common path for analysis must remain available at both ends of the Cpipe for the duration of the ADC sampling period; otherwise, a restart is required at both ends.

After the ADC process completes, there may be a shift in the jitter buffer fill level that corresponds to the other available paths with different latencies.

After successful Cpipe service startup, newly available paths and newly unavailable paths are automatically added to or removed from the collection of paths for the Cpipe. There is no requirement for a path to be available in both directions.


If ADC is enabled, symmetrical TDM service (delay symmetry) is also provided over the Cpipe. If an application requires hitless path redundancy without symmetrical TDM service, ADC does not need to be

enabled. The 7705 SAR supports AMP with ADC on the same adapter cards and platforms that support AMP; the exception is the 6-port E&M Adapter card, which does not support ADC.

The following figure shows an example of AMP hitless simultaneous transmission of TDM traffic over two paths with symmetrical TDM service.

Two different spoke SDP paths for the TDM SAP are configured at the host node (Node L). The **active-multipath** command is enabled at the endpoint to allow packets that encapsulate the ingress TDM traffic to be duplicated and sent over both paths (Path 1 and Path 2). For maximum resilience, it is recommended that the network links of the two paths be on different adapter cards, but this is not a requirement.

Figure 41: AMP hitless simultaneous transmission of TDM traffic over two paths with symmetrical TDM service



Network latency may be different for each path between the 7705 SAR nodes at each end. Path 1 can be different from Path 2 but both directions of each path must be symmetrical. To achieve symmetrical TDM service, each path is based on hop-by-hop strict LSPs; fast reroute (FRR) is not supported.

The configuration for the TDM service is the same at both the egress and ingress 7705 SAR nodes. The duplicated packets arrive over two different paths and are sent to the associated TDM adapter card. The packets are first processed through the new combiner for each TDM service. The combiner then feeds only a single copy of each packet to the jitter buffer while discarding unneeded packets.

ADC is mandatory as part of the symmetrical TDM service configuration. ADC analysis and adjustment are based on the traffic egressing the combiner. Jitter buffer overruns and underruns may still occur and result in service restarts, which is normal operation.



Note: Because the combiner is part of the jitter buffer implementation, the size of the jitter buffer selected by the operator must take this into account. See [Cpipe network latency measurement](#) for information about configuring the jitter buffer size.

4.2.5.6.2.1 Combiner operation

On startup, the combiner waits for traffic from the configured paths to be sent.

The all-paths-active state on the Cpipe is achieved when traffic is received from the configured paths and the combiner is able to compensate for the differential delay.

The spoke SDP path activity state for active-multipath (all-paths-active, not-all-paths-active, initialization, or down) can be viewed using CLI **show** commands or SNMP commands.

4.2.5.6.2.2 Access ingress fabric shaping

If the **destination-mode** fabric profile is configured at access ingress, the multipoint destination shaper of the **destination-mode** fabric profile is used for shaping the traffic. For more information about fabric shaping, see the 7705 SAR Quality of Service Guide.

4.2.5.7 Cpipe network latency measurement

Network latency measurement can be enabled at the service level on a Cpipe to record and display data on minimum, current, and maximum latency values. The feature is configured using the **config>service>cpipe>network-latency-measurement** command. The feature is enabled independently of AMP and ADC. If AMP is enabled, the end-to-end latency is measured on all configured paths.

If AMP is enabled, the operator must configure a jitter buffer size large enough to deal with worse-case jitter scenarios and the differential latency between the paths. The configured jitter buffer size is recommended to be equal to 2 x (maximum jitter of any path + maximum differential delay between paths). The network latency measurement function is useful in this case because having real-time data on the latency of each path helps to configure the jitter buffer size.

If network latency measurement is enabled, regular Cpipe packets are enhanced to include a proprietary 8-byte timestamp in every packet that is sent over the service. These timestamps are based on the router's internal time clock that is driven by PTP or GNSS. At the far-end router, the Cpipe packets are timestamped on arrival based on the router's time clock that is also driven by PTP or GNSS. The end-to-end latency calculation can then be made, over 1024 packet windows, for the Cpipe packets sent over the service path.

The network latency measurement configuration (enabled or disabled) must be the same on all segments of the service, including the end nodes; otherwise, the unconfigured path does not come up and "Network Latency Mismatch" is displayed for each spoke SDP. Only end-to-end latency measurements are displayed, not segment-by-segment. The service must be shut down before the service configuration can be changed.

This feature is supported on any Cpipe, that is, a single-path Cpipe, a Cpipe with PW redundancy, or a Cpipe with AMP. In each case, the measured latencies for every configured path are displayed, so there is a single set of latencies for a single-path Cpipe and up to four sets of latencies for a Cpipe with PW redundancy and a Cpipe with AMP. After a **no shutdown** command, all latencies are initialized to zero and updated every 1024 packets. Any latencies that remain at zero are displayed as "N/A".

A display of "N/A" could be caused by no traffic being received over the path or no timestamp being enabled at the near end, far end, or both ends. Otherwise, the most recent update is displayed.

The maximum latency that can be measured is 34.3 seconds. For any latency above 34.3 seconds, the current latency is displayed as "Too High". If there is a clocking issue that results in a situation where the far-end timestamp is earlier than the near-end timestamp, the latency is negative. In this case, the current latency is displayed as "Too Low". In either case, the minimum and maximum latencies are not updated, so the most recent minimum and maximum values are displayed.

VC switching supports latency measurements but does not support AMP. Inter-chassis backup (ICB) does not support latency measurements but the command is not blocked from the user in the CLI.

Latency measurements are supported on network interfaces on adapter cards and platforms that support a real-time clock. This includes all Ethernet adapter cards and the 7705 SAR-A, 7705 SAR-M, and 7705 SAR-X. On the 7705 SAR-A, 7705 SAR-M, and 7705 SAR-X, the network port can also be based on PPP/MLPPP.

The measured latencies for the Cpipe service can be displayed using the **show service id network-latency-measurement** command or under the tools menu using the **tools>dump>service>id>network-latency-measurement** command.

4.2.5.8 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.

4.2.5.9 Control word

The control word is mandatory for SAToP and CESoPSN. The bit structure is shown in the following figure and the table describes the bit fields. See [Pseudowire control word](#) for more information.

Figure 42: Control word bit structure




Table 27: Control word bit descriptions

Bit	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.
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: <ul style="list-style-type: none">When L bit = 0, and<ul style="list-style-type: none">M = 00 – Normal conditionsM = 01 – Reserved for future use

Bit	Description
	<p>M = 10 – RDI condition for the attachment circuit (AC)</p> <p>M = 11 – Reserved for CESoPSN</p> <ul style="list-style-type: none"> When L bit = 1, and <p>M = 00 – TDM data is invalid</p> <p>M = 01 – Reserved for future use</p> <p>M = 10 – Reserved for future use</p> <p>M = 11 – Reserved for future use</p>
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.

4.2.6 Transparent SDH/SONET over packet (TSoP)

Transparent SDH/SONET over packet (TSoP) is a method for transporting clear channel OC3/STM1 or clear channel OC12/STM4 traffic over a packet network using OC3/STM1 TSoP SFPs and OC12/STM4 TSoP SFPs. With TSoP, the entire signal is encapsulated in a pseudowire and transported over the network to a single destination, which simplifies operation. TSoP is modeled after the SAToP method for pseudowire transport of DS1, E1, DS3, or E3 circuits (RFC 4553).

TSoP SFPs are inserted into Ethernet SFP ports, and the 7705 SAR treats them as standard Ethernet SFPs. To set up the TSoP service, an Epipe must be created across the network connecting two OC3/STM1 TSoP SFPs or two OC12/STM4 TSoP SFPs. The TSoP SFPs implement DCR for service clock delivery. Both nodes must be synchronized against a common clock for DCR.

TSoP SFPs are supported on the 7705 SAR-8 Shelf V2 and 7705 SAR-18 on the following adapter cards:

- 8-port Gigabit Ethernet Adapter card
- 6-port Ethernet 10Gbps Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card (7705 SAR-18 only)

Each adapter card supports two OC3/STM1 or OC12/STM4 TSoP SFPs. A maximum of 16 TSoP SFPs are supported on the 7705 SAR-8 Shelf V2 or 7705 SAR-18.



Note: For a 7705 SAR-8 Shelf V2 with a maximum ambient temperature of 131°F (55°C), a maximum of eight TSoP SFPs are supported per adapter card.

4.2.7 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. See the 7705 SAR 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.

4.3 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](#)
- [Epipe with ATM SAPs](#)
- [MEF 8](#)
- [Ethernet OAM](#)
- [Control word](#)
- [MTU](#)
- [Raw and tagged modes](#)
- [IP filters \(Epipe\)](#)
- [MPLS entropy label](#)
- [Security zones and Epipes](#)

Epipe configuration information is found under the following topics:

- [Common configuration tasks](#)
- [Configuring VLL components](#)
- [Service management tasks](#)

See [Service support](#) for information about the adapter cards and chassis that support Ethernet VLL services.

4.3.1 Epipe service overview

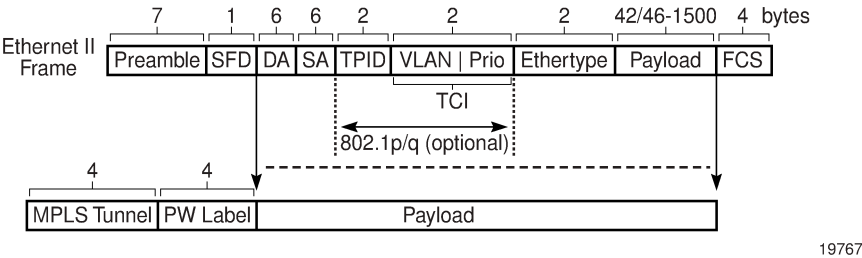
An Ethernet pseudowire (PW) is used to carry Ethernet/802.3 protocol data units (PDUs) over an MPLS or IP network, allowing service providers to offer emulated Ethernet services over existing MPLS or IP networks. For the 7705 SAR, Ethernet emulation is a point-to-point service.

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 Nokia implementation of an Ethernet VLL based on the IETF RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*.

The following figure shows a typical Ethernet VLL frame together with its MPLS tunnel encapsulation.


Figure 43: 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 44: Epipe service](#) 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 functions as 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 44: Epipe service



4.3.2 Ethernet access egress queuing and scheduling

Ethernet access egress queuing and scheduling is very similar to Ethernet access ingress behavior. When 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 about Ethernet access egress queuing and scheduling, see the 7705 SAR Quality of Service Guide.

4.3.3 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 are inserted in the egress frame if the SAP has a VLAN ID of 0 or VLAN IDs of 0.*. SAP 0.0 is not a valid combination.

In addition, the 7705 SAR-M can use a SAP-to-SAP Ethernet PW to provide an Ethernet-to-ATM interworking service. This is done by having one SAP on an Ethernet port and the other SAP on an ATM port or IMA bundle. Encapsulation options are specified in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. The Ethernet-to-ATM interworking service can be used to support:

- interworking of legacy bridged ATM traffic to Ethernet
- transport of Ethernet traffic over an existing ATM network


See [Configuring ATM encapsulation under Epipe service \(7705 SAR-M only\)](#) for more information.

4.3.4 Epipe with ATM SAPs

The 7705 SAR supports Epipe with ATM SAPs over an Ethernet SDP; this feature is available on the 7705 SAR-8 Shelf V2 or 7705 SAR-18. IP interworking is between an OC3 clear channel ATM over a 10-Gigabit or Gigabit Ethernet connection through an MPLS network. The SAP connection is from an ATM VC configured on a 4-port OC3/STM1 Clear Channel Adapter card. The Ethernet SDP connection is from a 6-port Ethernet 10Gbps Adapter card. The ATM SAP format can only be UNI. BPDU with LLC/SNAP is used as specified in RFC 2684.

The following figure shows an example of an Epipe network configuration with an ATM SAP on a 7705 SAR-8 Shelf V2. For a CLI configuration example, see [Configuring an Epipe with an ATM SAP](#).

Figure 45: Epipe network configuration with ATM SAP

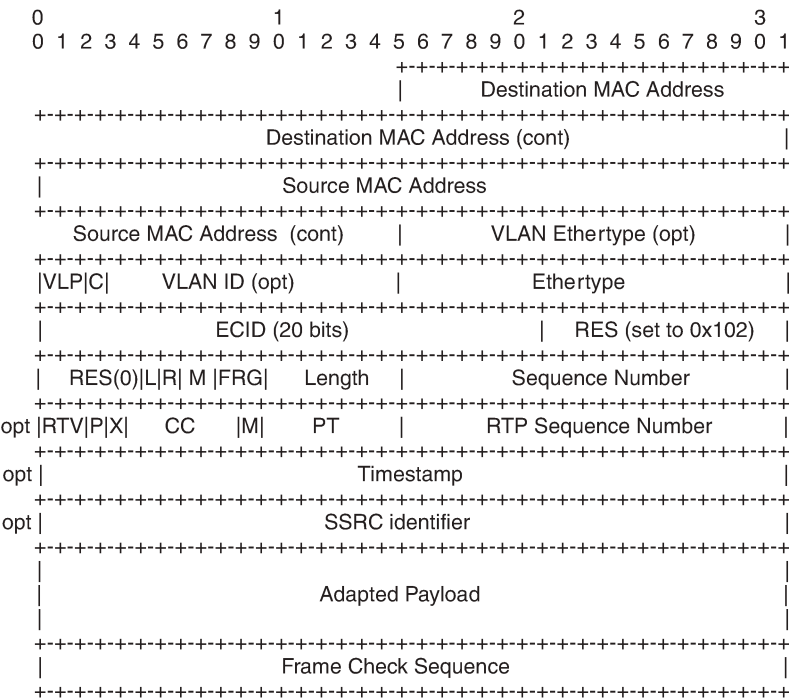


4.3.5 MEF 8

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.

The following figure shows an Ethernet-encapsulated TDM circuit. See [TDM PW encapsulation](#) for complete information about TDM PW encapsulation.

Figure 46: CESoETH encapsulation



23173

The 7705 SAR supports the following MEF 8 configuration scenarios:

- TDM SAP to Ethernet SAP (see [Figure 47: TDM SAP to Ethernet SAP](#))

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 spoke SDP (see [Figure 48: TDM SAP to spoke SDP](#))

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 47: TDM SAP to Ethernet SAP

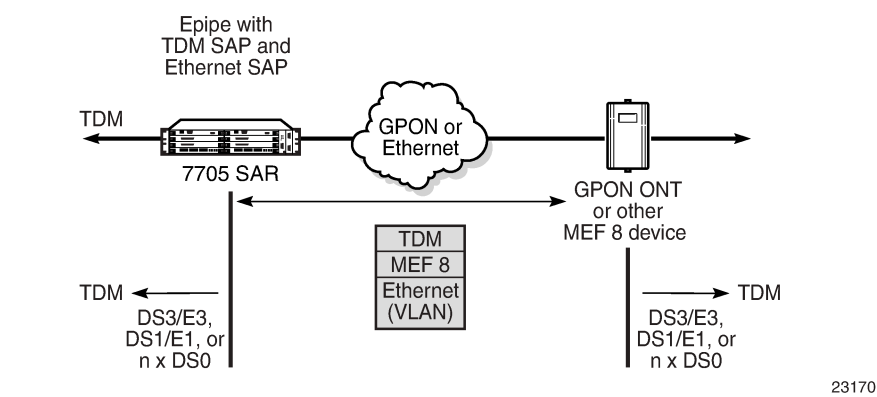
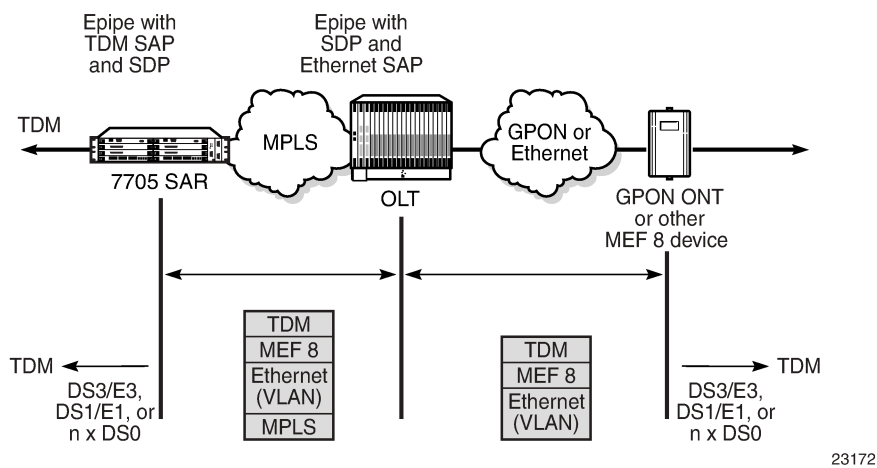


Figure 48: TDM SAP to spoke SDP



The following table shows the platforms, adapter cards, and modules that support MEF 8.

Table 28: 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	✓	✓ ²	✓	✓ ²

	TDM SAP to Ethernet SAP		TDM SAP to Ethernet via Epipe spoke SDP	
	Structured ¹	Unstructured	Structured ¹	Unstructured
All T1/E1 ports				
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
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
6-port E&M Adapter card	✓		✓	
6-port FXS Adapter card	✓		✓	
8-port FXO Adapter card	✓		✓	
12-port Serial Data Interface card, version 3	✓		✓	

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

4.3.5.1 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 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-A (variant with T1/E1 ports)
- 7705 SAR-M (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 about ACR, see the 7705 SAR Basic System Configuration Guide, "Adaptive clock recovery".

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

- 7705 SAR-A (variant with T1/E1 ports)
- 7705 SAR-M (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
- 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 about DCR, see the 7705 SAR Basic System Configuration Guide, "Differential clock recovery".

4.3.6 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 see the "Ethernet OAM" section in the 7705 SAR Interface Configuration Guide, and the "Ethernet OAM capabilities" section in the 7705 SAR 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, see the "Ethernet OAM" section in the 7705 SAR Interface Configuration Guide.

4.3.7 Control word

Ethernet (Epipe) services support an optional control word, with the exception of MEF 8 (CESoETH) services for which the control word is mandatory (for information about the MEF 8 service control word, see the [Control word](#) section in [Circuit emulation parameters and options](#)).

If the Epipe service control word is enabled, it is set to all zeros and ignored on egress.

See [Pseudowire control word](#) for more information.

4.3.8 MTU

Network-facing Ethernet ports must support a larger MTU than access-facing Ethernet ports 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 ports under various configurations, where the worst case is the largest MTU size required 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: Because 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 about MTU for VLL service.

4.3.9 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.

4.3.9.1 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 tags are 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.

4.3.9.2 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.

4.3.9.3 VLAN translation

VLAN ID translation is supported, as appropriate. [Table 31: VLAN tagging examples \(Epipe\)](#) (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.

4.3.9.4 Tagging rules for Epipe

[Table 29: Ingress SAP tagging rules](#) and [Table 30: Egress SAP tagging rules](#) 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 29: 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, the original tag is preserved. The VC tag is set on the spoke SDP, using the **vlan-vc-tag** command.

Table 30: 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, no VLAN tag is pushed.
3. If the SAP tag is 0, only the pop operation is performed.

The following table 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 31: VLAN tagging examples (Epipe)

Configuration settings				
Rx (Ing)/ Tx (Egr)	SAP/ SDP	Encap.	SAP VLAN tag	
			Outer	Inner
Rx	SAP	Null	N/A	N/A
Tx	SAP	Null	N/A	N/A
Result: Untagged, single-, and double-tagged frames are accepted. On egress, all frames are transmitted untouched.				

Configuration settings								
Rx (Ing)/ Tx (Egr)	SAP/ SDP	Encap.	SAP VLAN tag					
			Outer	Inner				
<div><div><div>Received Frames (Ingress)</div><div>Untagged</div><div><div>CRC</div><div>Payload</div><div>SA</div><div>DA</div></div></div><div>Single Tagged</div><div><div>CRC</div><div>Payload</div><div>VlanId=X</div><div>SA</div><div>DA</div></div></div> <div>Double Tagged</div> <div><div>CRC</div><div>Payload</div><div>VlanId=Y</div><div>VlanId=X</div><div>SA</div><div>DA</div></div>					<div><div>Transmitted Frames (Egress)</div><div>⇒ <div><div>CRC</div><div>Payload</div><div>SA</div><div>DA</div></div></div><div>⇒ <div><div>CRC</div><div>Payload</div><div>VlanId=X</div><div>SA</div><div>DA</div></div></div><div>⇒ <div><div>CRC</div><div>Payload</div><div>VlanId=Y</div><div>VlanId=X</div><div>SA</div><div>DA</div></div></div></div>			
Rx	SAP	Dot1q	252	N/A				
Tx	SAP	Null	N/A	N/A				
Result: Single- and double-tagged frames are accepted if the outermost VLAN ID matches. On egress, the outermost VLAN tag is popped.								
<div><div><div>Received Frames (Ingress)</div><div>Untagged</div><div><div>CRC</div><div>Payload</div><div>SA</div><div>DA</div></div></div><div>Single Tagged</div><div><div>CRC</div><div>Payload</div><div>VlanId=252</div><div>SA</div><div>DA</div></div></div> <div>Double Tagged</div> <div><div>CRC</div><div>Payload</div><div>VlanId=Y</div><div>VlanId=252</div><div>SA</div><div>DA</div></div>					<div><div>Transmitted Frames (Egress)</div><div>⇒ X (Frame is dropped on ingress)</div><div>⇒ <div><div>CRC</div><div>Payload</div><div>SA</div><div>DA</div></div></div><div>⇒ <div><div>CRC</div><div>Payload</div><div>VlanId=Y</div><div>SA</div><div>DA</div></div></div></div>			
Rx	SAP	QinQ	525	353				
Tx	SAP	Dot1q	789	N/A				
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.								

Transmitted Frames (Egress)

Untagged

CRC

Payload

SA

DA

Single Tagged

CRC

Payload

VlanId=X

SA

DA

Double Tagged

CRC

Payload

VlanId=Y

VlanId=X

SA

DA

24064

Configuration settings				
Rx (Ing)/ Tx (Egr)	SAP/ SDP	Encap.	SAP VLAN tag	
			Outer	Inner
<div><div><div>Received Frames (Ingress)</div><div>Untagged</div><div><div>CRC</div><div>Payload</div><div>SA</div><div>DA</div></div><div>⇒</div><div>X (Frame is dropped on ingress)</div></div><div><div>Single Tagged</div><div><div>CRC</div><div>Payload</div><div>VlanId=X</div><div>SA</div><div>DA</div></div><div>⇒</div><div>X (Frame is dropped on ingress)</div></div><div><div>Double Tagged</div><div><div>CRC</div><div>Payload</div><div>VlanId=353</div><div>VlanId=525</div><div>SA</div><div>DA</div></div><div>⇒</div><div><div>CRC</div><div>Payload</div><div>VlanId=789</div><div>SA</div><div>DA</div></div></div><div>24066</div></div>				
Rx	SAP	QinQ	525	353
Tx	SDP	VLAN	456	N/A
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.				
<div><div><div>Received Frames (Ingress)</div><div>Untagged</div><div><div>CRC</div><div>Payload</div><div>SA</div><div>DA</div></div><div>⇒</div><div>X (Frame is dropped on ingress)</div></div><div><div>Single Tagged</div><div><div>CRC</div><div>Payload</div><div>VlanId=X</div><div>SA</div><div>DA</div></div><div>⇒</div><div>X (Frame is dropped on ingress)</div></div><div><div>Double Tagged</div><div><div>CRC</div><div>Payload</div><div>VlanId=353</div><div>VlanId=525</div><div>SA</div><div>DA</div></div><div>⇒</div><div><div>CRC</div><div>Payload</div><div>VlanId=456</div><div>SA</div><div>DA</div></div></div><div>24067</div></div>				
Rx	SDP	Ether	N/A	N/A
Tx	SAP	QinQ	123	654
Result: Untagged, single-, and double-tagged frames are accepted. On egress, a double-push (outer and inner) is performed.				

Configuration settings				
Rx (Ing)/ Tx (Egr)	SAP/ SDP	Encap.	SAP VLAN tag	
			Outer	Inner
<div><div>Received Frames (Ingress)</div><div>Transmitted Frames (Egress)</div><div>Untagged</div><div>CRC Payload SA DA</div><div>⇒ CRC Payload VlanId=654 VlanId=123 SA DA</div><div>Single Tagged</div><div>CRC Payload VlanId=X SA DA</div><div>⇒ CRC Payload VlanId=X VlanId=654 VlanId=123 SA DA</div><div>Double Tagged</div><div>CRC Payload VlanId=Y VlanId=X SA DA</div><div>⇒ CRC Payload VlanId=Y VlanId=X VlanId=654 VlanId=123 SA DA</div><div>24068</div></div>				
Rx Tx	SAP SDP	Dot1q VLAN	Default (*) 741	N/A N/A
Result: Untagged, single-, and double-tagged frames are accepted. On egress, a push is performed.				
<div><div>Received Frames (Ingress)</div><div>Transmitted Frames (Egress)</div><div>Untagged</div><div>CRC Payload SA DA</div><div>⇒ CRC Payload VlanId=741 SA DA</div><div>Single Tagged</div><div>CRC Payload VlanId=X SA DA</div><div>⇒ CRC Payload VlanId=X VlanId=741 SA DA</div><div>Double Tagged</div><div>CRC Payload VlanId=Y VlanId=X SA DA</div><div>⇒ CRC Payload VlanId=Y VlanId=X VlanId=741 SA DA</div><div>24069</div></div>				
Rx Tx	SAP SDP	QinQ VLAN	852 963	0 N/A
Result: Single- and double-tagged frames are accepted if the outermost VLAN ID matches. On egress, the outermost VLAN tag is swapped.				

Configuration settings				
Rx (Ing)/ Tx (Egr)	SAP/ SDP	Encap.	SAP VLAN tag	
			Outer	Inner
<div><div><div>Received Frames (Ingress)</div><div>Untagged</div><div><div>CRC</div><div>Payload</div><div>SA</div><div>DA</div></div></div><div>⇒ X (Frame is dropped on ingress)</div><div>Single Tagged</div><div><div>CRC</div><div>Payload</div><div>VlanId=852</div><div>SA</div><div>DA</div></div><div>⇒ <div><div>CRC</div><div>Payload</div><div>VlanId=963</div><div>SA</div><div>DA</div></div></div><div>Double Tagged</div><div><div>CRC</div><div>Payload</div><div>VlanId=0</div><div>VlanId=852</div><div>SA</div><div>DA</div></div><div>⇒ <div><div>CRC</div><div>Payload</div><div>VlanId=0</div><div>VlanId=963</div><div>SA</div><div>DA</div></div></div></div> <div>24070</div>				
Rx	SAP	QinQ	Default (*)	Default (*)
Tx	SAP	QinQ	Default (*)	Default (*)
Result: Any double-tagged frames are accepted. On egress, all frames are transmitted untouched.				
<div><div><div>Received Frames (Ingress)</div><div>Untagged</div><div><div>CRC</div><div>Payload</div><div>SA</div><div>DA</div></div></div><div>⇒ X (Frame is dropped on ingress)</div><div>Single Tagged</div><div><div>CRC</div><div>Payload</div><div>VlanId=X</div><div>SA</div><div>DA</div></div><div>⇒ X (Frame is dropped on ingress)</div><div>Double Tagged</div><div><div>CRC</div><div>Payload</div><div>VlanId=Y</div><div>VlanId=X</div><div>SA</div><div>DA</div></div><div>⇒ <div><div>CRC</div><div>Payload</div><div>VlanId=Y</div><div>VlanId=X</div><div>SA</div><div>DA</div></div></div></div> <div>24071</div>				

4.3.10 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 email 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. See the 7705 SAR Router Configuration Guide, "Filter policies", for information about configuring filter policies and assigning them to network interfaces.

4.3.11 MPLS entropy label

For Cpipes, Epipes, and Ipipes, the router supports MPLS entropy labels as per RFC 6790. The entropy label provides greater granularity for load balancing on an LSR where load balancing is typically based on the MPLS label stack.

For more information, see the "MPLS entropy labels" section in the 7705 SAR MPLS Guide and the "LAG and ECMP hashing" section in the 7705 SAR Interface Configuration Guide.

4.3.12 Security zones and Epipes

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, see the 7705 SAR 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 2 endpoints or 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. Layer 3 zones support both NAT and firewall security policies. Layer 2 zones support only firewalls. 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>epipe>zone**
- **config>service>vpls>zone**
- **config>service>vprn>zone**
- **config>service>ies>zone**

Layer 2 and Layer 3 firewalls share system resources; that is, they share the maximum number of policies, profiles, and session ID space supported by the system.

A zone is created by adding at least one Layer 2 endpoint or Layer 3 interface to the zone configuration. Multiple zones can be created within each Layer 3 service or within the router context. Layer 2 services support only one zone. Layer 2 endpoints or Layer 3 interfaces from different services cannot be grouped into a single common zone. The following table lists the supported interfaces and endpoints that can be added to zones in each CLI context for NAT or firewall.

Table 32: Security zone interfaces and endpoints per context

CLI context	Interface/endpoint type	NAT	Firewall
Router	Layer 3	✓	✓
Epipe	SAP		✓
	Spoke-SDP termination		✓
VPLS	SAP		✓
	Spoke-SDP termination		✓
	Mesh SDP		✓
	EVPN		
VPRN	SAP	✓	✓
	Spoke-SDP termination	✓	✓
	IPSec private	✓	✓
	IPSec public	✓	
	Routed VPLS	✓	✓
IES	SAP	✓	✓
	Spoke-SDP termination	✓	✓
	IPSec public	✓	
	Routed VPLS	✓	✓



Note: A group of endpoints used for pseudowire redundancy cannot be added to a zone configured under an Epipe.

A zone configured under a Layer 2 service (VPLS or Epipe context) allows the 7705 SAR to perform Layer 3 firewall functionality on IPv4 packets.

NAT is not supported for zones configured under a Layer 2 service. A zone cannot be configured on a VPLS service with EVPN.

Unicast, multicast, and broadcast IPv4 packets are firewalled when they cross a Layer 2 service zone boundary.

Users can configure bypass policies to allow specific traffic, such as control plane protocols (OSPF, RIP, BGP, IGMP, PIM, LDP, RSVP, VRRP, DHCP, NTP, and so on) to bypass a firewall in a Layer 2 service. See the 7705 SAR Router Configuration Guide, "Bypass policies for firewalls in a Layer 2 service", for details. If not configured to bypass the zone, these packets are firewalled as normal unicast, multicast, or broadcast traffic and should be regulated by configuring firewall security policies for these protocols.

4.4 Frame relay VLL (Fpipe) services

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

- [Fpipe service overview](#)
- [Frame relay SAP-to-SAP service](#)
- [Frame relay traffic management](#)
- [Frame relay 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 about the adapter cards and chassis that support frame relay VLL services.

4.4.1 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 one-to-one 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 frame relay (Fpipe) pseudowire service in one-to-one mapping mode; see [Frame relay 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).

The following figure shows an example of the frame relay VLL end-to-end service.

Figure 49: FR VLL for end-to-end FR service



21995

Fpipe SAPs are supported on the following:

- 16-port T1/E1 ASAP Adapter card
- 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

Before an Fpipe can be configured on an MDA that supports multiple MDA modes, namely the 16-port T1/E1 ASAP Adapter card, 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**. See the 7705 SAR 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.

4.4.2 Frame relay 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.

4.4.3 Frame relay traffic management

Traffic management of frame relay VLLs is achieved through the application of ingress and egress QoS policies to SAPs.

4.4.3.1 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. See the 7705 SAR Quality of Service Guide for more information about the **de-1-out-profile** command.

4.4.4 Frame relay encapsulation

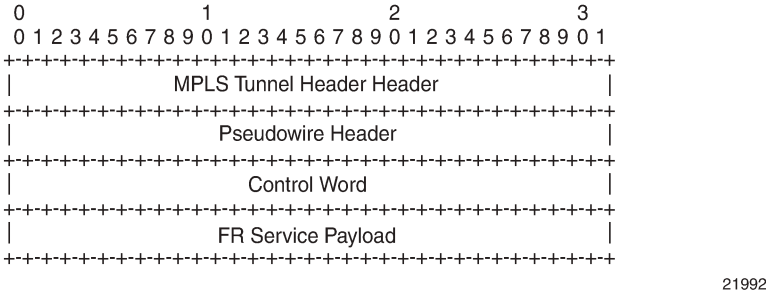
The following figure shows the standard FR frame.

Figure 50: FR frame



The following figure 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 51: 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 [Frame relay PW control word](#).

4.4.4.1 Frame relay PW control word

The control word (shown in the following figure) is required for FR one-to-one mode.

Figure 52: FR PW control word



The following table describes the bits used in the FR PW control word.

Table 33: Control word bit descriptions

Bit	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 Frame relay 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.

4.4.5 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*, 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.

4.5 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 about the adapter cards and chassis that support HDLC VLL services.

4.5.1 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
- 32-port T1/E1 ASAP Adapter card
- 12-port Serial Data Interface card
- T1/E1 ports on the 7705 SAR-M

Before HDLC SAPs can be configured on the 16-port T1/E1 ASAP Adapter card or 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 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 Hpipe. 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 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 Lpipe SAP. See [IP interworking VLL \(lpipe\) services](#) for more information.

4.5.2 HDLC VLL for end-to-end HDLC service

The following figure shows an example of how a mobile operator can deploy end-to-end HDLC services over an MPLS network.

Figure 53: HDLC VLL for end-to-end HDLC service



21994

In the figure, 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, Layer 2 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, see the PW reference diagram packet processing that supports the HDLC PW as described in RFC 4618.


4.5.3 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.

4.5.4 HDLC encapsulation

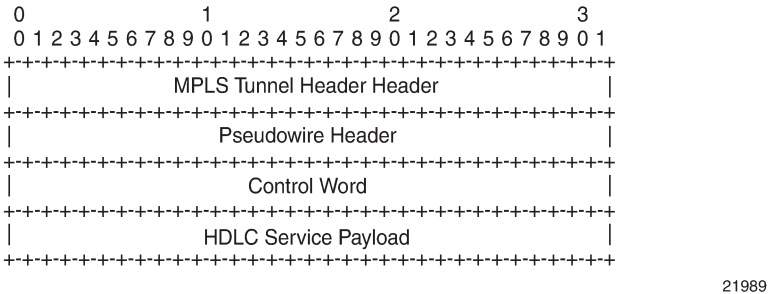
The following figure shows a typical HDLC VLL frame.

Figure 54: HDLC VLL frame



The following figure shows a typical HDLC VLL frame together with its MPLS tunnel encapsulation.

Figure 55: 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.


4.5.4.1 HDLC PW control word and payload size

HDLC VLLs support an optional control word (CW). The control word must be used for any HDLC (Hpipe) pseudowire service that transports packets that are less than 64 bytes.

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 the 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. The following figure shows the HDLC PW control word.

Figure 56: HDLC PW control word



The following table describes the bits used in the HDLC PW control word.

Table 34: Control word bit descriptions

Bit	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.

4.5.5 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.

4.5.5.1 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.

4.6 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 about the adapter cards and chassis that support IP interworking VLL services.

4.6.1 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.

The following figure 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 57: IP pseudowires between 7705 SAR nodes



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

4.6.2 IP interworking VLL datapath

To be able to forward IP packets between CE 1 and CE 3 in [Figure 57: IP pseudowires between 7705 SAR nodes](#), 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.

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.

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 lpipe service can have a maximum of two VLAN tags. Frames with more than two VLAN tags are dropped at the lpipe 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.

4.6.3 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](#)

4.6.3.1 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.

4.6.3.2 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 lpipe service.

4.6.3.3 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.

4.6.3.4 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.

4.6.3.5 CE IP address discovery for PPP/MLPPP

The address of a locally attached CE device can be learned via IPCP. If the lpipe 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.

4.6.3.6 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 lpipe. 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.

4.6.4 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.

4.6.5 Hardware support for interworking IP PWs

The following table lists the hardware that supports interworking IP PWs.

Table 35: 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
6-port Ethernet 10Gbps Adapter card	✓						
8-port Gigabit Ethernet Adapter card	✓						
10-port 1GigE/1-port 10GigE X-Adapter card							
— x1-10gb-sf+ mode							
— x10-1gb-sfp mode	✓						
16-port T1/E1 ASAP Adapter card							
— Channelized		✓			✓	✓	
— Clear channel		✓ ¹	✓	✓	✓	✓ ¹	
32-port T1/E1 ASAP Adapter card							
— Channelized		✓			✓	✓	
— Clear channel		✓ ¹	✓	✓	✓	✓ ¹	
4-port DS3/E3 Adapter card							

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
— 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		✓	✓				
4-port OC3/STM1 / 1-port OC12/STM4 Adapter card							
— on VT1.5/VC12		✓	✓	✓			
12-port Serial Data Interface card							
— on RS-232 ports							
— on V.35/X.21 ports		✓			✓	✓	
4-port SAR-H Fast Ethernet module	✓						
4-port T1/E1 and RS-232 Combination module		✓	✓	✓			
6-port SAR-M Ethernet module	✓						
7705 SAR-A	✓	✓	✓	✓			
7705 SAR-Ax	✓						
7705 SAR-H	✓						
7705 SAR-M	✓	✓ ¹	✓	✓	✓	✓ ¹	
7705 SAR-Wx	✓						

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
7705 SAR-X	✓	✓ ¹	✓	✓		✓ ¹	

Note:

1. Supported on framed DS1/E1 and unframed E1

4.6.6 Control word

IP interworking VLLs support an optional control word. The control word may be required to interoperate with devices at the far end. If the control word is enabled, it is set to all zeros and ignored on egress.

See [Pseudowire control word](#) for more information.

4.6.7 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](#)

4.6.7.1 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.


4.6.7.2 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 lpipe payload represented in the IP Datagram field in the following figure.

Figure 58: FR header with NLPID support for IP interworking PW



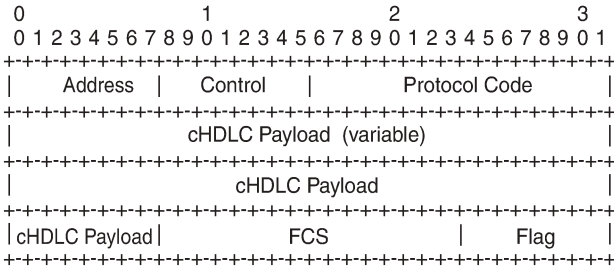
21997

4.6.7.3 cHDLC termination

The following features are supported for cHDLC termination:

- ports must be configured for access mode according to the cHDLC frame format. See the 7705 SAR Interface Configuration Guide for information about how to configure the **encap-type** command to **cisco-hdlc**.
- IP encapsulation (0x0800). [Figure 59: cHDLC header frame](#) shows the cHDLC header frame.
- keepalive support using SLARP (0x8035). [Figure 60: SLARP keepalive frame](#) shows the SLARP keepalive frame.
- configurable service MTU of 1 to 2048 bytes. When encapsulated in the IPv4 PDU, the cHDLC packet shown in the following figure is modified to remove the flag, address, control, protocol ID, and FCS.

Figure 59: cHDLC header frame

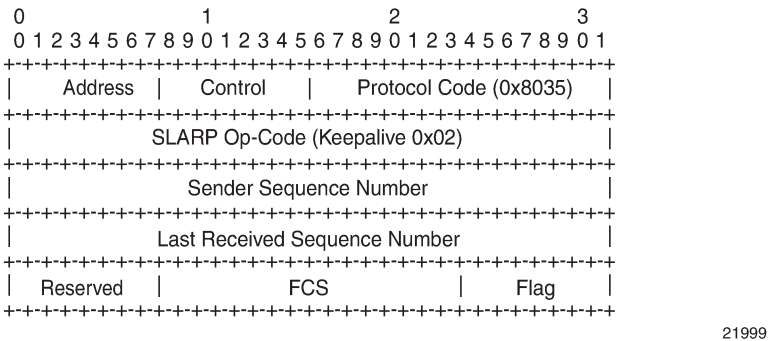


21998



Note: cHDLC encapsulation cannot be used to transmit HDLC frames into an lpipe.

Figure 60: SLARP keepalive frame



See the 7705 SAR Interface Configuration Guide for more information about how to configure access ports for interworking IP PWs.

4.6.7.4 PPP and MLPPP termination on 2-port OC3/STM1 Channelized Adapter cards


PPP termination is supported on DS1/E1 and *n* x DS0 (64 kb/s) channels.

MLPPP termination is supported on DS1/E1 channels.

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.

The following figure shows an example of a network using PPP and MLPPP encapsulation to IP PWs on a 2-port OC3/STM1 Channelized Adapter card.

Figure 61: 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.

See the 7705 SAR 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.

4.6.8 Traffic management

4.6.8.1 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.

4.6.9 Status signaling

4.6.9.1 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

4.6.10 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 email 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 lpipe 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 lpipe 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. See the 7705 SAR Router Configuration Guide, "Filter policies", for information about configuring filter policies and assigning them to network interfaces.

4.7 Pseudowire switching

Topics in this section include:

- [Overview](#)
- [Pseudowire switching with pseudowire redundancy](#)
- [Pseudowire switching behavior](#)
- [Pseudowire switching TLV](#)

4.7.1 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 62: Simplex to simplex pseudowire switching](#) 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.

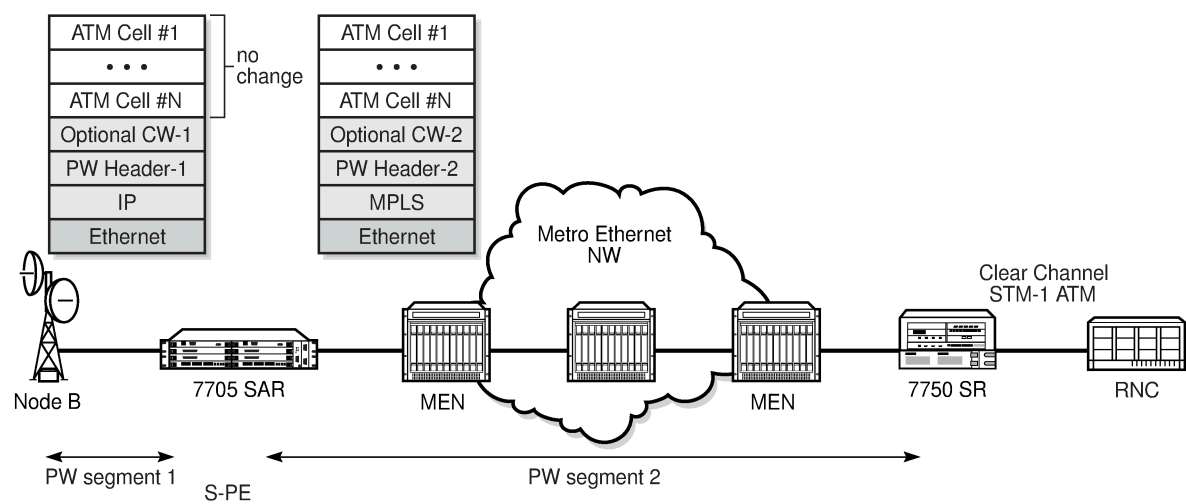
The switching pseudowire inserts a pseudowire switching point TLV indicating the system address in the label mapping message. This TLV is useful because it allows for troubleshooting of the path of the pseudowire, especially if multiple pseudowire switching points exist between the two PEs. The TLV also 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 toward 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, because 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.

The following figure 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 62: Simplex to simplex pseudowire switching



20772

Configuration examples can be found in [Configuring PW switching](#).

The following table shows the pseudowire switching options supported on the 7705 SAR.

Table 36: 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-cs	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.

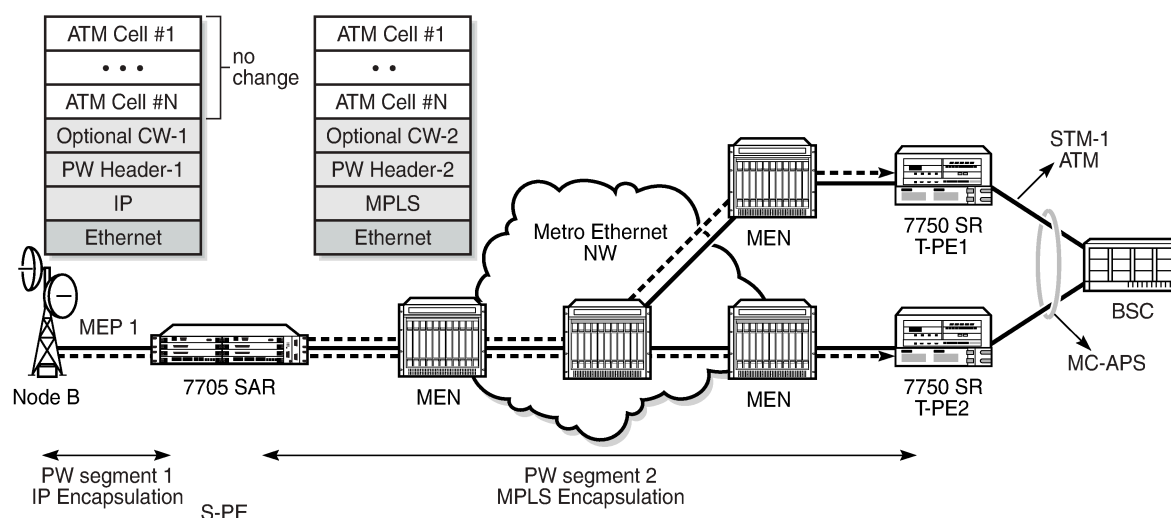
The 7705 SAR supports use of an optional control word on both pseudowire segments. Use of a control word is negotiated by the 7705 SAR and 7750 SR during the signaling phase. The 7705 SAR and 7750 SR negotiate during the signaling phase even if a control word is not used. If a control word is used, the 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.

4.7.2 Pseudowire switching with pseudowire redundancy

Pseudowire switching with pseudowire redundancy supports one redundant pseudowire with up to four redundant spoke SDPs. The following figure shows an example of a network with simplex to redundant pseudowire switching.

Figure 63: Simplex to redundant pseudowire switching




20773

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.

4.7.3 Pseudowire switching behavior

In the network in the following figure, T-PE nodes act as leading nodes and pseudowire switching nodes (S-PE nodes) act as followers 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 64: Pseudowire switching network



20905

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 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 spoke SDPs is up, the S-PE passes any received SAP or SDP-binding generated status notification message unchanged. For example, the status notification TLV is unchanged but the VC-ID in the FEC TLV is set to the value of the pseudowire segment to the next hop.
- When the local operational status for any of the spoke SDPs 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.

4.7.3.1 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.

4.7.4 Pseudowire switching TLV

The following figure shows the format of the pseudowire switching TLV and the table describes the fields.

Figure 65: Pseudowire switching TLV



21834

Table 37: 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

4.8 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 \(standby signaling\)](#)

4.8.1 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-M

Ethernet

Ethernet VLL service is supported on the following:

- 6-port Ethernet 10Gbps Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card
- Ethernet ports on the 7705 SAR-A
- Ethernet ports on the 7705 SAR-Ax
- Ethernet ports on the 7705 SAR-H
- Ethernet ports on the 7705 SAR-Hc
- Ethernet ports on the 7705 SAR-M
- Ethernet ports on the 7705 SAR-Wx
- Ethernet ports on the 7705 SAR-X
- 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
- 32-port T1/E1 ASAP Adapter card
- T1/E1 ports on the 7705 SAR-M

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
- 8-port C37.94 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
- RS-232 serial ports on the 7705 SAR-Hc
- T1/E1 ports on the 7705 SAR-M
- 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
- 32-port T1/E1 ASAP Adapter card
- T1/E1 ports on the 7705 SAR-M

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 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)
The 7705 SAR supports PPP to IP PW interworking and cHDLC to IP PW interworking on framed DS1/E1 and unframed E1.
- 32-port T1/E1 ASAP Adapter card (on PPP/MLPPP connections)
The 7705 SAR supports PPP to IP PW interworking and cHDLC to IP PW interworking on unframed E1.
- 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card (on PPP/MLPPP connections over DS1/E1 channels)
- all ports on the 7705 SAR-A (on PPP/MLPPP connections on the T1/E1 ports)
- 7705 SAR-H
- Ethernet ports on the 7705 SAR-Ax
- 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)
- Ethernet ports on the 7705 SAR-Wx
- T1/E1 ports on the 7705 SAR-X
The 7705 SAR-X supports PPP to IP PW interworking and cHDLC to IP PW interworking on framed DS1/E1 and unframed E1.
- 4-port SAR-H Fast Ethernet module
- 6-port SAR-M Ethernet module
The 7705 SAR supports PPP to IP PW interworking and cHDLC to IP PW interworking on unframed E1.



Note: MPLS and VLL service over MPLS are not supported on access ports.

4.8.2 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

4.8.2.1 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.

4.8.3 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 $N \geq 1$), 0x000A (for ATM N-to-1 VPC cell mode, where $N = 1$) and 0x0003 (for ATM transparent cell transport mode). See 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.

4.8.4 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.

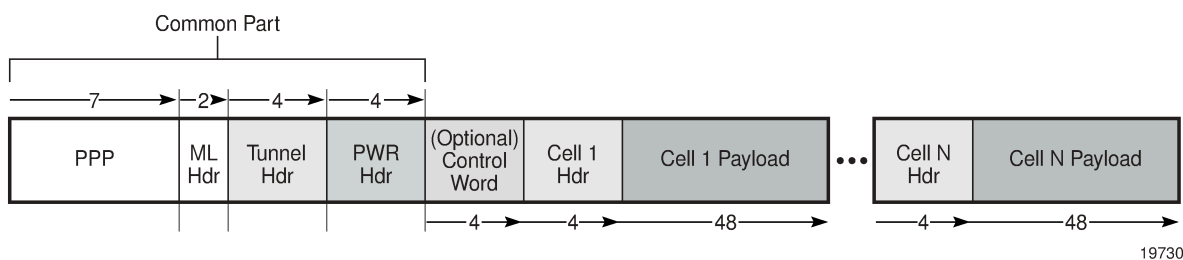
For VPCs, N is not user-configurable and must be equal to 1 ($N = 1$). VPC cell mode supports the 0x000A (ATM N-to-1 VPC Cell Mode) PW service type.

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.

The following figure shows the structure of an N-to-1 cell mode frame.

Figure 66: N-to-1 cell mode encapsulation



4.8.4.1 N-to-1 cell mode encapsulation (N = 1)

The following sections describe ATM cell mode encapsulation where N = 1.

4.8.4.1.1 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 groups 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.

4.8.4.1.2 VPI/VCI translation (N = 1)

Before traffic from different cell sites can be switched to an RNC, VPI and VCI translation may be required 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.

4.8.4.1.3 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 correct 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.

4.8.4.1.4 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. Apiece rate-limiting on egress is controlled by the traffic-descriptor profile, not the QoS policy queue rates.

4.8.4.1.5 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.

4.8.4.2 N-to-1 cell mode encapsulation (N > 1)

The following sections describe ATM cell mode encapsulation where N > 1.


4.8.4.2.1 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.

The following figure shows a typical network where N > 1 is deployed.

Figure 67: VCC cell mode encapsulation with N > 1



22321

4.8.4.2.2 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.

4.8.4.2.3 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—regardless 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.

4.8.4.2.4 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 may have 2 cells, SAP2 may have 5 cells, SAP3 may have 0 cells, and SAP4 may 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.

4.8.4.2.5 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.

4.8.4.2.6 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-M (variants with T1/E1 ports)

4.8.4.2.7 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.

4.8.5 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.

4.8.5.1 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.

4.8.5.1.1 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**.

4.8.5.1.2 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.

4.8.5.1.3 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.

4.8.5.2 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. See the 7705 SAR 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.

4.8.5.3 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*** command provides user-configurable controls for the interval and rate of statistics collection. See the 7705 SAR 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.

4.8.5.4 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 TLV 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.

4.8.6 QoS policies

When applied to 7705 SAR Apipe, Cpipe, Epipe, and Lpipe services, service ingress QoS policies only create the unicast queues defined in the policy.

With Apipe, Cpipe, Epipe, and Lpipe 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 Lpipe service. QoS policies on Apipes cannot perform any classification.

4.8.7 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 Ilc-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). See the 7705 SAR Router Configuration Guide, "Filter policies", for information about configuring IP filters.

4.8.8 MTU settings

There are several MTU values that must be set properly for a VLL service (Apipe, Cpipe, Epipe, Fpipe, Hpipe, or Lpipe) to work from end to end. The following figure locates the MTU point for each value and the table 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 68: MTU points on the 7705 SAR

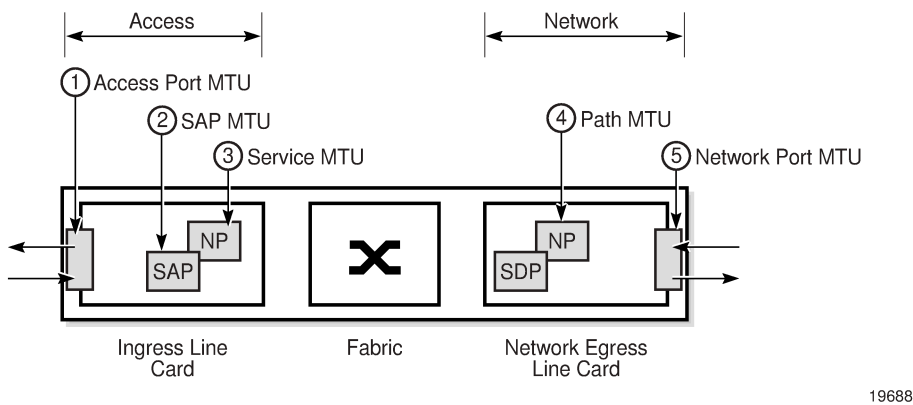


Table 38: MTU points and descriptions

	MTU point	Description	Command
1	Access port MTU	<p>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:</p> <ul style="list-style-type: none">• 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, TDM, serial, or SONET/SDH

	MTU point	Description	Command
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 68: MTU points on the 7705 SAR , 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, 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: <ul style="list-style-type: none"> 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
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)

The following table 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 39: 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 ²	Epipe, VPLS	1514	1	9670 ³	9724 ³
			lpipe	1500			
GigE SFP ¹	Access	null, dot1q, qinq ²	Epipe, VPLS	1514	1	9670 ³	9724 ³
			lpipe	1500			

Notes:

1. The maximum MTU value is supported only on cards that have buffer chaining enabled.
2. QinQ is supported only on access ports.
3. On the Packet Microwave Adapter card, the MWA ports support 4 bytes less than the Ethernet ports. MWA ports support a maximum MTU of 9720 bytes (null) or 9724 bytes (dot1q). MWA ports do not support QinQ.

For more information about port MTU, see "MTU configuration guidelines" in the 7705 SAR Interface Configuration Guide.

The following tables can be used as aids in calculating MTU values for various configurations and operating scenarios.

Table 40: MTU calculator – service creation (worst case) access ports and SAPs

Service creation									
		Access port default MTU		SAP					
		TDM/ATM	Eth	Epipe	lpipe	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								

Service creation									
		Access port default MTU		SAP					
		TDM/ATM	Eth	Epipe	lpipe	Apipe	Cpipe	Fpipe	Hpipe
	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 41: MTU calculator – service creation (worst case) network ports

Service creation									NW
		Network port default MTU							
							Epipe over MPLS encap	Epipe/lpipe 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

Service creation									NW
		Network port default MTU							
							Epipes over MPLS encap	Epipes/lpipes over GRE	LSR
		PPP	ML/PPP	Eth null	Eth dot1q	MPLS label	Best case	Worst case	Worst case
	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:**

- 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.

4.8.8.1 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. See 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. The following table shows the values that are expected to match.

Table 42: Matching MTU or payload values for signaled VLL services

	Apipe	Cpipe	Epipe	Fpipe	Hpipe	Lpipe
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

4.8.9 QinQ (VLL service)

Epipe and Lpipe VLL services support QinQ functionality. For details, see [QinQ support](#).

4.8.10 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:

- Apipes (ATM VLLs)
- Epipes (Ethernet VLLs)
- Cpipes (TDM VLLs in SAToP and CESoPSN circulation emulation mode)
- Fpipes (frame relay VLLs in one-to-one mapping mode)
- Hpipes (HDLC VLLs)
- Lpipes (IP interworking VLLs)

For Apipes, the control word is optional. It can be enabled to guarantee ordered packet/cell delivery.

For Epipes (with the exception of MEF 8 services) and Lpipes, the control word is optional. If it is enabled, it will be set to all zeros and ignored on egress.

For Cpipes, Epipe MEF 8 services, and Fpipes, the control word is mandatory and cannot be configured.

For Hpipes, the control word is optional when transporting packets that are more than 64 bytes but mandatory when transporting packets that are less than 64 bytes.

When 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 word is set to the length of the packet. Otherwise, the length field is set to 0. 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. If the control word is not set for packets less than 64 bytes, 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.

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 (Cpipes), Epipe MEF 8 services, and Fpipes always require the control word.)
- If the received Label Mapping message has C = 1 and the PW is locally configured such that the use of the control word is not supported, 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).

4.8.11 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 the following figure, 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 particular network. If 7x50-1 loses all of its connectivity to the network or to the RNC, 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), and T1/E1 multi-chassis line card redundancy (MC-LCR)
- PW redundancy with VPLS or PW switching

For more information about MC-LAG, see the 7705 SAR Interface Configuration Guide, "Multi-chassis LAG". For more information about MC-APS, see the 7705 SAR Interface Configuration Guide, "Automatic protection switching". For more information about MC-LCR, see the 7705 SAR Interface Configuration Guide, "T1/E1 line card redundancy".

Figure 69: Pseudowire redundancy



20225

PW redundancy applies to all VLL services available on the 7705 SAR: Apipe, Cpipe, Epipe, Fpipe, Hpipe, and Ipipe.

4.8.11.1 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. The following figure 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 the figure, implicit endpoints (a) 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 the figure, explicit endpoints (b) 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 70: Implicit and explicit endpoint objects



20226

The 7705 SAR supports the following types of endpoint objects:

- SAP – there can be only one SAP per PW endpoint (Endpoint X in the figure)
- 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 the figure).
- 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 the figure). If a primary spoke SDP is defined, 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 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 specific 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.

4.8.11.2 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.

4.8.11.2.1 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 to the primary path or to never revert to the primary path. When the primary PW goes down, the 7705 SAR selects the secondary spoke SDP that is operationally up and has the highest precedence setting.


4.8.11.2.2 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 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 the following figure, 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 71: Pseudowire redundancy with four spoke SDPs



20227


4.8.11.3 PW redundancy and inter-chassis backup

Inter-chassis backup (ICB) spoke SDPs are supported for use with Cpipe services in an MC-APS or MC-LCR 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, MC-LCR, or MC-LAG node to a protection node.

A failure on the access side triggers an access side MC-APS, MC-LCR, or MC-LAG switchover and a network-side pseudowire switchover. A failure on the network side triggers a pseudowire switchover but does not trigger an MC-APS, MC-LCR, or MC-LAG switchover.

The following figure shows a network experiencing an access side failure in an MC-LAG scenario.

Figure 72: Access side failure with ICB protection




23496

If the active access link of an MC-APS, MC-LCR, or MC-LAG group fails, the MC-APS, MC-LCR, 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.

The following figure shows a network experiencing a network side failure in an MC-LAG scenario.

Figure 73: Network side failure with ICB protection



23497

If the active network link of an MC-APS, MC-LCR, or MC-LAG group fails, pseudowire redundancy causes the network side to switch to the standby pseudowire. This does not trigger an MC-APS, MC-LCR, or MC-LAG switchover. If there is no ICB, the end-to-end transmission will be lost because 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, a SAP can be added only if it is part of an MC-LAG, MC-APS, or MC-LCR group. Similarly, if an MC-LAG, MC-APS, or MC-LCR SAP is added to an endpoint, the only other possible addition to that endpoint is an ICB spoke SDP.

4.8.11.4 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.

4.8.12 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 the following figure, 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 74: Active/standby mode for redundant pseudowires



20694

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.

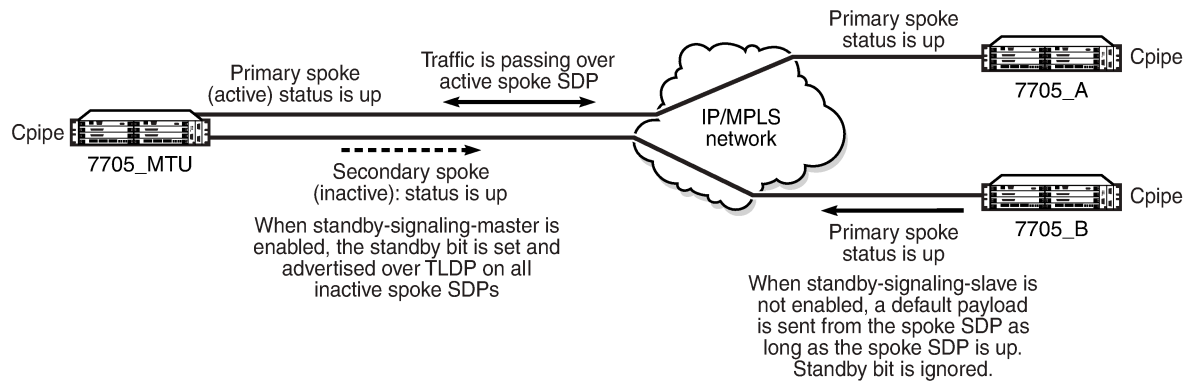
Standby signaling has two components: standby-signaling master and standby-signaling slave. [Figure 75: Standby-signaling-master enabled](#) shows an example where **standby-signaling-master** has been enabled on the 7705_MTU node. A Cpipe has been configured between the 7705_MTU and the 7705_A node and between the 7705_MTU and the 7705_B node. The spoke SDP toward 7705_A is the active PW (precedence primary), while the spoke SDP toward 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 131:556 endpoint "stdbyMaster" create
    precedence primary
  exit
```

Figure 75: Standby-signaling-master enabled



22796



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 (7705_B in the figure) ignores the bit and continues to transmit data toward the 7705_MTU as long as the PW is up; therefore, pseudowire redundancy does not work.

To stop data from being transmitted along standby spoke SDPs, the far-end endpoints must be enabled for **standby-signaling-slave**. [Figure 76: Standby-signaling-slave enabled](#) 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:

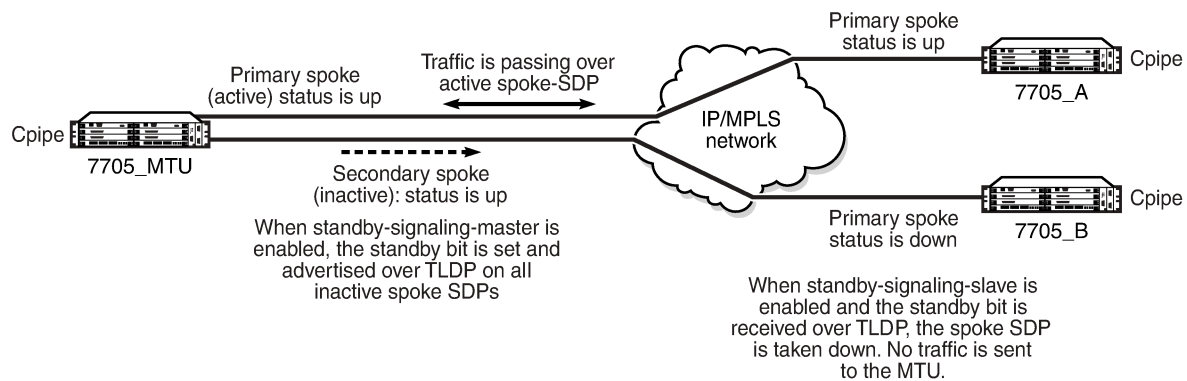
CLI syntax:

```
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 76: Standby-signaling-slave enabled



22797

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.

4.8.12.1 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 instead of 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.

4.8.12.1.1 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, 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 does not block its end of the PW and forwards 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.

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 option ensures that forwarding on a standby PW will be bidirectionally blocked to prevent PWs in standby mode from transmitting traffic via label withdrawal.

4.8.12.1.2 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 instead of 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 withdraws 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.

4.8.12.2 Pseudowire redundancy on serial data interface ports

Pseudowire redundancy is only supported on serial data interface ports when there is a Cpipe on a standby-signaling master with a single SAP and a endpoint with up to two spoke SDPs. The far-end slaves (standby-signaling slaves) each have a single SAP and a single spoke SDP back to the master.

Subrate speeds (< 64 kb/s) on RS-232 and X.21 ports are supported using HCM. HCM cannot determine signal quality until a circuit is established (that is, both endpoints of the circuit are connected); therefore, when **standby-signaling-slave** is enabled on subrate circuits, HCM framing will always be down on the inactive slave. The normal behavior is for the slave to send the port status to the master using the pseudowire status bit, indicating local attachment circuit (LAC) Tx/Rx faults. Because the slave cannot clear these faults, this prevents the master from switching back to the primary pseudowire as soon as possible (pseudowire redundancy reversion).

To enable pseudowire redundancy reversion in this case, the sending of LAC Tx/Rx fault messages from the slaves to the master is suppressed on RS-232 and X.21 ports configured for subrate speeds. The master is therefore not aware that the far-end port is down due to HCM being down and will switch back to the primary pseudowire as soon as other types of alarms are cleared.

These limitations apply to:

- RS-232 and X.21 ports on the 12-port Serial Data Interface card
- RS-232 ports on the 4-port T1/E1 and RS-232 Combination module (on the 7705 SAR-H)

- RS-232 ports on the 7705 SAR-Hc

4.9 Configuring a VLL service with CLI

This section provides the information required to configure virtual leased line (VLL) services using the CLI.

Topics in this section include:

- [Common configuration tasks](#)
- [Configuring VLL components](#)
- [Service management tasks](#)

4.10 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 the **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.

4.11 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 an example of the 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](#)

4.11.1 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#
```

4.11.1.1 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
-----

```

4.11.1.2 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#
```

4.11.1.3 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
...
-----
```

4.11.1.4 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
...
-----

```

4.11.2 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#
```

4.11.2.1 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]
    [threshold-repeat uSecs]
      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. See the 7705 SAR Interface Configuration Guide, "Configuration Command Reference", for information about 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
```

4.11.2.2 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#
```

4.11.3 Creating an Epipe service

This section describes how to create an Epipe service. Additional topics in this section include:

- [Configuring Epipe SAP parameters](#)
- [Configuring an Epipe with an ATM SAP](#)
- [Configuring Epipe SAP MEF 8 parameters](#)
- [Configuring Epipe SAP microwave link parameters for interworking with TDM2Ethernet](#)
- [Configuring ATM encapsulation under Epipe service \(7705 SAR-M only\)](#)
- [Configuring Epipe spoke SDP bindings](#)
- [Configuring a security zone within an Epipe](#)

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
-----
```

4.11.3.1 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 information about configuring ETH-CFM parameters on an Epipe SAP, see [ETH-CFM \(802.1ag and Y.1731\) tasks](#).

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
...
-----

```

4.11.3.2 Configuring an Epipe with an ATM SAP

Use the following CLI syntax to configure an Epipe with an ATM SAP. An OC3 clear channel ATM path is configured and then the Epipe SAP is configured.

CLI syntax:

```

configure port port-id
    description description
    sonet-sdh
        path
            atm
                cell-format uni
            mtu mtu
            no shutdown
        no shutdown

```

The following example shows an OC3 clear channel ATM path configuration and output.

Example:

```
*A:Sar8 Dut-B# configure port 1/2/3
config>port# description "1/OC03/KAN"
config>port# sonet-sdh
config>port>sonet-sdh# path
config>port>sonet-sdh>path# atm
config>port>sonet-sdh>path>atm# cell-format uni
config>port>sonet-sdh>path>atm# exit
config>port>sonet-sdh>path# mtu 2090
config>port>sonet-sdh>path# no shutdown
config>port>sonet-sdh>path# exit
config>port>sonet-sdh# exit
```

```
*A:Sar8 Dut-B>config>port# info
-----
description "1/OC03/KAN"
sonet-sdh
  path
    atm
      cell-format uni
      mtu 2090
      exit
      no shutdown
    exit
  exit
exit
no shutdown
```

CLI syntax:

```
config>service# epipe service-id [customer customer-id] create
description description-string
sap sap-id create
  description description-string
  egress
    qos policy-id
  no shutdown
  ingress
    qos policy-id
service-mtu octets
service-name service-name
spoke-sdp sdp-id:vc-id create
  no shutdown
no shutdown
```

The following example shows an Epipe SAP configuration and output.

Example:

```
*A:Sar8 Dut-B# configure service epipe 202165 customer 1 create
config>service>epipe$ description "00/ARDU/902/863/2949"
config>service>epipe$ sap 1/2/3:0/49 create
config>service>epipe>sap$ description "00/ARDU/902/863/2949"
config>service>epipe>sap$ egress
config>service>epipe>sap>egress$ qos 1421
config>service>epipe>sap>egress$ exit
config>service>epipe>sap$ no shutdown
config>service>epipe>sap$ ingress
config>service>epipe>sap>ingress$ qos 1421
config>service>epipe>sap>ingress$ exit
config>service>epipe>sap$ exit
config>service>epipe$ service-mtu 1622
```

```

config>service>epipe$ service-name "EPIPE 202165"
config>service>epipe$ spoke-sdp 3080:202165 create
config>service>epipe>spoke-sdp$ no shutdown
config>service>epipe>spoke-sdp$ exit
config>service>epipe$ no shutdown
config>service>epipe$ exit

```

```

*A: Sar8 Dut-B>config>service>epipe$ info
-----
description "00/AR DU/902/863/2949"
  service-mtu 1622
  service-name "EPIPE 202165"
  sap 1/2/3:0/49 create
    description "00/AR DU/902/863/2949"
    egress
      qos 1421
    exit
    ingress
      qos 1421
    exit
    no shutdown
  exit
  spoke-sdp 3080:202165 create
    description "00/AR DU/902/863/2949"
    no shutdown
  exit
  no shutdown
exit

```

4.11.3.3 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 about 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 about 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 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 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#

```

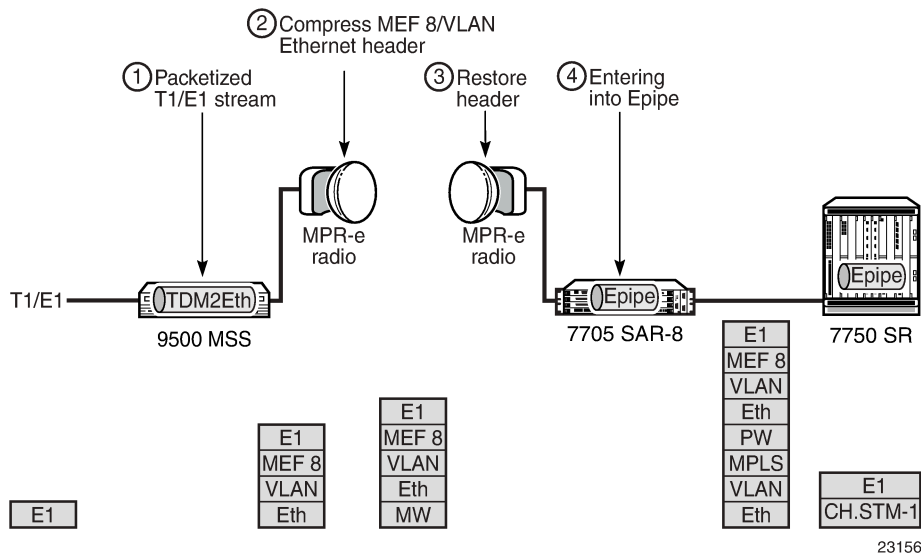
4.11.3.4 Configuring Epipe SAP microwave link parameters for interworking with TDM2Ethernet

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 77: Mixed microwave link scenario](#).

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 77: 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]
```

4.11.3.5 Configuring ATM encapsulation under Epipe service (7705 SAR-M only)

The 7705 SAR-M can provide the transport of Ethernet traffic over an ATM network through the configuration of a SAP-to-SAP Ethernet PW. For example, as shown in [Figure 78: Ethernet-to-ATM interworking on the 7705 SAR-M](#), the 7705 SAR-M provides radio access of Ethernet traffic into an existing ATM network.

The SAP facing the eNodeB is configured as a null- or dot1q-encapsulated Ethernet SAP. The SAP facing the ATM network is ATM/IMA-encapsulated.

ATM/IMA is configured on E1 ports of the 7705 SAR-M. The uplink port must be in access mode in order to host the ATM/IMA SAP. The 7705 SAR-M supports IMA groups with 1 to 16 member links as well as ATM encapsulation on a single E1 port.

The following exceptions on ATM encapsulation under an Epipe service apply:

- ATM SAPs are allowed on E1 and E1 IMA ports only
- ETH-CFM OAM on an Ethernet SAP is not supported
- ACL filtering is not supported
- QinQ encapsulated Ethernet SAP interworking to an ATM SAP is not supported
- ATM policing on an ATM SAP is not supported




Note: In addition to supporting ATM-encapsulated SAPs under an Epipe service, the 7705 SAR-M also supports:

- the termination of ATM-encapsulated SAPs to IES for node management purposes. If the only transport uplink for a site is ATM, it is necessary to offer in-band management over the existing ATM network.
- VC-mux routed PDU encapsulated SAPs for node management. The SAP must be attached to an IES service. No IP forwarding is supported on this management SAP.

For the application described in this section, an ATM-encapsulated SAP to IES is used for managing the 7705 SAR-M at the cell site. See [IES for in-band management](#) for details.

Figure 78: Ethernet-to-ATM interworking on the 7705 SAR-M



26463

The following examples show the Epipe SAP-to-SAP configurations for Ethernet-to-ATM interworking. The first example shows the ATM SAP configuration (ATM VPI/VCI is 30/80), and the second example shows the IMA configuration (IMA bundle). In both examples, SAP 1/1/2:2 is the Ethernet SAP facing the eNodeB.

Example:

```
epipe 105 customer 1 create
  sap 1/1/2:2 create
  exit
  sap 1/2/3.1:30/80 create
  exit
  no shutdown
exit
```

Example:

```
epipe 105 customer 1 create
  sap 1/1/2:2 create
  exit
```

```
sap bundle-ima-1/2.1:20/200 create
    atm
    exit
exit
no shutdown
exit
```

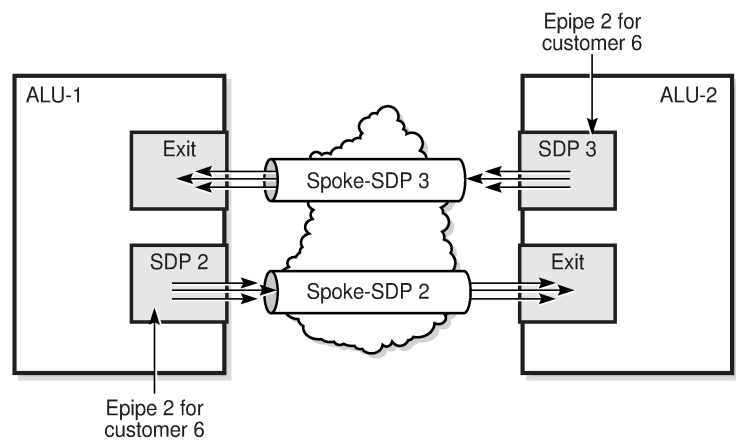
When ATM encapsulation under an Epipe service is used, the AAL-5 encapsulation type for the ATM SAP is "aal5mux-bridged-eth-nof*", as shown in the following **show>service>id service-id>sap sap-id detail** screen.

```
-----
ATM SAP Configuration Information
-----
Ingress TD Profile : 1                      Egress TD Profile : 1
Ingress TD Ovr    : N/A                    Egress TD Ovr    : N/A
Alarm Cell Handling: Enabled
OAM Termination   : Disabled               Periodic Loopback : Disabled
AAL-5 Encap       : aal5mux-bridged-eth-nof*
-----
```

4.11.3.6 Configuring Epipe spoke SDP bindings

The following figure 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 79: SDPs – unidirectional tunnels



19484

An SDP must exist before it can be used with a spoke SDP. For SDP configuration information, see [Configuring SDPs](#).
For information about configuring ETH-CFM parameters on a spoke SDP, see [ETH-CFM \(802.1ag and Y.1731\) tasks](#).

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
    pipe 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
    ...
    -----

```

4.11.3.7 Configuring a security zone within an Epipe

To configure firewall security functionality, you must:

- configure a 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 and specify the match criteria and the action to be applied to a packet if a match is found
- configure a firewall bypass policy in the **config>security** context and assign it to the Epipe, as shown in the following CLI syntax. This step is optional. If you do not configure a bypass policy, the protocol packets are firewalled based on the firewall security policies.
- configure a security zone and apply the policy ID to the zone, as shown in the following CLI syntax

CLI syntax:

```

config>service
    epipe service-id [customer customer-id] [create]
        fw-bypass-policy {bypass-id | name}
        zone zone-id [create]
        abort
        begin
        commit
        description description-string
        inbound

```

```

outbound
policy {policy-id | policy-name}
sap sap-id
shutdown
spoke-sdp sdp-id:vc-id

```

The following example displays the security zone configuration output.

```

*A: Sar8 Dut-A>config>service>epipe# info
-----
      stp
        shutdown
      exit
      fw-bypass-policy 1
      sap 1/2/2 create
        no shutdown
      zone 1 create
        name "Epipe zone"
        description "Sample zone"
        sap "1/2/3"
        policy "4"
        inbound
          limit
        exit
      exit
      outbound
        limit
      exit
      exit
      commit
    exit
  no shutdown
-----

```

4.11.4 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#
```

4.11.4.1 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
...
-----
```

4.11.4.2 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>fpipespoke-sdp 1:1 create
A:ALU-41>config>service>fpipespoke-sdp# no shutdown
A:ALU-41>config>service>fpipespoke-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
      spoke-sdp 1:1 create
      exit
      no shutdown
    exit
...
-----
A:ALU-41>config>service#

```

4.11.5 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
```

4.11.5.1 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
...
-----
```

4.11.5.2 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#
```

4.11.6 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 lpipe configuration:

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

4.11.6.1 Configuring lpipe SAP parameters

The following example shows an lpipe SAP configuration:

```
A:ALU-48>config>service# info
-----
...
    ipipe 202 customer 1 create
        sap 1/1/2:444 create
            description "eth_ipipe"
            ce-address 172.16.0.0
        exit
        spoke-sdp 16:516 create
            ce-address 172.16.0.1
        exit
        no shutdown
    exit
...
-----
```

The following shows a PPP to Ethernet local lpipe 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 172.16.0.0
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 172.16.0.1
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 172.16.0.0
    exit
    spoke-sdp 15:516 create
        ce-address 172.16.0.1
    exit
    exit
    no shutdown
    exit
exit
-----

```

The configuration of a Cisco HDLC SAP on an lpipe is similar to the above example, except that the *sap-id* has the form *slot/mdalport.channel-group* (1/2/2.1).

4.11.6.2 Configuring lpipe SDP bindings

The following shows an lpipe SDP configuration example:

```

A:ALU-48>config>service# info
-----
...
    sdp 16 mpls create
        far-end 10.4.4.4
        ldp
        path-mtu 1600
        keep-alive
        shutdown
    exit
    no shutdown
    exit
...
    lpipe 207 customer 1 create
        shutdown
        sap 1/1/2:449 create
            description "Remote_Ipipe"
            ce-address 172.16.0.10
        exit
        spoke-sdp 16:516 create
            ce-address 172.16.0.11
        exit

```

4.11.7 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.

When 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: SVCNMR #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
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
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
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
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
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-el 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
...
-----

```

4.11.8 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 lpipe 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
-----
```

4.11.9 Using the control word

The control word is mandatory for Cpipe SAToP and CESoPSN configurations, Epipe MEF 8 configurations, and Fpipe one-to-one mapping configurations.

The control word is optional for Apipe, Epipe (non-MEF 8), Hpipe, and lpipe services, but must be enabled for Hpipe pseudowire services when transporting packets that are less than 64 bytes. If the control word is enabled for Epipe or lpipe services, it will be set to all zeros and ignored on egress.

When the control word is enabled, the admin control word is set to preferred (enabled). 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, Epipe MEF 8, or Fpipe services. To disable the control word option on Apipe, Epipe (non-MEF 8), Hpipe, or Lpipe services, use the **no control-word** command.

Example:

```
config>service>apipe# spoke-sdp 1:2001 no control-word
config>service>apipe>spoke-sdp$ exit
```

4.11.10 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#

```

4.11.10.1 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
```

4.11.10.2 Configuring PW redundancy – ICB

Inter-chassis backup (ICB) spoke SDPs are supported for use with Cpipe services in an MC-APS configuration or MC-LCR 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, MC-LCR, or MC-LAG node to a protection node.

To configure ICB, create explicit endpoints on the access and network side on the working 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
```

The example below shows how to configure ICB on the working node in an MC-APS scenario:

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 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

```

The example below shows how to configure ICB on the protection node in an MC-APS scenario:

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 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

```

The CLI output below shows an example of ICB configured on the working node in an MC-LCR scenario:

```

cpipe 1 customer 1 vc-type cesopsn create
    description "Default Cpipe description for service id 1"
    service-mtu 1000
    endpoint "X" create
    exit
    endpoint "Y" create
    exit
    sap lcr-1/1.1 endpoint "X" create
        description "Default sap description for service id 1"
        no shutdown
    exit
    spoke-sdp 6003:1 endpoint "Y" create
        no shutdown
    exit
    spoke-sdp 6001:1001 endpoint "X" icb create
        no shutdown
    exit
    spoke-sdp 6001:2001 endpoint "Y" icb create
        no shutdown

```

```

        exit
    no shutdown
exit

```

The CLI output below shows an example of ICB configured on the protection node in an MC-LCR scenario:

```

cpipe 1 customer 1 vc-type cesopsn create
    description "Default Cpipe description for service id 1"
    service-mtu 1000
    endpoint "X" create
    exit
    endpoint "Y" create
    exit
    sap lcr-1/1.1 endpoint "X" create
        description "Default sap description for service id 1"
        no shutdown
    exit
    spoke-sdp 1003:1 endpoint "Y" create
        no shutdown
    exit
    spoke-sdp 1006:1001 endpoint "Y" icb create
        no shutdown
    exit
    spoke-sdp 1006:2001 endpoint "X" icb create
        no shutdown
    exit
    no shutdown
exit

```

4.12 Service management tasks

The service management tasks are similar for Apipe, Cpipe, Epipe, and Lpipe services. This section discusses the following service management tasks:

- [Modifying service parameters](#)
- [Disabling a service](#)
- [Re-enabling a service](#)
- [Deleting a service](#)

4.12.1 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-42>config>service>apipe>spoke-sdp# ingress
A:ALU-42>config>service>apipe>spoke-sdp>ingress# vc-label 2043
A:ALU-42>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#

```

4.12.2 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#
```

4.12.3 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 Lpipe 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
```

4.12.4 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 Lpipe 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
```

4.13 VLL services command reference

4.13.1 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
 - Epipe security configuration commands
 - Fpipe service configuration commands
 - Hpipe service configuration commands
 - Lpipe service configuration commands
- Show commands
- Clear commands

4.13.1.1 VLL services configuration commands

4.13.1.1.1 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
- 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).

4.13.1.1.2 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
  - 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).

4.13.1.1.3 Cpipe service configuration commands

```

config
- service
  - cpipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-type {satop-e1
| satop-t1 | satop-e3 | satop-t3 | satop-serial | satop-tpif | cesopsn | cesopsn-cas}] [vc-
switching]
  - no cpipe service-id
  - description description-string
  - no description

```

```

- [no] endpoint endpoint-name
- [no] active-multipath
- description description-string
- no description
- revert-time revert-time
- no revert-time
- [no] standby-signaling-master
- [no] standby-signaling-slave
- [no] network-latency-measurement
- sap sap-id [create]
- [no] sap sap-id
- accounting-policy acct-policy-id
- no accounting-policy
- cem
- active-multipath-timeout seconds
- no active-multipath-timeout
- asym-delay-control [samples {ksamples}] [min-repeat minutes] [threshold-
repeat uSecs]
- no asym-delay-control
- [no] mute-output
- [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] entropy-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).

4.13.1.1.4 Epipe service configuration commands

```

config
- service

```

```

- epipe service-id [customer customer-id] [create] [vpn vpn-id] [vc-switching]
- no epipe service-id
  - bgp (see Epipe commands for EVPN)
  - bgp-evpn (see Epipe commands for EVPN)
  - 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
  - load-balancing
    - [no] l4-load-balancing
    - [no] per-service-hashing
    - [no] teid-load-balancing
  - 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
      - encapsulation atm-encap-type
      - ingress
        - traffic-desc traffic-desc-profile-id
        - no traffic-desc
      - oam
        - [no] alarm-cells
  - 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] cflowd
  - [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
        - test-pattern {all-zeros | all-ones} [crc-enable]
        - no test-pattern
    - 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]
    - [no] entropy-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 about configuring ETH-CFM on Epipes, see the 7705 SAR OAM and Diagnostics Guide.

4.13.1.1.5 Epipe security configuration commands

```

config
- service
- epipe service-id
- fw-bypass-policy {bypass-id | name}
- no fw-bypass-policy
- zone {zone-id | name} [create]
- no zone {zone-id | 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}
- log {log-id | name}
- no log
- name 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] sap sap-id
- [no] shutdown
- [no] spoke-sdp sdp-id:vc-id
- [no] shutdown

```

4.13.1.1.6 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
- 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
  - 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

```

4.13.1.1.7 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
    - 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

```

4.13.1.1.8 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
  - 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]
- [no] entropy-label
- ingress
- vc-label ingress-vc-label
- no vc-label [ingress-vc-label]
- [no] shutdown
- precedence [precedence-value | primary]
- no precedence

```

4.13.1.2 Show commands

```

show
- eth-cfm (see the 7705 SAR 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
    - macsec
    - network-latency-measurement
    - 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}
- sap-using [ingress | egress] shaper-group shaper-group-name
- 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

```

4.13.1.3 Clear commands

```

clear
- service
  - id service-id
    - arp
    - network-latency-measurement
    - 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

```

4.13.2 Command descriptions

- [VLL service configuration commands](#)
- [Show commands](#)
- [Clear commands](#)

4.13.2.1 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](#)
- [Epipe security configuration commands](#)

4.13.2.1.1 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>epipe>zone
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 (such as #, \$, or spaces), 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>epipe>zone
config>service>epipe>zone>sap
config>service>epipe>zone>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.

4.13.2.1.2 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**

cpipe *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-type** {**satop-e1** | **satop-t1** | **satop-e3** | **satop-t3** | **satop-serial** | **satop-tpif** | **cesopsn** | **cesopsn-cas**}] [**vc-switching**]

no cpipe *service-id*

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.

The **no** form of the command deletes the configuration for the specified service.

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 number as the service ID.

Values 1 to 2147483647

vc-type

defines the type of VC signaled to the peer. This optional parameter is included when the Cpipe service is created.

Values

- satop-e1: unstructured SAToP E1 circuit emulation service
- satop-t1: unstructured SAToP DS1 circuit emulation service
- satop-e3: unstructured SAToP E3 circuit emulation service
- satop-t3: unstructured SAToP DS3 circuit emulation service
- satop-serial: unstructured SAToP serial circuit emulation service
- satop-tpif: unstructured SAToP circuit emulation service for teleprotection interface
- cesopsn: basic structured $n \times 64$ kb/s circuit emulation service
- cesopsn-cas: structured $n \times 64$ 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

epipe *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-switching**]
no epipe *service-id*

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 Nokia 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**]

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)

active-multipath

Syntax

[no] **active-multipath**

Context

config>service>cpipe>endpoint

Description

This command enables the Cpipe endpoint to have multiple active paths. When the command is enabled, Cpipe traffic can be transmitted to and received from up to four active paths at the same time.

If symmetrical latency is required over the active paths, the asymmetric delay control (ADC) feature must be enabled using the [asym-delay-control](#) command.

The **no** form of this command disables the multiple active paths capability.

Default

no active-multipath

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 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 **show 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 this 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

network-latency-measurement**Syntax**

[no] **network-latency-measurement**

Context

config>service>cpipe

Description

This command enables network latency measurement on the Cpipe.

The **no** form of this command disables network latency measurement on the Cpipe.

Default

no network-latency-measurement

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

load-balancing**Syntax**

load-balancing

Context

config>service>epipe

Description

This command accesses the context to configure load balancing.

l4-load-balancing

Syntax

[no] **l4-load-balancing**

Context

config>service>epipe>load-balancing

Description

This command enables or disables Layer 4 load balancing for the Epipe service. 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.

Adding the Layer 4 source and destination port fields to the hashing algorithm generates a higher degree of randomness and a more even distribution of packets across the available LAG ports.

You can add additional fields to generate more randomness and more equal distribution of packets with the [teid-load-balancing](#) command.

Hashing based on the **l4-load-balancing** and **teid-load-balancing** commands and hashing based on the **per-service-hashing** command are mutually exclusive.

The **no** form of the command disables Layer 4 load balancing.

Default

no l4-load-balancing

per-service-hashing

Syntax

[no] **per-service-hashing**

Context

config>service>epipe>load-balancing

Description

This command enables or disables hashing based on the service ID. The result of the hashing calculation is used to determine the distribution of packets.

Hashing based on the **per-service-hashing** command and hashing based on the **l4-load-balancing** and **teid-load-balancing** commands are mutually exclusive.

The **no** form of the command disables per-service hashing.

Default

no per-service-hashing

teid-load-balancing**Syntax**

[no] **teid-load-balancing**

Context

config>service>epipe>load-balancing

Description

This command enables or disables TEID load balancing for the Epipe service. 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.

Hashing based on the **teid-load-balancing** and **l4-load-balancing** commands and hashing based on the **per-service-hashing** command are mutually exclusive.

The **no** form of the command disables TEID load balancing.

Default

no teid-load-balancing

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.

If a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, 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

The following table displays MTU values for specific VC types.

Table 43: Service MTU values

VC-type	Example of service MTU	Advertised MTU
Ethernet	1514	1500

VC-type	Example of service MTU	Advertised MTU
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

4.13.2.1.3 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. See the 7705 SAR 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 lpipe service (RFC 1332)
- MLPPP bundle
- LAG
- Ethernet SAPs supporting null, dot1q, and qinq for lpipe service
- FR DLCI
- HDLC
- cHDLC (supported only for lpipe 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 the following table. The range of values for the parameters follow the table.

Table 44: SAP ID configurations

Type	Syntax	Example
port-id	<i>slot/mda/port[.channel]</i>	1/1/5
SCADA bridge	<i>slot/mda/bridge-id.branch-id</i>	1/5/16.10
null	[<i>port-id</i> <i>bundle-id</i> <i>lag-id</i> <i>aps-id</i> <i>mw-link-id</i>]	<i>port-id</i> : 1/1/3 <i>bundle-id</i> : bundle-ppp-1/1.1 <i>lag-id</i> : lag-1 <i>aps-id</i> : aps-1 <i>mw-link-id</i> : mw-link-1
dot1q	[<i>port-id</i> <i>lag-id</i> <i>aps-id</i> <i>mw-link-id</i>]:qtag1	<i>port-id</i> :qtag1: 1/1/3:100 <i>lag-id</i> : lag-1:10 <i>aps-id</i> : aps-1 <i>mw-link-id</i> : mw-link-1
qinq	[<i>port-id</i> <i>lag-id</i>]:qtag1.qtag2	<i>port-id</i> :qtag1.qtag2: 1/1/3:100.30 <i>lag-id</i> : lag-1:10.10
atm	[<i>port-id</i> <i>aps-id</i>][: <i>vpi/vci</i> <i>vpi</i> <i>vpi1.vpi2</i>] ¹	<i>port-id</i> : 1/1/1 or 1/1/1.1 (for T1/E1 channelized ports) <i>aps-id</i> : aps-1 <i>vpi/vci</i> : 16/26 <i>vpi</i> : 16 <i>vpi1.vpi2</i> : 16.22
lag	<i>lag-id</i>	lag-2
frame	[<i>port-id</i> <i>aps-id</i>]: <i>dlci</i>	<i>port-id</i> : 1/1/1

Type	Syntax	Example
		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
lcr-port-id	lcr-group-id/port.[channel]	lcr-1/1.1
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:

- 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	<i>slot/mda/port[.channel]</i>
SCADA bridge	<i>slot/mda/bridge-id.branch-id</i> <i>bridge-id</i> 1 to 16 <i>branch-id</i> 1 to 32
bundle-id	<i>bundle-type-slot/mda.bundle-num</i> bundle keyword <i>type</i> ima, ppp <i>bundle-num</i> 1 to 32
aps-id	<i>aps-group-id[.channel]</i> aps keyword <i>group-id</i> 1 to 24
lcr-port-id	<i>lcr-group-id/port.[channel]</i> lcr keyword <i>group-id</i> 1 to 6
mw-link-id	<i>mw-link-id</i> <i>id</i> 1 to 24
lag-id	<i>lag-id</i> lag keyword <i>id</i> 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	<i>tunnel-id.[private public]:tag</i> tunnel keyword <i>id</i> 1 to 16 ("1" is the only valid value) <i>tag</i> 0 to 4094

port-id
specifies the physical port ID in the *slot/mda/port* 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 subport. 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. The following table describes the allowed values for the port and encapsulation types.

Table 45: 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. A 0 qtag1 value also accepts untagged packets on the dot1q port, and a * qtag1 value accepts any VLAN ID that is not specifically configured on the port. ¹

Port type	Encap-type	Allowed values	Comments
Ethernet	QinQ	*, 0 to 4094	The SAP is identified by the two 802.1Q tags on the port. A 0 qtag1 or qtag2 value also accepts untagged packets on the qinq port, and 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.

cflowd

Syntax
[no] cflowd

Context
config>service>epipe>sap

Description

This command enables cflowd to collect traffic flow samples through a SAP for analysis. When cflowd is enabled on an Epipe service SAP, traffic can be sampled and processed by the system's cflowd engine and exported to cflowd version 10 collectors with the **l2-ip** template enabled.

When cflowd is enabled at the SAP level, all packets forwarded by the interface are subject to analysis according to the cflowd configuration.

For Layer 2 services, only ingress sampling is supported.

Default
no cflowd

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 addresses 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.

4.13.2.1.4 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 the following table.

Table 46: 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 46: SAP ID preconfiguration for SAP aggregation groups](#). See [Table 44: SAP ID configurations](#) 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

4.13.2.1.5 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 MEF8 standard.

If the **encap-type** of the SAP port is not **cem**, this command is blocked.

active-multipath-timeout

Syntax

active-multipath-timeout *seconds*

no active-multipath-timeout

Context

config>service>cpipe>sap>cem

Description

This command configures the period, in seconds, to wait for the remaining paths to become available after the first path is available. It is independent of whether ADC is enabled or disabled.

Default

10 s

Parameters

seconds
specifies duration of the timer

Values 1 to 60

asym-delay-control

Syntax

asym-delay-control [**samples** {*ksamples*}] [**min-repeat** *minutes*] [**threshold-repeat** *uSecs*]
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. This option is not configurable for ADC if AMP is enabled.

The optional **threshold-repeat** keyword sets the latency value that will be used by ADC repeat and on-demand ADC to determine if octets need to be added or removed. if the difference between the calculated average latency and the expected latency is greater than the **threshold-repeat** value, octets are added or dropped as necessary.

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

uSecs

the threshold used to determine if an octet correction is made after an ADC repeat or on-demand ADC analysis, in microseconds

Default 250

Values 125 to 4000, in 125- μ s increments

mute-output

Syntax

[no] mute-output

Context

config>service>cpipe>sap>cem

Description

This command blocks the VCB broadcast from transmitting out of this SAP. The Cpipe must be **shutdown** before this command can be issued. The **no** form of this command unmutes the SAP.

Default

no mute-output

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 \geq 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 24: Unstructured payload defaults](#). For SAToP serial, see [Table 21: SAToP serial payload size minimums and defaults](#). For CESoPSN without CAS, see [Table 25: Default and minimum payload sizes for CESoPSN without CAS](#). For CESoPSN with CAS, see [Table 26: Default values for the payload size for T1 and E1 CESoPSN with CAS](#).

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 $n \times 16$ E1 frames or $n \times 24$ 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. 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 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 = $2 \times l \times (\text{number of timeslots})$, where $l = 1$ to 20
- for DCR with a timestamp frequency of 19.44 MHz, the payload size = $8 \times l \times (\text{number of timeslots})$, where $l = 1$ 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.

For SAToP serial circuit emulation services, the configurable range is 2 to 1496 octets, in multiples of 2. Serial rates of 4800 b/s and lower only support a payload size of 2 bytes.

For SAToP TPIF circuit emulation services, the configurable range is 64 to 1514 octets, in multiples of 32. The default is 256 octets, which corresponds to a packetization delay of 1 ms.

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

4.13.2.1.6 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) Ethernet adapter card, the scheduler mode must be set to 16-priority. On a third-generation (Gen-3) Ethernet adapter card, the scheduler mode is always 4-priority. For information on adapter card generations, see the "Evolution of Ethernet Adapter Cards, Modules, and Platforms" section in the 7705 SAR 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**.

If the specified SAP is a LAG SAP, *agg-rate* and *cir-rate* can be configured regardless of the scheduler mode setting on Gen-2 or Gen-3 hardware. If the active port is on a Gen-3 card or platform, *agg-rate* and *cir-rate* are applicable. If the active port is on a Gen-2 card or platform, *agg-rate* and *cir-rate* apply when the scheduler mode is set to 16-priority. For details on the behavior of a mix-and-match LAG SAP, see the "LAG Support on Third-Generation Ethernet Adapter Cards, Ports, and Platforms" and "Network LAG Traffic Management" sections in the 7705 SAR Interface Configuration Guide.

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.

The following table 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.

Table 47: 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

Port/SAP type	Match-qinq-dot1p setting ¹	Existing packet tags	P-bits used for match
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.

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.

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).

The following table 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 48: 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
Dot1q/ 0 ¹	n/a	Outer tag
QinQ/ X.Y ¹	False	Two outer tags ³
	True	Outer tag ³
QinQ/ X.* ¹	True or False	Outer tag

Egress port type/SAP type	qinq-mark-top-only state	Egress P-bits marked or re-marked
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.

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, which only support 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.

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

4.13.2.1.7 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 Shelf V2 or 7705 SAR-18 MAC address

Values xx:xx:xx:xx:xx:xx**rtp**

enables the RTP header on the Epipe SAP

4.13.2.1.8 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

4.13.2.1.9 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 or from a failed MC-LCR MDA to a protection MDA. 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. In an MC-LCR configuration, if the active MDA fails, an MC-LCR switchover occurs which then triggers a pseudowire switchover.

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 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 Lpipe spoke SDP has additional parameters. See [spoke-sdp](#) for the Apipe, Cpipe, or Lpipe 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 the 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.

entropy-label

Syntax

[no] entropy-label

Context

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

Description

This command enables or disables the use of entropy labels for spoke SDPs.

If **entropy-label** is enabled, the entropy label (EL) and entropy label indicator (ELI) are inserted in packets for which at least one LSP in the stack for the far end of the tunnel used by the service has advertised entropy label capability.

If the tunnel is an RSVP-TE type, **entropy-label** can also be controlled by disabling **entropy-label-capability** under the **config>router>rsvp** or **config>router>mpls>lsp** contexts at the far-end LER.

When the **entropy-label** and **entropy-label-capability** commands are both enabled, the entropy label value inserted at the iLER is always based on the service ID.

Default

no entropy-label

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 lpipe SAP. The CE address specified at one end of an lpipe 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 lpipe 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>hpipe>spoke-sdp

config>service>ipipe>spoke-sdp

Description

This command indicates whether the control word is used. The value of the control word is negotiated with the peer.

The control word is mandatory for Cpipe SAToP and CESoPSN services, Epipe MEF 8 services, and Fpipe services (in one-to-one mapping configurations) and cannot be changed.

The control word is optional for Apipe services, Epipe non-MEF 8 services, Hpipe services, and lpipe services, but must be enabled for Hpipe pseudowire services when transporting packets that are less than 64 bytes. If the control word is enabled for Epipe or lpipe services, it will be set to all zeros and ignored on egress.

For Cpipe services, Epipe MEF 8 services, Fpipe services, and Hpipe services, when 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 word is set to the length of the packet. Otherwise, the length field is set to 0. 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 for packets less than 64 bytes, 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

no control-word

4.13.2.1.10 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 μ s, and a value of 400 equals 40000 μ s, or 40 ms.

Values 1 to 400

Default 400

4.13.2.1.11 ATM commands

atm

Syntax

atm

Context

config>service>apipe>sap

config>service>epipe>sap

Description

This command enables access to the context to configure ATM-related attributes. This command can only be used when a context (for example, a channel or SAP) supports ATM functionality, such as:

- configuring ATM ports or ATM port-related functionality on T1/E1 ASAP Adapter cards or OC3/STM1 Adapter cards
- configuring ATM-related configuration for ATM-based SAPs on T1/E1 ASAP Adapter cards or on OC3/STM1 Adapter cards

If ATM functionality is not supported for a context, the command returns an error.

egress

Syntax

egress

Context

config>service>apipe>sap>atm

config>service>epipe>sap>atm

Description

This command enables access to the context to configure egress ATM traffic policies for the SAP.

encapsulation

Syntax

encapsulation *atm-encap-type*

Context

config>service>epipe>sap>atm

Description

This command specifies the data encapsulation for an ATM PVCC-delimited 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 is dropped.

Default

For VLL 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)

ingress

Syntax

ingress

Context

config>service>apipe>sap>atm

config>service>epipe>sap>atm

Description

This command enables access to the context to configure ingress ATM traffic policies for the SAP.

oam

Syntax

oam

Context

```
config>service>apipe>sap>atm
```

```
config>service>epipe>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 or Epipe service is operationally up, OAM cells are transported over the Apipe or Epipe and are transparent to the 7705 SAR. This capability is in accordance with ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance.

traffic-desc

Syntax

```
traffic-desc traffic-desc-profile-id
```

```
no traffic-desc
```

Context

```
config>service>apipe>sap>atm>egress
```

```
config>service>epipe>sap>atm>egress
```

```
config>service>apipe>sap>atm>ingress
```

```
config>service>epipe>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 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)

Values 1 to 1000

alarm-cells

Syntax

[no] alarm-cells

Context

config>service>apipe>sap>atm>oam

config>service>epipe>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.

Apipes on the 7705 SAR do not support PVCC terminations. They allow OAM cells to be transported transparently end-to-end.

When this command is enabled on Apipe SAPs or ATM Epipe SAPs, AIS cells are generated when the Apipe, Epipe, or corresponding SAP is operationally down.

The **no** form of the command disables alarm-cells functionality for the Apipe or Epipe. When alarm-cells functionality is disabled, AIS cells are not generated as a result of the Apipe, Epipe, or corresponding SAP going into the operationally down state.

Default

alarm-cells

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

4.13.2.1.12 Epipe security configuration commands

fw-bypass-policy

Syntax

fw-bypass-policy {*bypass-id* | *name*}

no fw-bypass-policy

Context

config>service>epipe

Description

This command associates a bypass policy with this Epipe instance. The bypass policy must already be configured in the **config>security>bypass** context. All protocols, or protocols associated with specific source or destination ports, defined by the bypass policy bypass the firewall lookup table and are permitted across the zone associated with this Epipe instance without being firewalled.

Default

no fw-bypass-policy

Parameters

bypass-id

the firewall bypass ID number

Values 1 to 65535

name

the name of the firewall bypass policy

Values 1 to 32 characters

zone

Syntax

zone {*zone-id* | *name*} [**create**]

no zone {*zone-id* | *name*}

Context

config>service>epipe

Description

This command creates a security zone within an Epipe context. Each zone must have a unique ID. When a zone is created with a name, the system automatically assigns it the first available zone ID value.

The **no** form of this command deletes the zone. When a zone is deleted, all configuration parameters for the zone are also deleted.

Default

no zone

Parameters

zone-id

the zone ID number.

Values 1 to 65534

name

the name of the zone.

Values 1 to 32 characters (must start with a letter). If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

create

keyword required when first creating the security zone. When the zone is created, you can enter the context without the **create** keyword.

abort

Syntax

abort

Context

config>service>epipe>zone

Description

This command discards changes made to a security feature.

Default

n/a

begin

Syntax

begin

Context

config>service>epipe>zone

Description

This command enters the mode to create or edit security features.

Default

n/a

commit

Syntax

commit

Context

config>service>epipe>zone

Description

This command saves changes made to security features.

Default

n/a

inbound**Syntax**

inbound

Context

config>service>epipe>zone

Description

This command enables the context to configure limit parameters for inbound firewall sessions.

Default

n/a

outbound**Syntax**

outbound

Context

config>service>epipe>zone

Description

This command enables the context to configure limit parameters for outbound firewall sessions.

Default

n/a

limit**Syntax**

limit

Context

config>service>epipe>zone>inbound

config>service>epipe>zone>outbound

Description

This command enables the context to configure limits on concurrent sessions for inbound or outbound firewall sessions.

Default

n/a

concurrent-sessions

Syntax

concurrent-sessions {tcp | udp | icmp | other} *sessions*
no concurrent-sessions {tcp | udp | icmp | other}

Context

config>service>epipe>zone>inbound>limit
config>service>epipe>zone>outbound>limit

Description

This command configures the maximum number of concurrent firewall sessions that can be established per zone, for the specified protocol, 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 specified direction and protocol

Values 1 to 16383

log

Syntax

log {*log-id* | *name*}

no log

Context

config>service>epipe>zone

Description

This command applies a security log to the specified zone. The security log must already be configured in the **config>security>logging** context.

The **no** form of this command removes logging for the zone.

Default

n/a

Parameters

log-id

the identifier for the log

Values 1 to 100

name

the name of the log

Values 1 to 32 characters

name

Syntax

name *name*

no name

Context

config>service>epipe>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. If the zone name was already configured with the **zone** command, this command renames the zone.

Default

n/a

Parameters

name

the name of the zone

Values 1 to 32 characters (must start with a letter). If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

policy

Syntax

policy {*policy-id* | *name*}

no policy

Context

config>service>epipe>zone

Description

This command specifies the policy to be used by the security zone to build its matching criteria for incoming packets. The policy must already be configured in the **config>security** context.

The **no** form of this command deletes the specified policy.

Default

n/a

Parameters

policy-id

the number of the referenced policy

Values 1 to 65535

name

the name of the referenced policy

sap

Syntax

[no] **sap** *sap-id*

Context

config>service>epipe>zone

Description

This command assigns a SAP to the security zone.
The **no** form of this command removes the SAP from the zone.

Default

n/a

Parameters

sap-id
specifies the physical port identifier portion of the SAP definition. See [Table 44: SAP ID configurations](#) for a full list of SAP IDs.

spoke-sdp

Syntax

[no] spoke-sdp sdp-id:vc-id

Context

config>service>epipe>zone

Description

This command assigns a spoke SDP to the security zone.
The **no** form of this command removes the spoke SDP from the zone.

Default

n/a

Parameters

sdp-id
uniquely identifies the SDP
Values 1 to 17407

vc-id
identifies the virtual circuit
Values 1 to 4294967295

4.13.2.2 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 [Output example \(Apipe ATMVcc service\)](#) is an example of **id service-id all** information, and [Table 49: Service-ID all command field descriptions](#) describes the fields. Following the table are output examples for:

- [Output example \(Apipe ATMVpc service\)](#)
- [Output example \(Cpipe service\)](#)
- [Output example \(Epipe service\)](#)
- [Output example \(TDM SAP-to-Ethernet SAP MEF 8 Epipe service\)](#)
- [Output example \(Fpipe service\)](#)
- [Output example \(Hpipe service\)](#)
- [Output example \(Ipipe service\)](#)
- [Output example \(Ipipe frame relay to Ethernet service\)](#)
- [Output example \(Ipipe cHDLC to Ethernet service\)](#)

Output example (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  -(10.120.38.1)
-----
SDP Id           : 2:2
VC Type          : ATMVCC
Admin Path MTU   : 0
Far End          : 10.120.38.1
Type            : Spoke
VC Tag          : 0
Oper Path MTU    : 0
Delivery        : MPLS

Admin State      : Up
Acct. Pol       : None
Ingress Label    : 0
Ing mac Fltr    : n/a
Ing ip Fltr     : n/a
Admin ControlWord : Not Preferred
Admin BW(Kbps)  : 0
Last Status Change : 03/11/2008 19:58:19
Last Mgmt Change  : 03/28/2008 19:49:51
Endpoint        : N/A
PW Status Sig   : Enabled
Class Fwding State : Down
Flags           : StandbySigSlaveTxDown
Mac Move        : Ukwn
Local Pw Bits    : None
Peer Pw Bits     : pwFwdingStandby
Peer Fault Ip    : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None

Oper State       : Down
Collect Stats    : Disabled
Egress Label     : 0
Egr mac Fltr    : n/a
Egr ip Fltr     : n/a
Oper ControlWord : False
Oper BW(Kbps)    : 0
Signaling        : TLDP
Precedence       : 4

Blockable Level  : Unknown

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

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

Associated LSP LIST :
No LSPs Associated

-----
APIPE Service Destination Point specifics
-----
Admin Concat Limit : 1
Peer Concat Limit  : n/a
Oper Concat Limit  : 1
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 (Ingress)
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

Forwarding Engine Stats (Egress)
Dropped            : 0           n/a

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
Ingress TD Ovr    : N/A                    Egress TD Ovr    : N/A
Alarm Cell Handling: Enabled
OAM Termination   : Disabled                Periodic Loopback : Disabled
AAL-5 Encap       : n/a
-----
Service Endpoints
-----
No Endpoints found.
=====

```

Table 49: Service-ID all command field descriptions

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
MTU	Specifies the service MTU
Vc Switching	Specifies whether the service is configured as a PW switching point
NW Lat Measurement	Specifies whether network latency measurement is enabled or disabled on the Cpipe
SAP Count	Displays the number of SAPs specified for this service
SDP Bind Count	Displays the number of SDPs bound to this service

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)
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
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)

Label	Description
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 PathMTUtoo Small, 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, lacIngressFault, lacEgressFault, PwFwding Standby, and so on.
Peer Fault Ip	N/A
Peer Vccv CV Bits	Displays the setting of the pseudowire peer VCCV control verification bits (IsrPing)
Peer Vccv CC Bits	Displays the setting of the pseudowire peer VCCV control channel bits (pwe3ControlWord and/or mplsRouterAlert Label)
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
Hello Msg Len	Specifies the length of the SDP Echo Request messages transmitted on this SDP

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
Statistics/Grp Enc Stats	
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
E. Dro. Pkts.	Specifies the number of dropped egress packets
RSVP/Static LSPs:	
Associated LSP List	Lists the associated LSPs
Lsp Name	Specifies the name of the LSP
Time since Last Trans*	Specifies the time that the associated LSP has been in service
Segment Routing	
ISIS	Indicates the state of IS-IS segment routing: enabled or disabled
OSPF	Indicates the state of OSPF segment routing: enabled or disabled
TE-LSP	Indicates the state of TE-LSP segment routing: enabled or disabled
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)

Label	Description
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
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)

Label	Description
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 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
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

Label	Description
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
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 ActMultiPath	Specifies whether active multipath is enabled on the local SDP
Peer ActMultiPath	Specifies whether active multipath 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

Label	Description
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
IpipSdpBindCelpAd*	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
Description	Specifies the description for this SAP
Admin State	Specifies the desired state of this SAP
Oper State	Specifies the operating state of this SAP
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
Hold Meps Up	Specifies if hold-mep-up-on-failure is enabled or disabled for a SAP Epipe

Label	Description
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
Ingr IPv6 Fltr-Id	Specifies the ingress IPv6 filter policy ID assigned to the SAP
Egr IPv6 Fltr-Id	Specifies the egress IPv6 filter policy ID assigned to the SAP
tod-suite	n/a
qinq-pbit-marking	Indicates the qinq P-bit marking for the SAP: both or top
Ing Scheduler Mode	Indicates the ingress scheduler mode for the SAP
Egr Scheduler Mode	Indicates the egress scheduler mode for the SAP
Ing Agg Rate Limit	Indicates the PIR rate limit in the access ingress direction for the aggregate of the SAP queues
Egr Agg Rate Limit	Indicates the PIR rate limit in the access egress direction for the aggregate of the SAP queues
Ing Agg cir	Indicates the CIR rate limit in the access ingress direction for the aggregate of the SAP queues
Egr Agg cir	Indicates the CIR rate limit in the access egress direction for the aggregate of the SAP queues
Ing Shaper Group	Indicates the ingress shaper group for the SAP
Egr Shaper Group	Indicates the egress shaper group for the SAP

Label	Description
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
tod-suite	Indicates whether a time-based policy is applied to a multiservice site. Not applicable.
Ingr 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	

Label	Description
Last Cleared Time	Displays the date and time that a clear command was issued on statistics
Forwarding Engine Stats (Ingress)	
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
Forwarding Engine Stats (Egress)	
Dropped	Indicates the number of packets or octets dropped by the forwarding engine Note: This number is doubled when active-multipath is enabled.
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

Label	Description
Sap per Queue stats	
Ingress Queue n (Priority)	Specifies the index of the ingress QoS queue of this SAP, where n is the index number
Off. Combined	Indicates the combined total number of high-priority and low-priority packets or octets offered to the forwarding engine
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
Ingress Queue n (Profile)	Specifies the index of the ingress QoS queue of this SAP, where n is the index number
Off. ColorIn	Indicates the number of packets or octets colored as in-profile for the SAP (offered)
Off. ColorOut	Indicates the number of packets or octets colored as out-of-profile for the SAP (offered)
Off. Uncolor	Indicates the number of packets or octets that are unprofiled for the SAP (offered)
Dro. ColorOut	Indicates the number of packets or octets colored as out-of-profile that were dropped for the SAP
Dro. ColorIn/Uncolor	Indicates the number of packets or octets that were colored as in-profile or unprofiled that were dropped for the SAP
For. InProf	Indicates the number of forwarded packets or octets colored as in-profile (FC profile set to "in" or "no profile" and rate less than or equal to CIR)

Label	Description
For. OutProf	Indicates the number of forwarded packets or octets that were colored as out-of-profile (FC profile set to "out" or "no profile" and rate above CIR)
Egress Queue <i>n</i>	Specifies the index of the egress QoS queue of the SAP, where <i>n</i> 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
Ingress TD Ovr	The Ovr for the traffic descriptor applied to the ingress SAP
Egress TD Ovr	The Ovr for 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	The type of AAL5 encapsulation for this ATM SAP
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

Label	Description
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 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	Displays the number of times that the asymmetric delay control analysis is repeated
Samples	Displays the number of packets that are analyzed (k = 1024)
AMP timeout (s)	Displays the active-multipath-timeout value in seconds
AMP State	Displays the active multipath state (all-paths-active, not-all-paths-active, initialization, or down)
AMP Num Of Active Path	Displays the number of active multipath paths
AMP Active Sdp Id	Displays the active spoke SDP Note: This field is displayed only if the active multipath state is not-all-paths-active.
Cfg Alarm	Displays the alarms that have alarm reporting enabled
Alarm Status	Displays 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

Label	Description
Time (us)	(Column heading) Displays time (in microseconds) since the last counter reset for the statistic
Time (sec)	(Column heading) Displays the time (in seconds) since the last counter reset for the statistic
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
Jitter Buffer Depth	Displays the number of packets sitting in the jitter buffer at that instant for the Cpipe or MEF 8 Epipe
Jitter Buffer Delay	Displays the total amount of TDM PW data buffered
AMP Duplicate Dropped	Displays the number of AMP duplicate packets that have been dropped

Label	Description
AMP State Change Count	Displays the number of times the active multipath state has changed
ADC JB Sampling Complete	Displays the number of asymmetric delay control analysis periods completed
ADC JB Adjust	Displays the number of jitter buffer adjustments made for asymmetric delay control
ADC JB Sampling Avg Delay	Displays the average jitter buffer delay value of the samples taken during the last ADC analysis period
ADC JB Delay – target value	Displays the target jitter buffer delay value
Ingress Stats	Indicates that the following statistics are ingress statistics
Forwarded	Displays the number of forwarded packets
Dropped	Displays the number of dropped packets
Service Endpoints	
Endpoint name	Identifies the endpoint
Description	Describes the endpoint
Creation Origin	Indicates whether the service creation was manual or automatic
Revert time	Displays the revert time setting for the active spoke SDP
Act Hold Delay	Not applicable
Standby Signaling Master	Indicates whether the standby-signaling master is enabled (true or false)
Standby Signaling Slave	Indicates whether the standby-signaling slave is enabled (true or false)
Active-Multipath	Indicates whether active multipath is enabled (true or false)
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)
Oper Status	Indicates the operational status of the primary and secondary spoke SDPs

Output example (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  -(10.120.20.1)
-----
SDP Id           : 5:5              Type           : Spoke
VC Type          : ATMVPC           VC Tag         : 0
Admin Path MTU   : 0               Oper Path MTU  : 0
Far End          : 10.120.20.1      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
Last Mgmt Change : 04/01/2008 16:51:59
Signaling        : TLDP
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
Forwarding Engine Stats (Ingress)
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
-----

```

Forwarding Engine Stats (Egress)

```

Dropped      : 0                      n/a
-----

```

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
Ingress TD Ovr	: N/A	Egress TD Ovr	: N/A
Alarm Cell Handling	: Enabled		
OAM Termination	: Disabled	Periodic Loopback	: Disabled
AAL-5 Encap	: n/a		

Service Endpoints

No Endpoints found.

*A:ALU-A>show>service#

Output example (Cpipe service)

*A:7705:Dut-H# show service id 1 all

Service Detailed Information

Service Id	: 1		
Service Type	: Cpipe	VLL Type	: CESoPSN
Name	: (Not Specified)		
Description	: Default Cpipe description for service id 1		
Customer Id	: 1	Creation Origin	: manual
Last Status Change	: 01/20/2022 17:29:32		
Last Mgmt Change	: 01/20/2022 17:28:56		
Admin State	: Up	Oper State	: Up
MTU	: 1400		
Vc Switching	: False		
NW Lat Measurement	: enabled		
SAP Count	: 1	SDP Bind Count	: 2

Service Destination Points(SDPs)

Sdp Id 122:101 -(108.108.108.108)

Description	: (Not Specified)		
SDP Id	: 122:101	Type	: Spoke
Split Horiz Grp	: (Not Specified)		
VC Type	: CESoPSN	VC Tag	: 0
Admin Path MTU	: 0	Oper Path MTU	: 1550
Delivery	: MPLS		

```

Far End                : 108.108.108.108
Tunnel Far End         : n/a
Entropy Label          : Disabled
LSP Types              : RSVP

Admin State            : Up
MinReqd SdpOperMTU    : 1404
Acct. Pol              : None
Ingress Label          : 131068
Ingr Mac Fltr-Id       : n/a
Ingr IP Fltr-Id        : n/a
Admin ControlWord      : Preferred
Admin BW(Kbps)         : 0
Last Status Change     : 01/20/2022 17:29:23
Last Mgmt Change       : 01/20/2022 17:28:55
Endpoint               : ep1
PW Status Sig          : Enabled
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      : pwe3ControlWord mplsRouterAlertLabel
Ingress Qos Policy     : (none)
Ingress FP QGrp        : (none)
Ing FP QGrp Inst       : (none)
Egress Qos Policy      : (none)
Egress Port QGrp       : (none)
Egr Port QGrp Inst     : (none)
KeepAlive Information  :
Admin State            : Disabled
Hello Time             : 10
Max Drop Count         : 3
Statistics             :
I. Fwd. Pkts.          : 146156
I. Fwd. Octs.          : 8769360
E. Fwd. Pkts.          : 146157
E. Dro. Pkts.          : 0
Grp Enc Stats          :
I. Fwd. Pkts.          : 0
I. Dro. Inv. Spi.      : 0
E. Fwd. Pkts.          : 0
E. Dro. Enc. Pkts.     : 0
Oper State             : Up
Collect Stats          : Disabled
Egress Label           : 131068
Egr Mac Fltr-Id        : n/a
Egr IP Fltr-Id         : n/a
Oper ControlWord       : True
Oper BW(Kbps)          : 0
Signaling              : TLDP
Precedence             : 4
Oper State             : Disabled
Hello Msg Len          : 0
Hold Down Time         : 10
I. Dro. Pkts.          : 0
I. Dro. Octs.          : n/a
E. Fwd. Octets         : 8769420
I. Fwd. Octs.          : 0
I. Dro. OthEncPkt*     : 0
E. Fwd. Octs.          : 0

-----
RSVP/Static LSPs
-----
Associated LSP List :
Lsp Name            : Path2Strict
Admin State         : Up
Time Since Last Tr* : 00h03m11s
Oper State          : Up

-----
Segment Routing
-----
ISIS                : disabled
OSPF                 : disabled
TE-LSP              : disabled

-----
CPIPE Service Destination Point specifics
-----
Local Bit-rate       : 1
Local Payload Size   : 8
Local Jitter Buffer   : 8
Local Asym Delay     : enabled
Local ActMultiPath   : enabled
Peer Bit-rate        : 1
Peer Payload Size    : 8
Peer Jitter Buffer    : 8
Peer Asym Delay      : enabled
Peer ActMultiPath    : enabled

```

```

Local Sig Pkts      : No Sig.
Local CAS Framing   : No CAS
Local RTP Header    : No
Local Differential   : No
Local Timestamp     : 0
Peer Sig Pkts      : No Sig.
Peer CAS Framing    : No CAS
Peer RTP Header     : No
Peer Differential    : No
Peer Timestamp      : 0
-----
Sdp Id 123:1  -(108.108.108.108)
-----
Description         : (Not Specified)
SDP Id              : 123:1
Split Horiz Grp     : (Not Specified)
VC Type             : CESoPSN
Admin Path MTU      : 0
Delivery            : MPLS
Far End             : 108.108.108.108
Tunnel Far End      : n/a
Entropy Label       : Disabled
Type                : Spoke
VC Tag              : 0
Oper Path MTU       : 1550
LSP Types           : RSVP
Admin State         : Up
MinReqd SdpOperMTU : 1404
Acct. Pol           : None
Ingress Label       : 131069
Ingr Mac Fltr-Id    : n/a
Ingr IP Fltr-Id     : n/a
Admin ControlWord   : Preferred
Admin BW(Kbps)      : 0
Last Status Change  : 01/20/2022 17:29:29
Last Mgmt Change    : 01/20/2022 17:28:56
Endpoint            : ep1
PW Status Sig       : Enabled
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   : pwe3ControlWord mplsRouterAlertLabel
Ingress Qos Policy  : (none)
Ingress FP QGrp     : (none)
Ing FP QGrp Inst    : (none)
Egress Qos Policy   : (none)
Egress Port QGrp    : (none)
Egr Port QGrp Inst  : (none)
KeepAlive Information :
Admin State         : Disabled
Hello Time          : 10
Max Drop Count      : 3
Statistics          :
I. Fwd. Pkts.       : 139321
I. Fwd. Octs.       : 8359260
E. Fwd. Pkts.       : 139319
E. Dro. Pkts.       : 0
Grp Enc Stats       :
I. Fwd. Pkts.       : 0
I. Dro. Inv. Spi.   : 0
E. Fwd. Pkts.       : 0
E. Dro. Enc. Pkts.  : 0
Oper State          : Up
Collect Stats       : Disabled
Egress Label        : 131069
Egr Mac Fltr-Id     : n/a
Egr IP Fltr-Id      : n/a
Oper ControlWord    : True
Oper BW(Kbps)       : 0
Signaling           : TLDP
Precedence          : 4
Oper State          : Disabled
Hello Msg Len       : 0
Hold Down Time      : 10
I. Dro. Pkts.       : 0
I. Dro. Octs.       : n/a
E. Fwd. Octets      : 8359140
I. Fwd. Octs.       : 0
I. Dro. OthEncPkt*  : 0
E. Fwd. Octs.       : 0
-----
RSVP/Static LSPs
-----
Associated LSP List :
Lsp Name            : Path1Strict
Admin State         : Up
Time Since Last Tr* : 00h03m04s
Oper State          : Up
-----

```

Segment Routing

```

-----
ISIS                : disabled
OSPF                 : disabled
TE-LSP               : disabled
-----

```

CPIPE Service Destination Point specifics

```

-----
Local Bit-rate      : 1                Peer Bit-rate      : 1
Local Payload Size  : 8                Peer Payload Size  : 8
Local Jitter Buffer  : 8                Peer Jitter Buffer  : 8
Local Asym Delay    : enabled           Peer Asym Delay    : enabled
Local ActMultiPath  : enabled           Peer ActMultiPath  : enabled
Local Sig Pkts      : No Sig.           Peer Sig Pkts      : No Sig.
Local CAS Framing    : No CAS            Peer CAS Framing    : No CAS
Local RTP Header     : No               Peer RTP Header     : No
Local Differential   : No               Peer Differential   : No
Local Timestamp      : 0                Peer Timestamp      : 0
-----

```

Number of SDPs : 2

* indicates that the corresponding row element may have been truncated.

Service Access Points

SAP 1/4/15.1

```

-----
Service Id          : 1
SAP                  : 1/4/15.1          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   : 01/20/2022 17:29:32
Last Mgmt Change     : 01/20/2022 17:28:55
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
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
Ing Shaper Group     : n/a              Egr Shaper Group     : n/a
Endpoint             : N/A
Acct. Pol            : None              Collect Stats         : Disabled
-----

```

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    : 01/20/2022 17:30:06
                      Packets              Octets
Forwarding Engine Stats (Ingress)
-----

```

```

Dropped                : 0                0
Off. HiPrio            : 143160           1145280
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            : 143160           1145280
For. OutProf           : 0                0
Forwarding Engine Stats (Egress)
Dropped                : 0                n/a
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            : 143160           1145280
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           : 8                  Jitter Buffer (ms) : 8
Jitter Buffer (packets) : 8                  Playout Threshold (packets): 5
Use RTP Header         : No                 Differential       : No
Timestamp Freq         : 0                  CAS Framing        : No CAS
Effective PDVT         : +/-4.0 ms

AsymDelayControl       : enabled           ThresholdRepeat (usec) : 250
RepeatPeriod           : disabled          Samples            : 2K
AMP timeout(s)         : 60
AMP State              : all-paths-active   AMP Num Of Active Path : 2
Cfg Alarm              : stray malformed pktloss overrun underrun
Alarm Status           :

-----
CEM SAP Statistics
-----
Packets    Time(us)    Time(sec)    Events
Egress Stats
Forwarded   : 147643
Dropped     : 0
Missing     : 0
Reordered Forwarded : 0
Underrun    : 0                0
Overrun     : 0                0
Misordered Dropped : 0
Malformed Dropped : 0
LBit Dropped : 0
Error       :                0
Severely Error :                0

```



```

Unavailable          : 0
Failure Count        : 0
Jitter Buffer Depth  : 4
Jitter Buffer Delay   : 5000
AMP Duplicate Dropped : 147643
AMP State Change Count : 0
ADC JB Sampling Complete : 0
ADC JB Adjust        : 0
ADC JB Sampling Avg Delay : 4243
ADC JB Target Delay   : 4000
Ingress Stats
Forwarded            : 147643
Dropped              : 0
-----
Service Endpoints
-----
Endpoint name        : ep1
Description           : (Not Specified)
Creation Origin       : manual
Revert time          : 0
Act Hold Delay        : 0
Standby Signaling Master : false
Standby Signaling Slave : false
Active-Multipath      : true
-----
Members
-----
Spoke-sdp: 122:101 Prec:4      Oper Status: Up
Spoke-sdp: 123:1 Prec:4       Oper Status: Up
=====
*A:7705:Dut-H#

```

Output example (Epipe service)

```

A:ALU-1>show>service>id# all
=====
Service Detailed Information
=====
Service Id          : 2
Service Type        : Epipe
Name                : epipel
Description          : (Not Specified)
Customer Id         : 1
Creation Origin      : manual
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
TEID Hashing        : Disabled
L4 Hashing          : Disabled
Force QTag Fwd      : n/a
-----
BGP Information
-----
Service Destination Points(SDPs)
-----
Sdp Id 1:11  -(10.10.10.10)
-----
SDP Id             : 1:11
Type               : Spoke

```

```

VC Type           : Ether
Admin Path MTU    : 0
Far End           : 10.10.10.10
VC Tag            : n/a
Oper Path MTU     : 0
Delivery          : MPLS

Admin State       : Up
Acct. Pol         : None
Ingress Label     : 0
Ing mac Fltr      : n/a
Ing ip Fltr       : n/a
Admin ControlWord : Not Preferred
Admin BW(Kbps)    : 0
Last Status Change : 07/13/2009 18:50:40
Last Mgmt Change  : 07/13/2009 18:50:40
Endpoint          : N/A
PW Status Sig     : Enabled
Class Fwding State : Down
Flags             : SvcAdminDown SdpOperDown
                  : NoIngVCLabel NoEgrVCLabel
                  : PathMTUTooSmall
Time to RetryReset : 476014240 seconds
Mac Move          : Ukwn
Local Pw Bits     : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None

VC Tag            : n/a
Oper State        : Down
Collect Stats     : Disabled
Egress Label      : 0
Egr mac Fltr      : n/a
Egr ip Fltr       : n/a
Oper ControlWord  : False
Oper BW(Kbps)     : 0
Signaling         : TLDP
Force Vlan-Vc     : Disabled
Precedence        : 4

Retries Left      : -1
Blockable Level   : Unknown

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

Statistics :
I. Fwd. Pkts.    : 0
I. Fwd. Octs.    : 0
E. Fwd. Pkts.    : 0
I. Dro. Pkts.    : 0
I. Dro. Octs.    : 0
E. Fwd. Octets   : 0
-----
Eth-Cfm Configuration Information
-----
Md-index         : 1
Ma-index         : 1
MepId            : 2
LowestDefectPri  : macRemErrXcon
Defect Flags     : None
Mac Address      : a4:58:ff:00:00:00
CcmTx            : 0
LbRxReply        : 0
LbRxBadMsdu     : 0
LbNextSequence  : 1
LtRxUnexplained  : 0
Direction        : Down
Admin            : Disabled
CCM-Enable       : Disabled
HighestDefect    : none
CcmLtmPriority   : 7
CcmSequenceErr   : 0
LbRxBadOrder     : 0
LbTxReply        : 0
LtNextSequence   : 1

Associated LSP LIST :
No LSPs Associated

-----
Number of SDPs : 1
-----
Service Access Points
-----
-----
SAP 1/5/1
-----

```

```

Service Id      : 2
SAP             : 1/5/1
Admin State     : Up
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
Split Horizon Group: (Not Specified)

QinQ Ethertype  : 0x8100

Hold Meps Up    : Disabled
LLF Admin State : Down
Admin MTU        : 1514
Ingr IP Fltr-Id : n/a
Ingr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a
tod-suite        : None
Ing Scheduler Mode : 4-priority
Ing Agg Rate Limit : 999000
Ing Agg cir       : 333000
Ing Shaper Group  : test_sgl
Endpoint         : N/A
Q Frame-Based Acct : Disabled
Vlan-translation : None

LLF Oper State   : Clear
Oper MTU         : 1514
Egr IP Fltr-Id   : n/a
Egr Mac Fltr-Id   : n/a
Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both
Egr Scheduler Mode : 4-priority
Egr Agg Rate Limit : max
Egr Agg cir       : 0
Egr Shaper Group  : default

Acct. Pol        : None
Collect Stats    : Disabled

Ignore Oper Down : Disabled

Loopback         : None
Swap Mac Addr    : Disabled
Loopback Time Left: unspecified
-----
QoS
-----
Ingress qos-policy : 1
Ingress FP QGrp    : n/a
Ing FP QGrp Inst    : n/a
Shared Q plcy       : n/a
Egress qos-policy  : 1
Egress Port QGrp    : n/a
Egr Port QGrp Inst : n/a
Multipoint shared   : Disabled
-----
Sap Statistics
-----
Last Cleared Time   : N/A
                    Packets
Forwarding Engine Stats (Ingress)
Dropped             : 0
Off. HiPrio         : 0
Off. LowPrio        : 0
                    Octets

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

Forwarding Engine Stats (Egress)
Dropped             : 0
                    n/a

Queueing Stats(Egress QoS Policy 1)
Dro. InProf         : 0
Dro. OutProf        : 0
For. InProf         : 0
For. OutProf        : 0

```

```

-----
Sap per Queue stats
-----
                Packets                Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio      : 0                    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#

```

Output example (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 (Ingress)		
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

Forwarding Engine Stats (Egress)

Dropped	0	n/a
---------	---	-----

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
Off. LoPrio      : 5891091
Dro. HiPrio      : 0
Dro. LoPrio      : 0
For. InProf      : 0
For. OutProf     : 5891091
0
907228014

Egress Queue 1
For. InProf      : 5891109
For. OutProf     : 0
Dro. InProf      : 0
Dro. OutProf     : 0
0
930795222

-----
SAP 1/1/1.1
-----
Service Id       : 1
SAP              : 1/1/1.1
Description      : Default sap description for service id 1
Admin 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
Ingr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a
tod-suite        : None
Ing Scheduler Mode : 4-priority

Oper MTU         : 1514
Egr IP Fltr-Id   : n/a
Egr Mac Fltr-Id  : n/a
Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : n/a
Egr Scheduler Mode: 4-priority

Endpoint         : N/A
Vlan-translation : None

Acct. Pol        : None
Application Profile: None
Collect Stats    : Disabled

-----
QOS
-----
Ingress qos-policy : 1
Shared Q plcy      : n/a
Egress qos-policy  : 1
Multipoint shared  : Disabled

-----
Sap Statistics
-----
Last Cleared Time : N/A

Packets          Octets
Forwarding Engine Stats (Ingress)
Dropped          : 0
Off. HiPrio      : 5886512
Off. LowPrio     : n/a
0
753473536
n/a

```

```

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio      : 0
Dro. LowPrio     : n/a
For. InProf      : 5886512
For. OutProf     : 0

Forwarding Engine Stats (Egress)
Dropped          : 0

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

Ingress Queue 1 (Priority)
Off. HiPrio      : 5886512
Off. LoPrio      : n/a
Dro. HiPrio      : 0
Dro. LoPrio      : n/a
For. InProf      : 5886512
For. OutProf     : 0

Egress Queue 1
For. InProf      : n/a
For. OutProf     : n/a
Dro. InProf      : n/a
Dro. OutProf     : n/a
-----
CEM SAP Configuration Information
-----
Endpoint Type    : NxDS0
Payload Size     : 128
Jitter Buffer (packets): 5
Use RTP Header   : No
Timestamp Freq   : 0
Effective PDVT   : +/-2.0 ms
Local Ecid       : 2
Remote MacAddr   : 10:00:50:00:00:02
Bit-rate         : 16
Jitter Buffer (ms) : 5
Playout Threshold (packets): 3
Differential     : No
CAS Framing      : No CAS
Remote Ecid      : 2

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
Overrun          : 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#

```

Output example (Fpipe service)

```

=====
*A:ALU-A>show>service# id 10021 all
=====
Service Detailed Information
=====
Service Id       : 10021
Service Type     : Fpipe
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
MTU              : 1508
Vc Switching     : False
SAP Count        : 1
Oper State       : up
SDP Bind Count   : 1
-----
Service Destination Points(SDPs)
-----
Sdp Id 1:1021 - (100.100.100.148)
-----
Description      : (Not Specified)
SDP Id           : 1:1021
Split Horiz Grp  : (Not Specified)
VC Type          : FRDLICI
Admin Path MTU   : 0
Far End          : 100.100.100.148
Admin State      : Up
Acct. Pol        : None
Ingress Label    : 131056
Ing mac Fltr     : n/a
Ing ip Fltr      : n/a
Ing ipv6 Fltr    : n/a
Admin ControlWord : Preferred
Admin BW(Kbps)   : 0
Last Status Change: 07/07/2011 17:04:33
Last Mgmt Change : 07/07/2011 17:03:06
Endpoint         : N/A
Class Fwding State: Down
Flags            : None
Mac Move         : Ukwn
Local Pw Bits    : None
Peer Pw Bits     : None
Peer Fault Ip    : None
Peer Vccv CV Bits : lspPing
Type             : Spoke
VC Tag           : 0
Oper Path MTU    : 1550
Delivery         : LDP
Oper State       : Up
Collect Stats    : Disabled
Egress Label     : 131056
Egr mac Fltr     : n/a
Egr ip Fltr      : n/a
Egr ipv6 Fltr    : n/a
Oper ControlWord : True
Oper BW(Kbps)    : 0
Signaling        : TLDP
Precedence       : 4
Blockable Level  : Unknown

```



```

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
Forwarding Engine Stats (Ingress)
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

Forwarding Engine Stats (Egress)
Dropped           : 0                n/a

```

```

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.
=====

```

Output example (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
Last Status Change: 07/11/2011 14:16:35
Last Mgmt Change   : 07/11/2011 14:16:14
Endpoint           : N/A
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
Hello Time         : 10
Drop Count         : 3
Statistics :
I. Fwd. Pkts.      : 0
I. Fwd. Octs.      : 0
E. Fwd. Pkts.      : 0
Associated LSP LIST :
Lsp Name           : LSPToDut-B
Admin State        : Up
Time Since Last Tr*: 00h01m53s
Oper BW(Kbps)      : 0
Signaling          : TLDP
Precedence         : 4
Oper State         : Disabled
Hello Msg Len      : 0 Max
Hold Down Time     : 10
I. Dro. Pkts.      : 0
I. Dro. Octs.      : 0
E. Fwd. Octets     : 0
Oper State         : Up
-----
Number of SDPs : 1
-----
Service Access Points
-----
SAP 1/10/9.1
-----
Service Id : 501
SAP : 1/10/9.1
Description : Default sap description for service id 501
Admin 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
Ingr IP Fltr-Id : n/a
Ingr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a
tod-suite : None
Ing Agg Rate Limit : max
Endpoint : N/A
Acct. Pol : None
Oper MTU : 1514
Egr IP Fltr-Id : n/a
Egr Mac Fltr-Id : n/a
Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : n/a
Egr Agg Rate Limit: max
Collect Stats : Disabled
-----
QoS
-----
Ingress qos-policy : 1
Shared Q plcy : n/a
I. Sched Pol : (Not Specified)
E. Sched Pol : (Not Specified)
Egress qos-policy : 1
Multipoint shared : Disabled
-----
Sap Statistics
-----
Last Cleared Time : N/A
Packets
Octets

```

Forwarding Engine Stats (Ingress)		
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
Forwarding Engine Stats (Egress)		
Dropped	: 0	n/a
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
Ingress Queue 3 (Profile)		
Off. ColorIn	: 0	0
Off. ColorOut	: 0	0
Off. Uncolor	: 0	0
Dro. ColorOut	: 0	0
Dro. ColorIn/Uncolor	: 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.		
=====		

Output example (Ipipe service)

*A:ALU-A# show service id 1301 all	
=====	
Service Detailed Information	
=====	
Service Id	: 1301
Service Type	: Ipipe
Name	: epipel

```

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*: 10.1.10.4

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     : 10.1.10.3
SAP MAC Address   : 00:1a:f0:bd:ab:b0               Mac Refresh Inter*: 14400

```

Ipipe SAP ARP Entry Info

10.1.10.3 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 (Ingress)		
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

Forwarding Engine Stats (Egress)

Dropped	: 0	n/a
---------	-----	-----

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#
```

Output example (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 : 10.12.20.1          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   : 10.12.10.1
Configured CE IP : 10.12.10.1      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 (Ingress)
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

Forwarding Engine Stats (Egress)
Dropped           : 0          n/a

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

Ingress Queue 3 (Profile)
Off. ColorIn      : 0          0
Off. ColorOut     : 0          0
Off. Uncolor      : 0          0
Dro. ColorOut     : 0          0
Dro. ColorInUncolor : 0          0
For. InProf       : 0          0
For. OutProf      : 0          0

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#

```

Output example (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
CE Addr Discovery : disabled
SDP Bind Count : 1
-----
Service Destination Points(SDPs)
-----
Sdp Id 333:2 -(10.10.10.2)
-----
Description : RSVP_SdpToDut-B
SDP Id : 333:2
Split Horiz Grp : (Not Specified)
VC Type : Ipipe
Admin Path MTU : 0
Far End : 10.10.10.2
Type : Spoke
VC Tag : 0
Oper Path MTU : 1550
Delivery : MPLS

Admin State : Up
Acct. Pol : None
Ingress Label : 131069
Ing mac Fltr : n/a
Ing ip Fltr : n/a
Admin ControlWord : Not Preferred
Admin BW(Kbps) : 0
Last Status Change : 07/08/2011 07:50:47
Last Mgmt Change : 07/08/2011 07:50:44
Endpoint : N/A
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
Oper State : Up
Collect Stats : Disabled
Egress Label : 131068
Egr mac Fltr : n/a
Egr ip Fltr : n/a
Oper ControlWord : False
Oper BW(Kbps) : 0
Signaling : TLDP
Precedence : 4

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

Statistics :
I. Fwd. Pkts. : 125
I. Fwd. Octs. : 31832
E. Fwd. Pkts. : 489
I. Dro. Pkts. : 0
I. Dro. Octets : 0
E. Fwd. Octets : 31296

Associated LSP LIST :
Lsp Name : LSPToDut-B
Admin State : Up
Time Since Last Tr*: 00h03m46s
Oper State : Up
-----
IPIPE Service Destination Point specifics
-----
Configured CE IP Addr : 10.12.20.1
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   : 10.12.10.1
Configured CE IP : 10.12.10.1      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
Forwarding Engine Stats (Ingress)
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

```

Forwarding Engine Stats (Egress)

```

Dropped           : 0              n/a

```

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 50: Service-ID base field descriptions](#) describes the fields.

Output example (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
=====

Output example (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 50: Service-ID base field descriptions

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

Label	Description
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
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 51: Service egress label field descriptions](#) describes the fields.

Output example

```
*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 51: Service egress label field descriptions

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
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 52: Service ingress label field descriptions](#) describes the fields.

Output example

```
*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 52: Service ingress label field descriptions

Label	Description
Svc ID	Identifies the service

Label	Description
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 53: Service-ID endpoint field descriptions](#) describes the fields.

Output example

```
*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 53: Service-ID endpoint field descriptions

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
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 54: Service-ID labels field descriptions](#) describes the fields.

Output example

```
*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 54: Service-ID labels field descriptions

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

macsec

Syntax

macsec

Context

show>service>id

Description

This command displays MACsec security information for the specified service.

Output

The following output is an example of MACsec information, and [Table 55: Service-ID MACsec field descriptions](#) describes the fields.

Output example

```
*A:ALU-12# show service id 1 macsec
=====
MACsec (Summary), Service 1
=====
SAP          MACsec    MACsec    Encap    CA tags    CA-name
              port      sub-port  match    in-clear
-----
1/1/3        1/1/3      1         all      0          cal
=====
*A:ALU-12#
```

Table 55: Service-ID MACsec field descriptions

Label	Description
SAP	The service SAP
MACsec port	The port enabled for MACsec
MACsec sub-port	The subport enabled for MACsec
Encap match	The traffic encapsulation type to match: all traffic, untagged-only traffic, single-tag or dot1q traffic, double-tag or QinQ traffic
CA tags in-clear	The number of tags in clear text for this CA
CA-name	The name of the MACsec connectivity association for this SAP

network-latency-measurement

Syntax

network-latency-measurement

Context

show>service>id

Description


This command displays minimum, current, and maximum latency measurement values.


Output

The following outputs are examples of network latency measurement information, and [Table 56: Service-ID network latency measurement field descriptions](#) describes the fields.

Output example

```
A:SAR18-11-2# show service id 100 network-latency-measurement
=====
Path Average Latency Measurements (us)
-----
Spoke-sdp      Min      Current  Max      Last update
121:100        450      500      600      02/21/23 20:17:14
122:200        350      600      800      02/21/23 20:17:14
....
```

-  **Note:**
- The "Path Average Latency Measurements" output is only shown if the Cpipe has a timestamp. If a path is up but has not completed the first latency measurement, "N/A" is displayed. If a path that was up goes down, the last valid value is displayed.
 - For Cpipes with PW redundancy, only the active path average latency measurement is shown; other paths are not shown.
 - If a path is administratively down, "N/A" is displayed.
 - If a path is transported over a non-Ethernet interface (that is, timestamping is not available), "N/A" is displayed.

 **Note:** The maximum latency that can be measured is 34.3 seconds. For any latency above 34.3 seconds, the current latency is displayed as "Too High". If there is a clocking issue that results in a situation where the far-end timestamp is earlier than the near-end timestamp, the latency is negative. In this case, the current latency is displayed as "Too Low". In either case, the minimum and maximum latencies are not updated, so the most recent minimum and maximum values are displayed.

The following output is an example of network latency measurement information where the far-end timestamp is earlier than the near-end timestamp, resulting in a negative latency that is displayed as "Too Low".

Output example

```
A:SAR18-11-2# show service id 100 network-latency-measurement
=====
Path Average Latency Measurements (us)
-----
Spoke-sdp      Min      Cur      Max      Last update
321:1200       N/A      Too Low  N/A      Never
=====
```

Table 56: Service-ID network latency measurement field descriptions

Label	Description
Path Average Latency Measurements (us)	
Spoke-sdp	Identifies the spoke SDP associated with the Cpipe
Min	The minimum latency value at the time of the last update

Label	Description
Current	The current latency value at the time of the last update
Max	The maximum latency value at the time of the last update
Last-update	The time that the latency values were last updated

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 44: SAP ID configurations](#) 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 [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 [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 [sap-aggregation-group](#) command for information about how to show statistics for aggregation group members.

Output

The [Output example \(Apipe\)](#) is an example of service-id SAP information. [Table 57: Service-ID SAP field descriptions](#) describes the fields for the Apipe example. [Table 57: Service-ID SAP field descriptions](#) includes additional fields that appear in the outputs of the other VLL services, which are listed below and are included after the table:

- [Output example \(Cpipe\)](#)
- [Output example \(Epipe\)](#)
- [Output example \(Fpipe\)](#)
- [Output example \(Hpipe\)](#)
- [Output example \(Ipipe\)](#)
- [Output example \(Ipipe with frame relay SAP\)](#)
- [Output example \(Ipipe with cHDLC SAP\)](#)
- [Output example \(Apipe in a SAP aggregation group\)](#)

Output example (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 (Ingress)
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

Forwarding Engine Stats (Egress)
Dropped             : 0                      n/a

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
Ingress TD Ovr     : N/A                    Egress TD Ovr     : N/A
Alarm Cell Handling: Enabled
OAM Termination    : Disabled                Periodic Loopback : Disabled
AAL-5 Encap        : aal5mux-bridged-eth-nof*
=====

```

```
*A:ALU-12>show>service>id#
```

Table 57: Service-ID SAP field descriptions

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
Flags	Specifies the conditions that affect the operating status of this SAP Display output includes ServiceAdminDown, PortOper Down, 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

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
Ing Shaper Group	The ingress shaper group for the SAP
Egr Shaper Group	The egress shaper group for the SAP
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 (Ingress)	
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

Label	Description
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
Forwarding Engine Stats (Egress)	
Dropped	Indicates the number of packets or octets dropped by the forwarding engine
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 (Priority)	Specifies the index of the ingress QoS queue of this SAP, where n is the index number
Off. Combined	Indicates the combined total number of high-priority and low-priority packets or octets offered to the forwarding engine
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

Label	Description
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
Ingress Queue <i>n</i> (Profile)	Specifies the index of the ingress QoS queue of this SAP, where <i>n</i> is the index number
Off. ColorIn	Indicates the number of packets or octets colored as in-profile for the SAP (offered)
Off. ColorOut	Indicates the number of packets or octets colored as out-of-profile for the SAP (offered)
Off. Uncolor	Indicates the number of packets or octets that are unprofiled for the SAP (offered)
Dro. ColorOut	Indicates the number of packets or octets colored as out-of-profile that were dropped for the SAP
Dro. ColorIn/Uncolor	Indicates the number of packets or octets that were colored as in-profile or were unprofiled that were dropped for the SAP
For. InProf	Indicates the number of forwarded packets or octets that were colored as in-profile (FC profile set to "in" or "no profile" and rate less than or equal to CIR)
For. OutProf	Indicates the number of forwarded packets or octets that were colored as out-of-profile (FC profile set to "out" or "no profile" and rate above CIR)
Egress Queue <i>n</i>	Specifies the index of the egress QoS queue of the SAP, where <i>n</i> 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

Label	Description
Ma-index	Displays the value of the MA index
Admin	Displays the administrative state of the MEP (enabled or disabled)
Mepld	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
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

Label	Description
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
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
Ingress TD Ovr	The Ovr for the traffic descriptor applied to the ingress SAP
Egress TD Ovr	The Ovr for 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	The type of AAL5 encapsulation for this ATM SAP
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

Label	Description
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
ThresholdRepeat (usec)	Specifies the latency value that will be used by ADC repeat and on-demand ADC, in microseconds
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)
AMP timeout (s)	Specifies the active-multipath-timeout value in seconds
AMP State	Specifies the active multipath state (all-paths-active, not-all-paths-active, initialization, or down)
AMP Num Of Active Path	Specifies the number of active multipath paths
AMP Active Sdp Id	Specifies the active spoke SDP Note: This field is displayed only if the active multipath state is not-all-paths-active.
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
Time (µs)	(Column heading) Displays the number of microseconds elapsed for the statistic since the last counter reset

Label	Description
Time (sec)	(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
Jitter Buffer Depth	Displays the number of packets sitting in the jitter buffer at that instant for the Cpipe or MEF 8 Epipe
Jitter Buffer Delay	Displays the total amount of TDM PW data buffered
AMP Duplicate Dropped	Displays the number of AMP duplicate packets that have been dropped
AMP State Change Count	Displays the number of times the active multipath state has changed

Label	Description
ADC JB Sampling Complete	Displays the number of asymmetric delay control analysis periods completed
ADC JB Adjust	Displays the number of jitter buffer adjustments made for asymmetric delay control
ADC JB Sampling Avg Delay	Displays the average jitter buffer delay value of the samples taken during the last ADC analysis period
ADC JB Sampling Avg Delay	Displays the average jitter buffer delay value of the samples taken during the last ADC analysis period
Ingress Stats	Indicates that the following statistics are ingress statistics
Forwarded	Displays the number of forwarded packets
Dropped	Displays the number of dropped packets

Output example (Cpipe)

```

*A:7705:Dut-H# show service id 1 sap 1/4/15.1 detail
=====
Service Access Points(SAP)
=====
Service Id      : 1
SAP             : 1/4/15.1          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 : 01/24/2022 20:47:03
Last Mgmt Change  : 01/24/2022 20:46:25
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
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
Ing Shaper Group   : n/a             Egr Shaper Group : n/a
Endpoint          : N/A

Acct. Pol        : None              Collect Stats    : Disabled
-----
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   : 01/24/2022 20:47:38
Packets             :
Octets              :

```

```

Forwarding Engine Stats (Ingress)
Dropped                : 0                0
Off. HiPrio            : 108180            865440
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            : 108180            865440
For. OutProf           : 0                0
Forwarding Engine Stats (Egress)
Dropped                : 0                n/a
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            : 108180            865440
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           : 8                  Jitter Buffer (ms) : 8
Jitter Buffer (packets) : 8                  Playout Threshold (packets): 5
Use RTP Header         : No                Differential      : No
Timestamp Freq         : 0                  CAS Framing       : No CAS
Effective PDVT         : +/-4.0 ms

AsymDelayControl       : enabled                ThresholdRepeat (usec) : 250
RepeatPeriod           : disabled                Samples           : 2K
AMP timeout(s)         : 60
AMP State              : all-paths-active    AMP Num Of Active Path : 2
Cfg Alarm              : stray malformed pktloss overrun underrun
Alarm Status          :

-----
CEM SAP Statistics
-----
Packets                Time(us)                Time(sec)                Events
Egress Stats
Forwarded              : 115575
Dropped                : 0
Missing               : 0
Reordered Forwarded   : 0
Underrun              : 0                                0
Overrun               : 0                                0
Misordered Dropped    : 0
Malformed Dropped     : 0
LBit Dropped          : 0
Error                 :                                0

```

```

Severely Error      : 0
Unavailable         : 0
Failure Count       : 0
Jitter Buffer Depth : 4
Jitter Buffer Delay  : 4250
AMP Duplicate Dropped : 115575
AMP State Change Count : 0
ADC JB Sampling Complete : 0
ADC JB Adjust       : 0
ADC JB Sampling Avg Delay : 4052
ADC JB Target Delay  : 4000
Ingress Stats
Forwarded           : 115575
Dropped             : 0
=====
*A:7705:Dut-H#

```

Output example (Epipe)

```

*A:csasim2>show>service>id# sap 1/3/1 detail
=====
Service Access Points(SAP)
=====
Service Id      : 3
SAP             : 1/3/1
Admin State     : Up
LLF Admin State : Up
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
Ingr IP Fltr-Id : n/a
Ingr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a
tod-suite       : None
Ing Scheduler Mode : 16-priority
Ing Agg Rate Limit : 50000
Ing Agg cir      : 5000
Ing Shaper Group : default
Endpoint        : N/A
Q Frame-Based Acct : Disabled
Vlan-translation : None

Acct. Pol       : None
Collect Stats   : Disabled

-----
QOS
-----
Ingress qos-policy : 1
Shared Q plcy      : n/a
Egress qos-policy  : 1
Multipoint shared  : Disabled

-----
Sap Statistics
-----
Last Cleared Time : 05/07/2008 21:32:32

Forwarding Engine Stats (Ingress)
Dropped : 0
Off. HiPrio : 2655264
Packets
Octets

```

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	
Forwarding Engine Stats (Egress)			
Dropped	: 0	n/a	
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	
Ingress Queue 3 (Profile)			
Off. ColorIn	: 0	0	
Off. ColorOut	: 0	0	
Off. Uncolor	: 0	0	
Dro. ColorOut	: 0	0	
Dro. ColorIn/Uncolor	: 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
LbNextSequence   : 1
LtRxUnexplained   : 0

Md-index         : 1
Ma-index         : 1
MepId            : 1
LowestDefectPri   : macRemErrXcon
Defect Flags     : None
Mac Address      : 00:00:00:00:00:00
CcmTx            : 0
DmrRepliesTx     : 0
LmrRepliesTx     : 0
Dual-Loss Test   : Disabled
Eth-Ais:         : Disabled
Eth-Tst:         : Disabled
LbRxReply        : 0
LbRxBadMsdu      : 0
LbNextSequence   : 1
LtRxUnexplained   : 0

LbTxReply         : 0
LtNextSequence    : 1

Direction         : Down
Admin             : Disabled
CCM-Enable        : Enabled
HighestDefect     : none

CcmSequenceErr    : 0

LbRxBadOrder      : 0
LbTxReply         : 0
LtNextSequence    : 1
=====

```

See [Table 57: Service-ID SAP field descriptions](#) for Epipe field descriptions.

Output example (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
Description      : (Not Specified)
Admin 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
Ingr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a
tod-suite       : None
Ing Agg Rate Limit : max
Endpoint        : N/A
FRF-12          : Disabled

Oper MTU         : 1600
Egr IP Fltr-Id   : n/a
Egr Mac Fltr-Id  : n/a
Egr IPv6 Fltr-Id : n/a
qinq-pbit-marking : both
Egr Agg Rate Limit: max

Acct. Pol       : None
Collect Stats   : Disabled

-----
QoS
-----
Ingress qos-policy : 1
Shared Q plcy      : n/a
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
Egress qos-policy  : 1
Multipoint shared  : Disabled

-----
Sap Statistics
-----
Last Cleared Time   : N/A

```

	Packets	Octets
Forwarding Engine Stats (Ingress)		
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
Forwarding Engine Stats (Egress)		
Dropped	: 0	n/a
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
Ingress Queue 3 (Profile)		
Off. ColorIn	: 0	0
Off. ColorOut	: 0	0
Off. Uncolor	: 0	0
Dro. ColorOut	: 0	0
Dro. ColorIn/Uncolor	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0
Egress Queue 1		
For. InProf	: 0	0
For. OutProf	: 1803	897894
Dro. InProf	: 0	0
Dro. OutProf	: 0	0
=====		

See [Table 57: Service-ID SAP field descriptions](#) for Fpipe field descriptions.

Output example (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
Shared Q plcy      : n/a
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
Egress qos-policy  : 1
Multipoint shared  : Disabled
-----
Sap Statistics
-----
Last Cleared Time   : N/A
Packets
Forwarding Engine Stats (Ingress)
Dropped             : 0
Off. HiPrio         : 0
Off. LowPrio        : 167
Octets
15030

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio         : 0
Dro. LowPrio        : 0
For. InProf         : 0
For. OutProf        : 167
Octets
15030

Forwarding Engine Stats (Egress)
Dropped             : 0
Octets
n/a

Queueing Stats(Egress QoS Policy 1)
Dro. InProf         : 0
Dro. OutProf        : 0
For. InProf         : 0
For. OutProf        : 165
Octets
14850
-----
Sap per Queue stats
-----
Packets
Octets
Ingress Queue 1 (Priority)
Off. HiPrio         : 0
Off. LoPrio         : 167
Dro. HiPrio         : 0
Dro. LoPrio         : 0
For. InProf         : 0
For. OutProf        : 167
Octets
15030
Egress Queue 1
For. InProf         : 0
For. OutProf        : 165
Dro. InProf         : 0
Dro. OutProf        : 0
Octets
14850
=====

```

See [Table 57: Service-ID SAP field descriptions](#) for Hpipe field descriptions.

Output example (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   : 10.1.10.3
SAP MAC Address : 00:1a:f0:bd:ab:b0 Mac Refresh Inter*: 14400
```

```
-----
Ipipe SAP ARP Entry Info
-----
```

```
10.1.10.3      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
Forwarding Engine Stats (Ingress)
Dropped           : 0
Off. HiPrio       : 20683584      1613319552
Off. LowPrio      : 1840167       143533026

Queueing Stats(Ingress QoS Policy 13)
Dro. HiPrio       : 0
Dro. LowPrio      : 0
For. InProf       : 11271525      879178950
For. OutProf      : 11252226      877673628

Forwarding Engine Stats (Egress)
Dropped           : 0
n/a

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 57: Service-ID SAP field descriptions](#) for lpipe field descriptions.

Output example (lpipe 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

-----
lpipe SAP Configuration Information
-----
Ce IP Address   : 10.12.10.1
Configured CE IP : 10.12.10.1          Discovered CE IP : n/a

-----
lpipe SAP ARP Entry Info
-----
No lpipe 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 (Ingress)
Dropped          : 0
Off. HiPrio      : 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

Forwarding Engine Stats (Egress)
Dropped          : 0          n/a

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 57: Service-ID SAP field descriptions](#) for Ipipe with frame relay field descriptions.

Output example (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   : 10.12.10.1
Configured CE IP : 10.12.10.1      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   : 10.12.10.1
Configured CE IP : 10.12.10.1      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 (Ingress)
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

Forwarding Engine Stats (Egress)
Dropped       : 0              n/a

```

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 57: Service-ID SAP field descriptions](#) for lpipe with cHDLc field descriptions.

Output example (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 57: Service-ID SAP field descriptions](#) 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 58: Service-ID SAP aggregation group field descriptions](#) describes the fields.

Output example

```
*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
-----
Forwarding Engine Stats (Ingress)
-----
Dropped      : 0
Off. HiPrio  : 77780484
Off. LowPrio : n/a
-----
Queueing Stats(Ingress QoS Policy 1)
-----
Dro. HiPrio  : 0
Dro. LowPrio : n/a
For. InProf  : 0
-----
Packets
-----
Octets
-----
```



```

For. OutProf      : 77780484      5755755816
Forwarding Engine Stats (Egress)
Dropped          : 0              n/a

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 58: Service-ID SAP aggregation group field descriptions

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
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

Label	Description
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 (Ingress)	
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
Forwarding Engine Stats (Egress)	
Dropped	Indicates the number of packets or octets dropped by the forwarding engine
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 Aggregation Group per Queue stats	

Label	Description
Ingress Queue <i>n</i> (Unicast) (Priority)	Specifies the index of the ingress QoS queue of the SAP, where <i>n</i> 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 <i>n</i>	Specifies the index of the egress QoS queue of the SAP, where <i>n</i> 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 eth-ring [ring-id *eth-ring-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}

sap-using [ingress | egress] shaper-group *shaper-group-name*

Context

show>service

Description

This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters

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

shaper-group

specifies the shaper group for which to display matching SAPs

eth-ring-id

specifies the Ethernet ring ID for which to display matching SAPs

Values 1 to 128

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

shaper-group-name

identifies the ingress or egress shaper group 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 44: SAP ID configurations](#) for a full list of SAP IDs.

Output

The following outputs are examples of service SAP-using information, and [Table 59: Service SAP-using field descriptions](#) describes the fields.

Output example

```
*A:ALU-48# show service sap-using
=====
Service Access Points
=====
```

PortId	SvcId	Ingr. QoS	Ingr. Fltr	Egr. QoS	Egr. Fltr	Adm	Opr
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	Ingr. QoS	Ingr. Fltr	Egr. QoS	Egr. Fltr	Adm	Opr
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 eth-ring ring-id 1

=====

Service Access Points (Ethernet Ring)

=====

SapId	SvcId	Eth-Ring	Path	Admin State	Oper State	Blocked	Control/Data
1/1/1:1.0	1	1	b	Up	Up	No	Ctrl
lag-32:1.0	1	1	a	Up	Up	No	Ctrl
1/1/1:2.2	5	1	b	Up	Up	No	Data
lag-32:2.2	5	1	a	Up	Up	No	Data
1/1/1:7.7	15	1	b	Up	Up	No	Ctrl
lag-32:7.7	15	1	a	Up	Up	No	Ctrl

*A:ALU-48# show service sap-using egress atm-td-profile 1

=====

Service Access Point Using ATM Traffic Profile 1

=====

PortId	SvcId	Ing. QoS	Ing. Fltr	Egr. QoS	Egr. Fltr	Adm	Opr
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: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#

```

```

*A:7705custDoc:Sar18>show>service# sap-using ingress shaper-group test_sg1
=====
Service Access Points Using Ingress Shaper Group "test_sg1"
=====
PortId          SvcId      Scheduler Mode    Shaper Policy      Opr
-----
1/2/1           30         4-priority      test_shaper_policy  Down
-----
Number of SAPs : 1
=====

*A:Sar18 Dut-B>config>service>epipe>sap>ingress#

```

Output example (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
  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
  1/2/9.1:2/301           Up   Up
-----
Number of Sap Aggregation Groups : 2
=====

```

Table 59: Service SAP-using field descriptions

Label	Description
GroupName	Displays the SAP aggregation group identifier
SvcId	Identifies the service
Scheduler Mode	Identifies the scheduler mode: 4-priority or 16-priority
Shaper Policy	Identifies the shaper policy that the shaper group belongs to
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
SapID	Identifies the SAP
Eth-Ring	Identifies the Ethernet ring
Path	Identifies the Ethernet ring path, either A or B
Admin State	Displays the administrative state
Oper State	Displays the operational state
Blocked	Identifies whether the SAP is blocked

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 60: Service SCADA bridge SAP field descriptions](#) describes the fields.

Output example

```
*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 60: Service SCADA bridge SAP field descriptions

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 61: SDP field descriptions](#) describes the fields.

Output example (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          : CEsPSN              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 61: SDP field descriptions

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

Label	Description
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
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, Sdp OperDown, 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, psn EgressFault, lacIngressFault, lacEgressFault

Label	Description
Peer Fault Ip	N/A
Peer Vccv CV Bits	Displays the setting of the pseudowire peer VCCV control verification bits (IspPing)
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
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)
Mepld	Displays the MEP-ID

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
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	

Label	Description
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
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 62: SDP-using field descriptions](#) describes the fields.

Output example

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 62: SDP-using field descriptions

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 63: Service-using field descriptions](#) describes the fields.

Output example (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 63: Service-using field descriptions

Label	Description
ServiceID	Identifies the service

Label	Description
Type	Identifies the service type
Adm	Specifies the administrative state of this service
Opr	Specifies the operational state of this service
CustomerId	Specifies the ID of the customer
Service Name	The service name

4.13.2.3 Clear commands

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*

arp

Syntax

arp

Context

clear>service>id

Description

This command clears the ARP entries from an Ipipe service.

counters

Syntax

counters

Context

clear>service>statistics>id

Description

This command clears all traffic queue counters associated with the service ID.

network-latency-measurement

Syntax

network-latency-measurement

Context

clear>service>id

Description

This command clears all network latency statistics associated with the service ID.

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

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 44: SAP ID configurations](#) 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

5 VPLS

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](#)

5.1 VPLS overview

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 (such as Epipe and Ipipe) 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 because 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 a 6-port Ethernet 10Gbps Adapter card with CLI identifier **a6-eth-10G** or **a6-eth-10G-e** installed on a 7705 SAR-8 Shelf V2 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 Shelf V2 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-M with CLI identifier **i7-1gb**
 - on a 4-port SAR-H Fast Ethernet module
 - on a 6-port SAR-M Ethernet module
 - on a 7705 SAR-A
 - on a 7705 SAR-Ax
 - on a 7705 SAR-Wx
 - 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 Shelf V2 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 scalability can be enabled with the **vpls-high-scale** command, allowing the 7705 SAR-8 Shelf V2 and 7705 SAR-18 to support up to 255 VPLS instances on access and network links on the following cards:

- 2-port 10GigE (Ethernet) Adapter card
- 8-port Gigabit Ethernet Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card
- Packet Microwave Adapter card
- 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card

See the 7705 SAR Basic System Configuration Guide, "System Command Reference", for more information.

5.1.1 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. See [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. The VLAN range covering SAPs to be managed by a mVPLS instance is set under a

specific mVPLS SAP definition. The 7705 SAR supports RSTP on the designated VLAN for fast detection of failures. See [VPLS and spanning tree protocol](#) for more information.

5.1.2 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

5.1.3 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](#).

5.1.4 VPLS packet walkthrough

This section describes an example of VPLS processing based on the network shown in the following figure.

Figure 80: VPLS service architecture

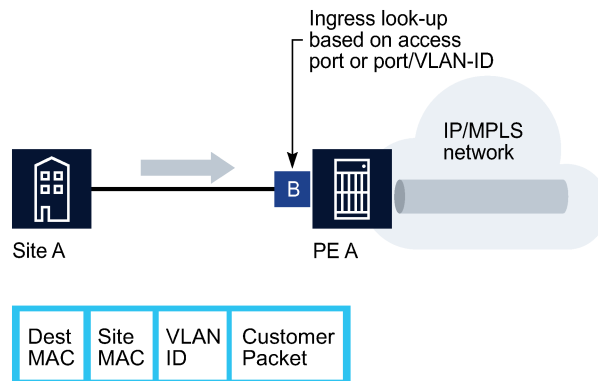


21559

1. PE router A ([Figure 81: Access port ingress packet format and lookup](#)):

- 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 81: Access port ingress packet format and lookup



17230


For a known MAC Address ([Figure 82: Network port egress packet format and flooding](#)):

- 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 when 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 82: Network port egress packet format and flooding](#)):

- 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 because this VPLS service does not exist on that PE router.

Figure 82: Network port egress packet format and flooding



21560

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 switching 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:

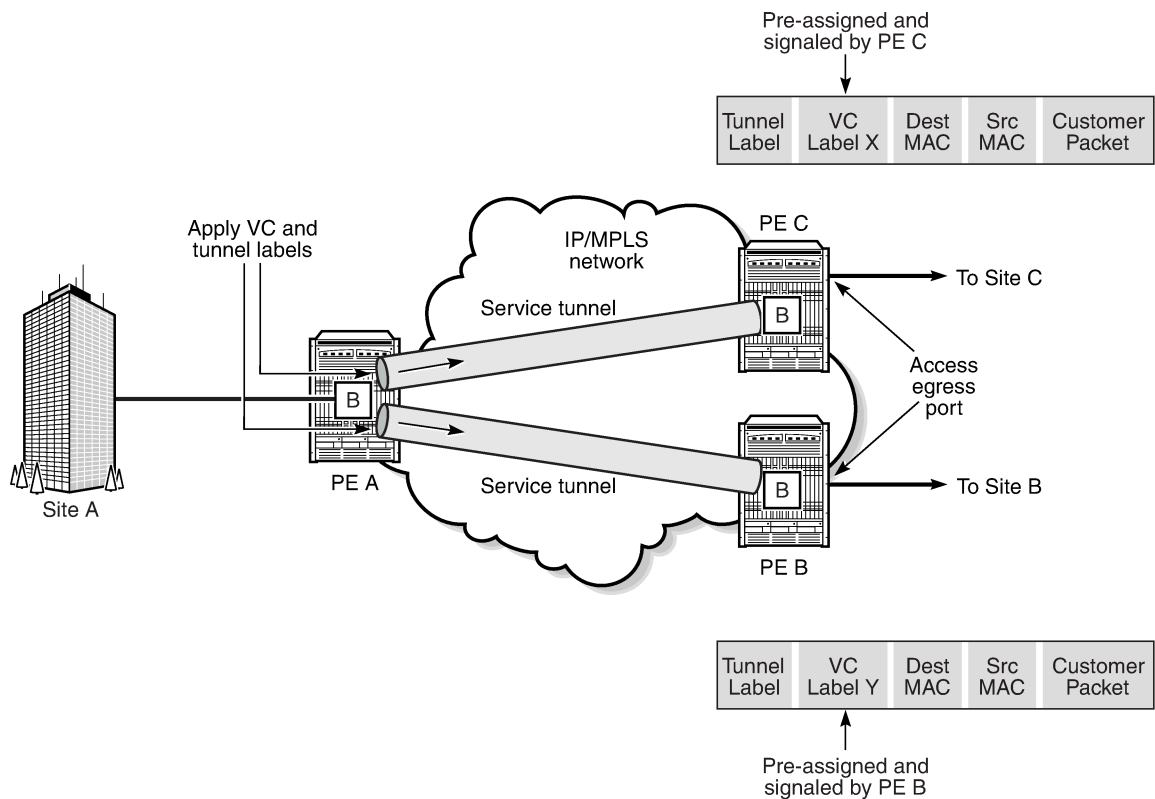
- a. PE C strips the transport label of the received packet to reveal the inner VC label. The VC label identifies the VPLS service instance to which the packet belongs.
- b. PE C learns the source MAC address in the packet and creates an entry in the FDB table that associates the MAC address to PE A 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 83: Access port egress packet format and lookup](#)):

- 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 83: Access port egress packet format and lookup):
- e. PE C will flood packets, as applicable.

Figure 83: Access port egress packet format and lookup




21560

5.1.5 Bridged mobile backhaul

The following figure 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 84: Typical pseudowire-based mobile backhaul



21553


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 85: Local VPLS on 7705 SAR in mobile backhaul](#), 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.

For the VPLS-3G example in the figure, 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, the 7705 SAR switches the received Ethernet frame to its destination. If the destination MAC address is not known, 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 85: Local VPLS on 7705 SAR in mobile backhaul



21554

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 the figure, 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).

5.1.6 Multi-tenant unit termination


Figure 85: Local VPLS on 7705 SAR in mobile backhaul 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 multi-tenant unit (MTU), the 7705 SAR is spoke SDP-terminated to a VPLS node (7750 SR node in the figure), eliminating the necessity to have a full mesh architecture for all VPLS-enabled nodes. Therefore the mesh requirement is "pushed" to the core nodes only (that is, to the 7750 SR nodes).

The 7750 SR nodes can be replaced by 7705 SAR nodes. Figure 86: Spoke-SDP termination to VPLS using 7705 SAR-18 routers 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 86: Spoke-SDP termination to VPLS using 7705 SAR-18 routers, the node that initiates the redundant PWs is the cell site 7705 SAR). Therefore, 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 86: Spoke-SDP termination to VPLS using 7705 SAR-18 routers



21555

5.2 VPLS features

Topics in this section include:

- [VPLS enhancements](#)
- [Fabric mode](#)
- [Subscriber VLAN](#)
- [ATM encapsulated residential SAPs](#)
- [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 \(IGMP and MLD snooping\)](#)
- [PIM snooping for VPLS](#)
- [MPLS entropy label](#)

- [Ethernet OAM](#)
- [Security zones and VPLS](#)

5.2.1 VPLS enhancements


The Nokia 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

The following figure illustrates VPLS enhancements using an example of ATM DSLAM backhaul, where the 7705 SAR may not be used solely for DSLAM backhaul purposes or not all the services might be bound to VPLS. In the figure, colocated IP DSLAM (ISAM) traffic can also be transported by the 7705 SAR.

Figure 87: ATM and IP DSLAM backhaul

Ingress



21599

5.2.2 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). VPLS service is therefore 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.

5.2.3 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 may 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. 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).

Because 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.

5.2.4 ATM encapsulated residential SAPs

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)

5.2.5 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

5.2.6 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.

5.2.7 Pseudowire control word


The **control-word** command enables the use of the control word individually on each mesh SDP or spoke SDP. The control word is optional and is disabled by default. 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.

5.2.8 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. The following figure illustrates the signaling.

Figure 88: PPPoE initialization and agent ID push function



21598

The following figure 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 89: Agent circuit ID information



21381

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.

5.2.9 MAC filters

MAC filters offer the ability to transport Ethernet frames that match specific 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 specific 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 "Filter policies" in the 7705 SAR Router Configuration Guide.

5.2.10 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

These MAC table management features are required for each instance of a SAP or spoke SDP within a particular VPLS instance.

5.2.10.1 FDB size

Users can configure the MAC FDB size limits 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, 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

5.2.10.2 FDB size alarms

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

5.2.10.3 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 ingress 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 about two times their respective age timers to be aged out.

5.2.10.3.1 Disable MAC aging

Users can disable MAC aging timers to 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.

5.2.10.3.2 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.

5.2.10.4 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.

5.2.11 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 about QoS policies for broadcast, multicast, and unknown (BMU) traffic, see "Filter policies" in the 7705 SAR Quality of Service Guide.

5.2.12 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 always follows 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).

5.2.13 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 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.

5.2.13.1 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.

5.2.14 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 $2 \times n$ 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 about the following topics:

- [Application examples](#)
- [Group and addressing support](#)
- [IP multicast in r-VPLS](#)
- [Multicast router ports](#)

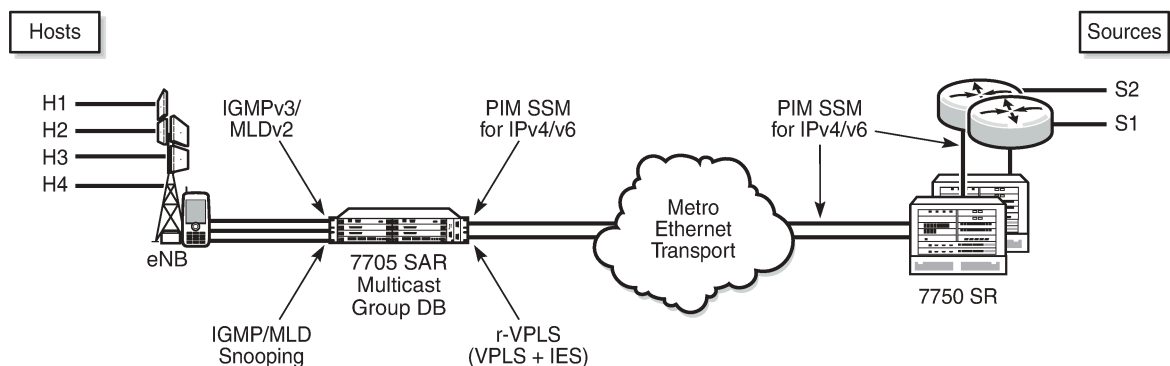
- [Tagged access traffic](#)
- [Hardware support](#)

For more information about multicast, see "IP multicast" in the 7705 SAR Routing Protocols Guide.

5.2.14.1 Application examples

The following figure 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 7705 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 7705 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 90: Video over mobile backhaul



25374

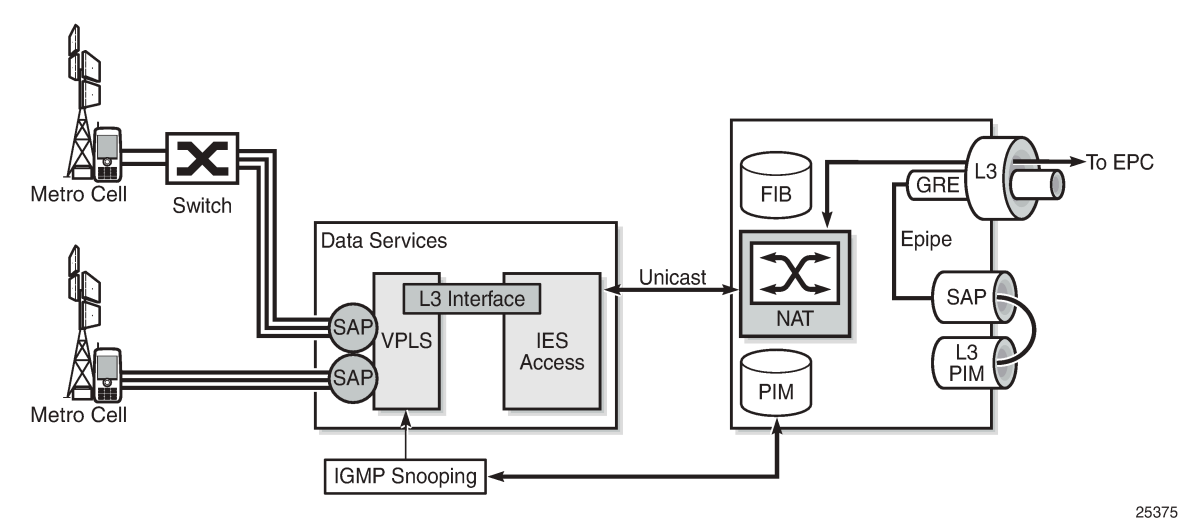
Figure 91: Metro cell multicast shows another example of IGMP and MLD snooping, where service providers can offer evolved multimedia broadcast multicast services (eMBMS) on their metro cell network (data services).

In the figure, 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.

The figure 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.

Figure 91: Metro cell multicast



25375

5.2.14.2 Group and addressing support

This section contains information about the following topics:

- [IPv4 multicast support](#)
- [IPv6 Layer 2 multicast support and group address](#)

5.2.14.2.1 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.

5.2.14.2.2 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 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.

5.2.14.3 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. When 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.

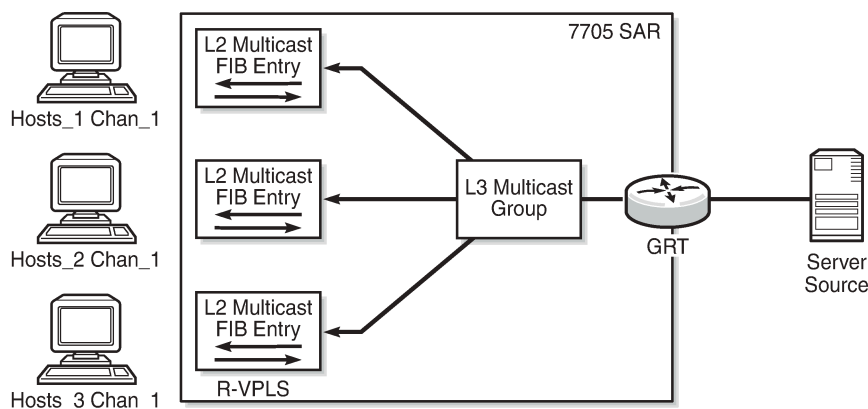


Note: IP multicast on r-VPLS is supported only if IGMP or MLD snooping is enabled on the VPLS associated with the IES. That is, configuring IGMP or MLD on the IES that is associated with the VPLS without configuring snooping on the VPLS will not flood multicast traffic out the VPLS.

[Figure 92: Multiple hosts in an r-VPLS receiving the same channel](#) and [Figure 93: Multiple hosts in an r-VPLS receiving different channels](#) illustrate how Layer 2 multicast interacts with Layer 3 multicast.

In the following figure, there are three hosts and one channel. The Layer 3 multicast group forwards source traffic to each port for which there is a corresponding Layer 2 multicast instance of a Layer 2 FIB entry. All three Layer 2 FIB entries are within the same r-VPLS. To configure this scenario, create three Layer 2 FIB entries 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 92: Multiple hosts in an r-VPLS receiving the same channel




25376

In the following figure, 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). Therefore, 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 the following figure, 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 93: Multiple hosts in an r-VPLS receiving different channels



25377

5.2.14.3.1 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.

5.2.14.4 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**.

5.2.14.5 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.

5.2.14.6 Hardware support

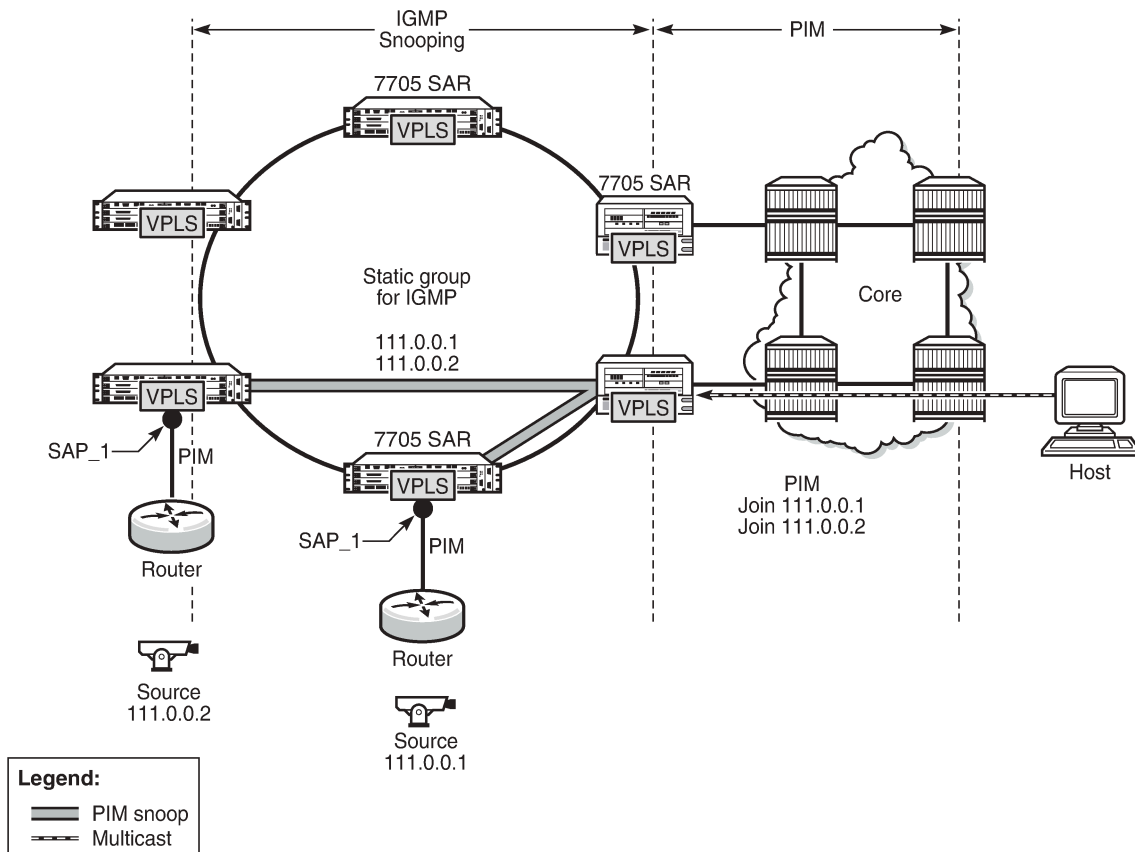
Multicast VPLS and r-VPLS are supported on the following 7705 SAR hardware:

- 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-Ax, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-Wx, and 7705 SAR-X
- 4-port SAR-H Fast Ethernet module
- hardware that supports PPP/MLPPP for uplink spoke SDPs

5.2.15 PIM snooping for VPLS

PIM snooping is used in Layer 2 networks to stitch together the PIM session of two disjointed Layer 3 networks. In most provider networks, strategic industry (SI) applications, or mobile backhaul applications, the access routers are connected to the core Layer 3 network via a Layer 2 network. For multicast scenarios, PIM can be used to build the multicast data trees (MDTs) on the Layer 3 routers. However, PIM is a Layer 3 protocol and Layer 2 networks do not understand PIM messages. This creates an inefficient multicast domain in the Layer 2 network, as all packets will be broadcast. PIM snooping in a Layer 2 network can be used to stitch the PIM session from the access routers to the core Layer 3 network. The following figure illustrates this scenario.

Figure 94: PIM snooping example




25920

Because PIM snooping stitches the PIM session between two routers, it may be desirable to start the PIM session from the same router (a 7705 SAR) that initiates the VPLS PIM snooping service. In this case, an external loop cable can be used to connect the network Layer 3 interface, which has PIM configured, to the SAP of the VPLS that is performing the PIM snooping (as shown in the following figure), thereby simulating routed VPLS.

PIM snooping supports both (*,G) as well as (S,G) to address scenarios where the sources only support PIM ASM.

The 7705 SAR supports PIM snooping for IPv4 and IPv6.

Figure 95: Simulating r-VPLS



25921

5.2.15.1 Snooping versus proxy mode

Snooping mode does not stop or intercept a PIM message. Snooping listens to network traffic between hosts and routers, and maintains a table that maps multicast streams between sources and hosts.

Proxy mode intercepts PIM messages and generates a single message when necessary. Proxy allows a switch to send PIM messages on behalf of routers. For example, when multiple routers are connected to the same PIM-enabled switch and all the routers want to join to the same source, the proxy can intercept the messages and generate a single message in order to minimize the flood of PIM messages in the Layer 2 domain.

Configure snooping or proxy mode by using the **config>service>vpls> pim-snooping>mode** command. By default, proxy mode is enabled in VPLS.

5.2.16 MPLS entropy label

The router supports the MPLS entropy label as per RFC 6790. The entropy label provides greater granularity for load balancing on an LSR where load balancing is typically based on the MPLS label stack.

For more information, see "MPLS entropy labels" in the 7705 SAR MPLS Guide and "LAG and ECMP hashing" in the 7705 SAR Interface Configuration Guide.

5.2.17 Ethernet OAM

The 7705 SAR supports Ethernet OAM functions for ETH-CFM (according to the 802.1ag and Y.1731 standards) and for Y.1731 performance monitoring on VPLS Ethernet SAPs and SDPs. VPLS OAM MAC diagnostics are also supported.

For information about Ethernet OAM, see [ETH-CFM \(802.1ag and Y.1731\)](#); see also "Ethernet OAM capabilities" and "VPLS MAC diagnostics" in the 7705 SAR OAM and Diagnostics Guide.

5.2.18 Security zones and VPLS

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, see "Configuring security parameters" in the 7705 SAR Router Configuration Guide.

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 2 endpoints or 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. Layer 3 zones support both NAT and firewall security policies. Layer 2 zones support only firewalls. 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>epipe>zone**
- **config>service>vpls>zone**
- **config>service>vprn>zone**
- **config>service>ies>zone**

Layer 2 and Layer 3 firewalls share system resources; that is, they share the maximum number of policies, profiles, and session ID space supported by the system.

A zone is created by adding at least one Layer 2 endpoint or Layer 3 interface to the zone configuration. Multiple zones can be created within each Layer 3 service or within the router context. Layer 2 services support only one zone. Layer 2 endpoints or Layer 3 interfaces from different services cannot be grouped into a single common zone. The following table lists the supported interfaces and endpoints that can be added to zones in each CLI context for NAT or firewall.

Table 64: Security zone interfaces and endpoints per context

CLI context	Interface/endpoint type	NAT	Firewall
Router	Layer 3	✓	✓
Epipe	SAP		✓
	Spoke-SDP termination		✓
VPLS	SAP		✓
	Spoke-SDP termination		✓
	Mesh SDP		✓
	EVPN		
VPRN	SAP	✓	✓

CLI context	Interface/endpoint type	NAT	Firewall
	Spoke-SDP termination	✓	✓
	IPSec private	✓	✓
	IPSec public	✓	
	Routed VPLS	✓	✓
IES	SAP	✓	✓
	Spoke-SDP termination	✓	✓
	IPSec public	✓	
	Routed VPLS	✓	✓



Note: A group of endpoints used for pseudowire redundancy cannot be added to a zone configured under an Epipe.

A zone configured under a Layer 2 service (VPLS or Epipe context) allows the 7705 SAR to perform Layer 3 firewall functionality on IPv4 packets.

NAT is not supported for zones configured under a Layer 2 service. A zone cannot be configured on a VPLS service with EVPN.

Unicast, multicast, and broadcast IPv4 packets are firewalled when they cross a Layer 2 service zone boundary.

If routed VPLS is configured on a VPLS service, and traffic is traversing between two Layer 2 endpoints, the firewall security policies for the Layer 2 service are used. If the traffic is traversing between Layer 2 and Layer 3, the Layer 3 security policy in IES or VPRN is used. The system ignores the firewall rules for the Layer 2 service and instead uses the Layer 3 NAT and firewall rules. Firewall rules for Layer 2 services are always ignored when traffic is traversing between Layer 2 and Layer 3, even when there are no Layer 3 firewall rules.

Users can configure bypass policies to allow specific traffic, such as control plane protocols (OSPF, RIP, BGP, IGMP, PIM, LDP, RSVP, VRRP, DHCP, NTP, and so on) to bypass a firewall in a Layer 2 service. For information, see "Bypass policies for firewalls in a Layer 2 service" in the 7705 SAR Router Configuration Guide. If not configured to bypass the zone, these packets are firewalled as normal unicast, multicast, or broadcast traffic and should be regulated by configuring firewall security policies for these protocols.

5.3 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 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.

5.3.1 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](#)

5.3.1.1 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.

5.3.1.2 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 scans for a VPLS service context using the service name and takes the following actions:

- if the service name is not currently in use by a service, the IP interface is placed in an operationally down: Non-existent service name or inappropriate service type state
- if the service name is currently in use by a VPLS service without the **allow-ip-int-binding** flag set, the IP interface is placed in the operationally down: VPLS service allow-ip-int-binding flag not set state
- 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 is attached to the VPLS service

5.3.1.3 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.

5.3.1.4 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.

5.3.2 IP interface MTU and fragmentation

The VPLS service is affected by two MTU values: port MTUs and the VPLS service MTU. The MTU on each physical port defines the largest Layer 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, a 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, 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.

5.3.3 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.

5.3.3.1 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 toward. The following tables describe how the ARP cache and MAC FIB entry states interact at ingress and the corresponding egress behavior.

Table 65: 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 66: 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

5.3.4 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.

When 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.

5.3.4.1 VPLS feature restrictions with allow-ip-int-binding

If the allow-ip-int-binding flag is set on a VPLS service, the following features are disabled:

- residential SHG
- DHCP

- mVPLS
- mac-subnet-length
- GRE SDP (cannot be bound to the VPLS)



Note: The DHCP relay functionality under **config>service>ies>if>dhcp** or **config>service>ies>if>ipv6>dhcp6-relay** can be used to dynamically assign IP addresses to the clients connected to routed VPLS SAPs.

5.3.5 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.

5.3.6 VPLS ingress IP filter override

If an IP interface is attached to a VPLS service context, the VPLS SAP or SDP configured IP or MAC 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 filter override is configured, the IP or MAC filter configured on the SAP or SDP applies to non-routed packets. If filter override is not configured, the IP or MAC filter configured on the SAP or SDP applies to both routed and non-routed packets.

5.3.7 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 and IPv6)
- ARP
- ND protocol
- DHCP (IPv4 and IPv6)

5.4 VPLS and spanning tree protocol

Topics in this section include:

- [VPLS redundancy](#)
- [VPLS access redundancy](#)
- [MAC flush message processing](#)

The Nokia 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 Nokia 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 about command usage, descriptions, and CLI syntax, see [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.

5.4.1 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. Because node redundancy is implicit in non-hierarchical VPLS services (using a full mesh of SDPs between PEs), node redundancy for spoke SDPs needs to be provided separately.

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

5.4.1.1 Spoke SDP redundancy for metro interconnection

When two or more meshed VPLS instances, such as in the following figure, are interconnected by redundant spoke SDPs, a loop in the topology results. 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 96: H-VPLS with spoke SDP redundancy



21877

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.

5.4.1.2 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 96: H-VPLS with spoke SDP redundancy](#), 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 SDP.

The secondary spoke SDP is in a blocking state (both on receive and transmit) as long as the primary spoke SDP is available. If the primary spoke SDP becomes unavailable (due to link failure, PEs failure, and so on), the MTU immediately switches traffic to the backup spoke SDP and starts receiving traffic from the standby spoke SDP. 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 SDP when a spoke SDO 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.

5.4.2 VPLS access redundancy

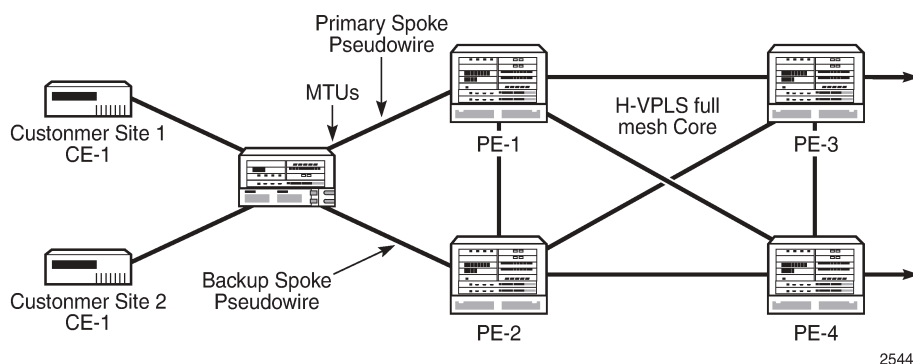
Another application of hierarchical VPLS uses MTUs that are not MPLS-enabled and must have Ethernet links to the closest PE node. To protect against failure of the PE node, an MTU can be dual-homed and have two SAPs on two PE nodes.

On the 7705 SAR, the mechanism used to resolve a loop in an access circuit uses STP-based access, with or without mVPLS.

5.4.2.1 STP-based redundant access to VPLS

In the configuration shown in the following figure, 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 97: 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 STP instance.

5.4.3 MAC flush message processing

The previous sections describe 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 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 VPLS service. The message is sent to all LDP peers connected in the context of the 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 an MC-LAG SAP or MC-APS SAP transitions from an operationally up state to an operationally down state. The message is sent to all LDP peers connected in the context of the VPLS service.

5.4.3.1 Dual homing to a VPLS service

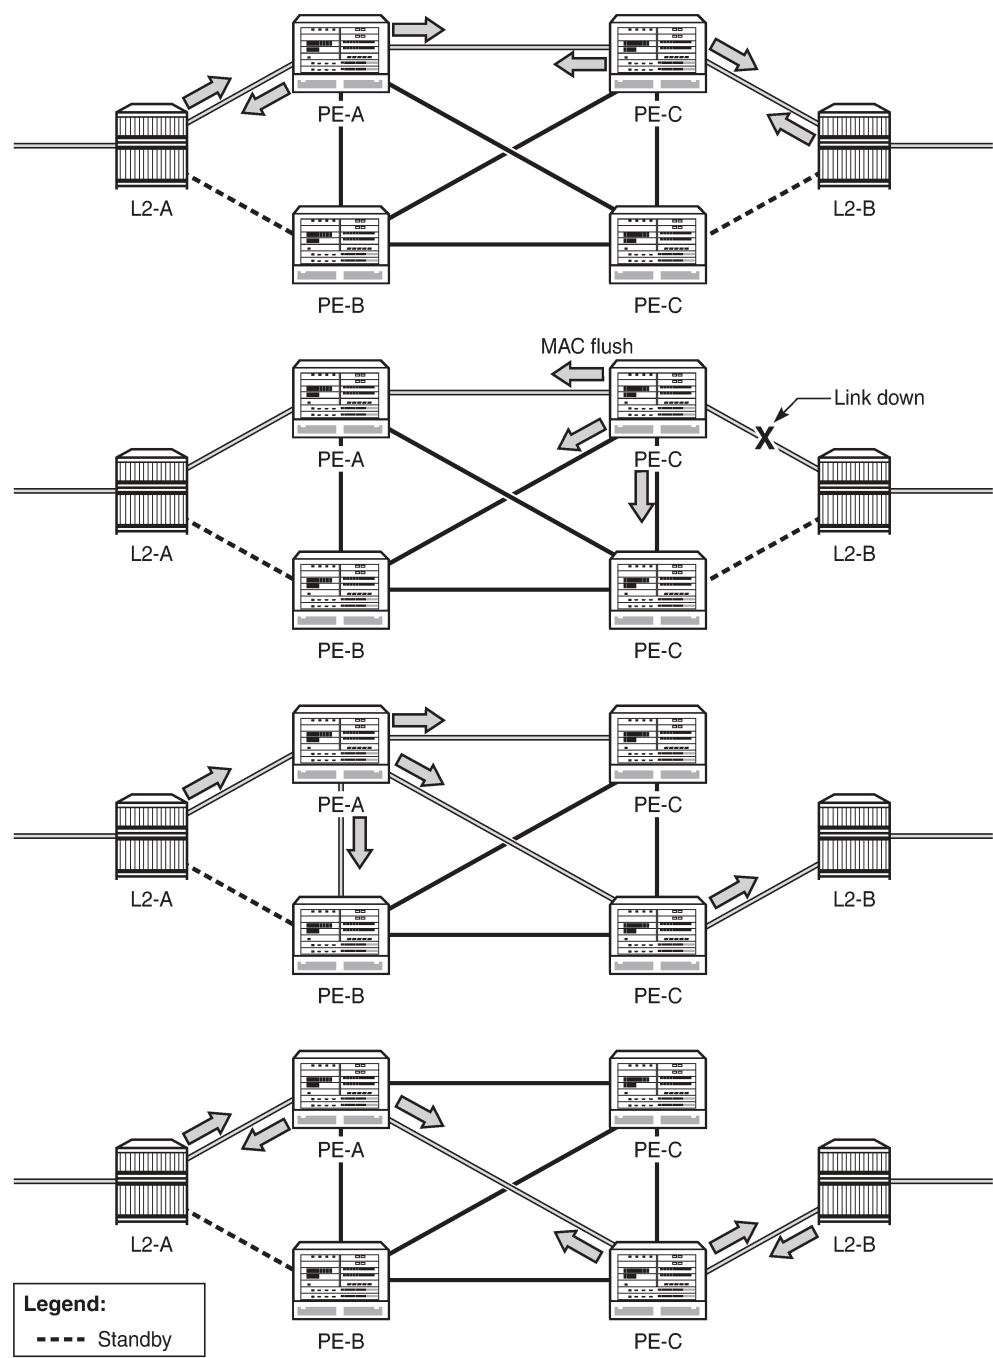
The following figure illustrates a dual-homed connection to a VPLS service (PE-A, PE-B, PE-C, PE-D) and the operation in case of link failure (between PE-C and L2-B). Upon detection of a link failure, PE-C

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, 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, 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 particular CE device opens the alternative link (L2-B switch in the figure) 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 98: Dual-homed CE connection to VPLS




25445

5.5 ATM PVC access and termination on a VPLS service

The application shown in the following figure 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 99: Example of ATM PVC access and termination on a VPLS



21558

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 SAP over an ATM interface on the 7705 SAR. The ATM-related encapsulation is stripped, and the frames (without FCS) are forwarded toward destination SAPs either locally or using SDPs associated with the VPLS service (as dictated by destination MAC address VPLS processing). In the egress direction, the received untagged frames are encapsulated into RFC 2427 or RFC 2684 (no Q-tags are added, no FCS in the forwarded frame) and sent over the ATM VC toward 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 toward the destinations 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 toward the customer CPE.

Because 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). Therefore, MAC addresses used must be unique across all the VLANs used by the customer for a specific VPLS service instance. If a customer wants a per-VLAN separation, 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.

5.6 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 \(VPLS\)](#)

5.6.1 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](#))

5.6.2 VLAN processing

The SAP encapsulation definition on Ethernet ingress ports defines which VLAN tags are used to determine the service that the packet belongs 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.

5.6.2.1 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 about tagging rules, see [Tagging rules for Epipe](#).

5.6.3 QinQ (VPLS)

VPLS supports QinQ functionality. For details, see [QinQ support](#).

5.7 Configuration notes

The following guidelines and restrictions 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

5.8 Configuring a VPLS service with CLI

This section provides information to configure VPLS services using the CLI.

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](#)
- [Configuring PIM snooping for VPLS](#)
- [Configuring a security zone within a VPLS](#)
- [Service management tasks](#)

5.9 Basic configuration

The following fields require specific input (there are no defaults) to configure a basic VPLS service:

- customer ID (see [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
  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#

```

5.10 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:
 - a. Select nodes and ports.
 - b. (optional) Select QoS policies other than the default (configured in the **config>qos** context).
 - c. (optional) Select filter policies (configured in the **config>filter** context).
 - d. (optional) Select accounting policy (configured in the **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.

5.11 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](#)

5.11.1 Creating a VPLS service

Use the following CLI syntax to create a VPLS service:

CLI syntax:

```
config>service#
vpls service-id [customer customer-id] [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#
```

5.11.2 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
...
-----
```

5.11.3 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#

```

5.11.4 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.

5.11.4.1 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

```

5.11.4.2 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

```

5.11.4.3 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

5.11.4.4 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

5.11.4.5 Hold count

The **hold-count** command configures the peak number of BPDUs that can be transmitted in 1 second.

CLI syntax:

```
config>service>vpls service-id# stp
hold-count count-value
```

Range: 1 to 10

Default: 6

Restore Default: no hold-count

5.12 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.

For information about configuring ETH-CFM parameters on VPLS (Ethernet) SAPs, see [ETH-CFM \(802.1ag and Y.1731\) tasks](#).

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](#)

5.12.1 Local VPLS SAPs

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

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#
```

5.12.2 Distributed VPLS SAPs

To configure a distributed VPLS service, you must configure service entities on originating and far-end nodes. You must use the same service ID on all ends (for example, create a VPLS service ID 9000 on 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 SDPs](#). 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
    ...
-----

```

5.12.3 Configuring SAP-specific STP parameters

When a VPLS has STP enabled, each SAP within the VPLS has STP enabled by default. The operation of STP on each SAP is governed by:

- [SAP STP administrative state](#)
- [SAP virtual port number](#)
- [SAP priority](#)
- [SAP path cost](#)
- [SAP edge port](#)
- [SAP auto edge](#)
- [SAP link type](#)

5.12.3.1 SAP STP administrative state

The administrative state of STP within a SAP controls how BPDUs are transmitted and handled when received. The allowable states are:

- **SAP admin up**

The default administrative state is up for STP on a SAP. BPDUs are handled in the normal STP manner on a SAP that is administratively up.

- **SAP admin down**

An administratively down state allows a service provider to prevent a SAP from becoming operationally blocked. BPDUs will not originate out the SAP toward the customer.

If STP is enabled on 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)
```

5.12.3.2 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.

Because 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
```

5.12.3.3 SAP priority

SAP priority allows a configurable "tiebreaking" 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

5.12.3.4 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. Because 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

5.12.3.5 SAP edge port

The SAP **edge-port** command is used to reduce the time it takes a SAP to reach the forwarding state when the SAP is on the edge of the network and therefore has no further STP bridge to handshake with.

The **edge-port** command is used to initialize the internal OPER_EDGE variable. At any time, when OPER_EDGE is false on a SAP, the normal mechanisms are used to transition to the forwarding state. 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

5.12.3.6 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

5.12.3.7 SAP link type

The SAP **link-type** command instructs STP on the maximum number of bridges behind this SAP.

If there is only a single bridge, transitioning to the forwarding state is 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

5.12.4 STP SAP operational states

The operational state of STP on a SAP controls how BPDUs are transmitted and handled when received. Defined states are:

- [Operationally disabled](#)
- [Operationally discarding](#)

- [Operationally learning](#)
- [Operationally forwarding](#)

5.12.4.1 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.

5.12.4.2 Operationally discarding

A SAP in the discarding state only receives and sends BPDUs, building the local correct STP state for each SAP while not forwarding actual user traffic.



Note: In previous versions of the STP standard, the discarding state was called a blocked state.

5.12.4.3 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.

5.12.4.4 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.

5.12.5 Configuring VPLS SAPs with split horizon

To configure a VPLS service with a split horizon group, add the **split-horizon-group** parameter when creating the SAP. Traffic arriving on a SAP within a split horizon group will not be copied to other SAPs in the same split horizon group.

The following example 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#

```

5.13 Configuring SDP bindings

This section contains the following topics:

- [Configuring mesh SDP bindings](#)
- [Configuring spoke SDPs](#)
- [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.

The following figure 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 100: SDPs – unidirectional tunnels



21561

5.13.1 Configuring mesh SDP bindings

Use the following CLI syntax to create a mesh SDP binding with a distributed VPLS service. SDPs must be configured before binding. See [Configuring SDPs](#) for information about creating SDPs. For information about configuring ETH-CFM parameters on VPLS (Ethernet) mesh SDPs, see [ETH-CFM \(802.1ag and Y.1731\) tasks](#).

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
```

5.13.2 Configuring spoke SDPs

Topics in this section include:

- [Configuring spoke SDP bindings](#)
- [Configuring spoke SDP-specific STP parameters](#)

5.13.2.1 Configuring spoke SDP bindings

Use the following CLI syntax to create a spoke SDP binding with a distributed VPLS service. SDPs must be configured before binding. See [Configuring SDPs](#) for information about creating SDPs. For information about configuring ETH-CFM parameters on VPLS (Ethernet) spoke SDPs, see [ETH-CFM \(802.1ag and Y.1731\) tasks](#).

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
    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#

```

5.13.2.2 Configuring spoke SDP-specific STP parameters

When a VPLS has STP enabled, each spoke SDP in 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](#)

5.13.2.2.1 Spoke SDP STP administrative state

The administrative state of STP within a spoke SDP controls how BPDUs are transmitted and handled when received. The allowable states are:

- **spoke SDP admin up**

The default administrative state is up for STP on a spoke SDP. BPDUs are handled in the normal STP manner on a spoke SDP that is administratively up.

- **spoke SDP admin down**

An administratively down state allows a service provider to prevent a spoke SDP from becoming operationally blocked. BPDUs will not originate out the spoke SDP toward the customer.

If STP is enabled on 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
```

5.13.2.2.2 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.

Because 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
```

5.13.2.2.3 Spoke SDP priority

Spoke SDP priority allows a configurable "tiebreaking" 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
```

5.13.2.2.4 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. Because 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
```

5.13.2.2.5 Spoke SDP edge port

The spoke SDP **edge-port** command is used to reduce the time it takes a spoke SDP to reach the forwarding state when the spoke SDP is on the edge of the network and therefore has no further STP bridge to handshake with.

The **edge-port** command is used to initialize the internal OPER_EDGE variable. At any time, when OPER_EDGE is false on a spoke SDP, the normal mechanisms are used to transition to the forwarding state. 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
```

5.13.2.2.6 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
```

5.13.2.2.7 Spoke SDP link type

The spoke SDP **link-type** command instructs STP on the maximum number of bridges behind this spoke SDP.

If there is only a single bridge, transitioning to the forwarding state is 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}
```

5.13.3 Configuring VPLS spoke SDPs with split horizon

To configure spoke SDPs with a split horizon group, add the **split-horizon-group** parameter when creating the spoke SDP. Traffic arriving on a SAP or spoke SDP within a split horizon group will not be copied to other SAPs or spoke SDPs in the same split horizon group.

The following example 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#
```

5.13.4 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
```

5.14 Configuring routed VPLS

To establish routed VPLS (r-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-routed-override-filter** 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
```

5.15 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#

```

5.16 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 **mls-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 192.168.0.0/16
      ipv6
        address 2001:db8:a::123
      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-1>config>router>pim#
```

5.17 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. See the "Filter policies" section in the 7705 SAR 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 10.0.0.0
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 max-num-grp-sources
[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#

```

5.18 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 **ml-d-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 192.0.2.0
            starg
        exit
        group 192.0.2.1
            source 192.0.2.10
            source 192.0.2.11
        exit
    exit
-----
*A:ALU-1>config>service>vpls>sap>igmp-snooping#

```

5.19 Configuring PIM snooping for VPLS

Use the **pim-snooping** command to connect a source in a Layer 2 access network to the host in a Layer 3 core network.

Use the following CLI syntax to configure PIM snooping for VPLS and to configure the maximum number of multicast groups for PIM snooping for VPLS SAPs and spoke SDPs.

CLI syntax:

```

config>service>vpls#
pim-snooping
    [no] group-policy grp-policy-name [grp-policy-name....up to 5 max]
    [no] hold-time seconds
    [no] ipv4-multicast-disable
    [no] ipv6-multicast-disable
    mode [snoop | proxy]

```

CLI syntax:

```

config>service>vpls>sap#
pim-snooping
    [no] max-num-groups max-num-groups

```

CLI syntax:

```

config>service>vpls>spoke-sdp#
pim-snooping
    [no] max-num-groups max-num-groups

```

The following displays a VPLS configuration with PIM snooping.

```

*A:ALU>config>service>vpls>pim-snooping# info detail
-----
        mode proxy
        hold-time 90
        no group-policy
        no ipv4-multicast-disable
        ipv6-multicast-disable
-----
*A:ALU>config>service>vpls>pim-snooping#

*A:ALU>config>service>vpls>sap>pim-snooping# info detail
-----
        no max-num-groups

```

```
-----
*A:ALU>config>service>vpls>sap>pim-snooping#
```

5.20 Configuring a security zone within a VPLS

To configure firewall security functionality, you must:

- configure a 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 and specify the match criteria and the action to be applied to a packet if a match is found
- configure a firewall bypass policy in the **config>security** context and assign it to the VPLS, as shown in the following CLI syntax. This step is optional. If you do not configure a bypass policy, the protocol packets are firewalled based on the firewall security policies.
- configure a security zone and apply the policy ID to the zone, as shown in the following CLI syntax

CLI syntax:

```
config>service
  vpls service-id [customer customer-id] [create]
    fw-bypass-policy {bypass-id | name}
    zone zone-id [create]
      abort
      begin
      commit
      description description-string
      inbound
      mesh-sdp sdp-id:vc-id
      outbound
      policy {policy-id | policy-name}
      sap sap-id
      shutdown
      spoke-sdp sdp-id:vc-id
```

The following example displays the security zone configuration output.

```
*A: Sar8 Dut-A>config>service>vpls# info
```

```
-----
      stp
        shutdown
      exit
      fw-bypass-policy 1
      sap 1/2/2 create
        no shutdown
      zone 1 create
        name "VPLS zone"
        description "Sample zone"
        sap "1/2/3"
        policy "4"
        inbound
          limit
        exit
      exit
      outbound
```

```

        limit
        exit
    exit
    commit
exit
no shutdown
-----

```

5.21 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](#)

5.21.1 Modifying VPLS service parameters

You can change existing service parameters. The changes are applied immediately.

To display a list of services, use the **show service service-using vpls** command. Enter the parameters, such as description, SAP, SDP, or service-MTU command syntax, 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#

```

5.21.2 Modifying management VPLS parameters

To modify the range of VLANs on an access port that are to be managed by an existing management VPLS, the new range should first 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
```

5.21.3 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
```

5.21.4 Disabling a management VPLS

You can shut down a management VPLS without deleting the service parameters. When a management VPLS is disabled, all associated user VPLS services are also disabled (to prevent loops). If this is not wanted, 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
```

5.21.5 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
```

5.21.6 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
```

5.21.7 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
```

5.22 VPLS command reference

5.22.1 Command hierarchies

- [VPLS service configuration commands](#)
 - [Global commands](#)
 - [SAP commands](#)
 - [Mesh SDP commands](#)
 - [Spoke SDP commands](#)
 - [Routed VPLS commands](#)
 - [VPLS security configuration commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

5.22.1.1 VPLS service configuration commands

5.22.1.1.1 Global commands

```

config
- service
- vpls service-id [customer customer-id] [m-vpls] [create]
- no vpls service-id
- bgp (see VPLS commands for EVPN)
- bgp-evpn (see VPLS commands for EVPN)
- 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
- load-balancing
  - [no] l4-load-balancing
  - [no] per-service-hashing
  - [no] spi-load-balancing
  - [no] teid-load-balancing
- 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] pim-snooping
  - group-policy grp-policy-name [grp-policy-name...(up to 5 max)]
  - no group-policy
  - hold-time seconds
  - no hold-time
  - [no] ipv4-multicast-disable
  - [no] ipv6-multicast-disable
  - mode mode
- [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
- static-mac (see Epipe commands for EVPN)
- 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
- system
  - fp
  - options
    - vpls-high-scale
    - [no] shutdown

```

5.22.1.1.2 SAP commands

```

config
- service
  - vpls service-id [customer customer-id] [create]
  - no vpls service-id
  - sap sap-id [split-horizon-group group-name] [eth-ring ring-index] [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] cflowd
  - [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
  - filter mac mac-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
- eth-cfm
  - mep mep-id domain md-index association ma-index [direction {up | down}]
  - no mep mep-id domain md-index association ma-index
    - [no] ais-enable
      - client-meg-level [level [level ...]]
      - no client-meg-level
      - interval {1 | 60}
      - no interval
      - priority priority-value
      - no priority
    - [no] ccm-enable
    - ccm-ltm-priority priority
    - no ccm-ltm-priority
    - description description-string
    - no description
    - [no] eth-test-enable
      - bit-error-threshold bit-errors
      - test-pattern {all-zeros | all-ones} [crc-enable]
      - no test-pattern
    - low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon |
xcon | noXcon}
      - mac-address mac-address
      - no mac-address
      - one-way-delay-threshold seconds
      - [no] shutdown
- [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 max-num-grp-sources
  - 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
- pim-snooping
- max-num-groups num-groups
- no max-num-groups
- [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
```

5.22.1.1.3 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]
- [no] entropy-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
- 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 max-num-grp-sources
- 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]

```

5.22.1.1.4 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]
      - [no] entropy-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
    - 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 max-num-grp-sources
- 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
- pim-snooping
  - max-num-groups num-groups
  - no max-num-groups
- 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]

```



Note: For information on configuring ETH-CFM on VPLS SAPs and SDPs, see the 7705 SAR OAM and Diagnostics Guide.

5.22.1.1.5 Routed VPLS commands

```

config
- service
  - vpls service-id
    - allow-ip-int-binding
    - no allow-ip-int-binding
    - service-name service-name
    - no service-name

```

5.22.1.1.6 VPLS security configuration commands

```

config
- service
  - vpls service-id
    - fw-bypass-policy {bypass-id | name}
    - no fw-bypass-policy
    - zone {zone-id | name} [create]
    - no zone {zone-id | 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}
      - log {log-id | name}
      - no log
      - [no] mesh-sdp sdp-id:vc-id
        - [no] shutdown
      - name 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] sap sap-id
  - [no] shutdown
- [no] shutdown
- [no] spoke-sdp sdp-id:vc-id
  - [no] shutdown

```

5.22.1.2 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 sap-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
  - macsec
  - 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]
  - pim-snooping
    - group [grp-ip-address] [source ip-address] [type {starg | sg}] [detail]
[family]
  - neighbor [{sap sap-id | sdp sdp-id:vc-id} [address ip-address]] [detail]
[family]
  - port [sap sap-id | sdp sdp-id:vc-id] [group [grp-ip-address]] [detail]
[family]

```

```

- statistics [sap sap-id] [sdp sdp-id:vc-id] [family]
- status [family]
- sap
- 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}
- sap-using [ingress | egress] shaper-group shaper-group-name

```

5.22.1.3 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]
  - mesh-sdp sdp-id[:vc-id] ingress-vc-label
  - 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
  - pim-snooping
    - database [[sap sap-id | sdp sdp-id:vc-id] [group grp-ip-address] [source src-
ip-address]] [family]
    - neighbor [ip-address | sap sap-id | sdp sdp-id:vc-id] [family]
    - statistics [sap sap-id | sdp sdp-id:vc-id] [family]
  - 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]

```

5.22.1.4 Debug commands

```

debug
- service
- id service-id
- [no] event-type {config-change | svc-oper-status-change | sap-oper-status-
change | sdpbind-oper-status-change}
- [no] 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] pim-snooping
- [no] adjacency
- all [group grp-ip-address] [source src-ip-address] [detail]
- no all
- database [group grp-ip-address] [source src-ip-address] [detail]
- no database
- jp [group grp-ip-address] [source src-ip-address] [detail]
- no jp
- packet [hello | jp] [sap sap-id | sdp sdp-id:vc-id]
- [no] packet
- port [sap sap-id | sdp sdp-id:vc-id] [detail]
- no port
- red [detail]
- no red
- [no] 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

```

See the 7705 SAR OAM and Diagnostics Guide for information about CLI commands and syntax for OAM and diagnostics commands.

5.22.2 Command descriptions

- [VPLS service configuration commands](#)
- [Routed VPLS configuration commands](#)
- [VPLS security configuration commands](#)
- [VPLS show commands](#)
- [VPLS clear commands](#)
- [VPLS debug commands](#)

5.22.2.1 VPLS service configuration commands

- [Generic commands](#)
- [VPLS service commands](#)
- [VPLS SAP commands](#)
- [VPLS SAP ATM commands](#)
- [VPLS IGMP and MLD snooping commands](#)
- [VPLS STP commands](#)
- [VPLS filter and QoS policy commands](#)
- [Service billing commands](#)
- [VPLS SAP DHCP commands](#)
- [VPLS SDP commands](#)

5.22.2.1.1 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
config>service>vpls>zone
config>service>vpls>zone>mesh-sdp
config>service>vpls>zone>sap
config>service>vpls>zone>spoke-sdp
```

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
config>service>vpls>zone

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 (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

5.22.2.1.2 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 service (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.

When a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified will result in an error.

More than one VPLS service may be created for a single customer ID.

By default, no VPLS instances exist until they are explicitly created.

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs and SDPs defined within the service ID have been 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 ignore 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 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

pim-snooping

Syntax

[no] **pim-snooping**

Context

config>service>vpls

config>service>vpls>sap

config>service>vpls>spoke-sdp

Description

This command enables PIM snooping for the VPLS service. When enabled, it is enabled for all SAPs except default SAPs. A default SAP is a SAP that has a wildcard VLAN ID, such as sap 1/1/1:*.

The **no** form of the command disables PIM snooping and removes the PIM snooping configuration.

group-policy

Syntax

group-policy *grp-policy-name* [*grp-policy-name...*(up to 5 max)]

no group-policy

Context

config>service>vpls>pim-snooping

Description

This command identifies one or more route policies for multicast groups applied to this VPLS entity. The sources of the multicast traffic must be members of the VPLS.

Routing policies are configured in the **config>router>policy-options** context. The router policy must be defined before it can be imported. The *grp-policy-name* variable is the same as the *name* variable in the **policy-options>policy-statement** *name* command.

For details on route policies, see the "Route Policies" section in the 7705 SAR Router Configuration Guide.

The **no** form of the command removes the policy association from the VPLS configuration.

Default

n/a

Parameters

grp-policy-name

the group policy name

hold-time

Syntax

hold-time *seconds*

no hold-time

Context

config>service>vpls>pim-snooping

Description

This command configures the length of time during which the PIM snooping switch snoops all the PIM states in the VPLS. During the **hold-time**, multicast traffic is flooded in the VPLS. At the end of the **hold-time**, multicast traffic is forwarded using the snooped states. The snooped state consists of the forwarding state for the (S,G) and (*,G) groups, the incoming interface, and the outgoing interfaces.

When PIM snooping is enabled in VPLS, there is a period of time when the PIM snooping switch may not have built the complete snooping state. The switch cannot build states until the routers connected to the VPLS refresh their PIM messages.

This parameter is applicable only when PIM snooping is enabled.

Parameters

seconds

specifies the PIM snooping hold time, in seconds

Values 0 to 300

Default 90

ipv4-multicast-disable

Syntax

[no] ipv4-multicast-disable

Context

config>service>vpls>pim-snooping

Description

This command disables PIM snooping for IPv4 multicast traffic within a VPLS service. By default, IPv4 multicast traffic is enabled.

The **no** form of the command enables PIM snooping for IPv4 multicast traffic within a VPLS service. To fully remove PIM snooping from a VPLS service, the **no pim-snooping** command must be used.

Default

enabled (no ipv6-multicast-disable)

ipv6-multicast-disable

Syntax

[no] ipv6-multicast-disable

Context

config>service>vpls>pim-snooping

Description

This command disables PIM snooping for IPv6 multicast traffic within a VPLS service. By default, IPv6 multicast traffic is disabled.

The **no** form of the command enables PIM snooping for IPv6 multicast traffic within a VPLS service. To fully remove PIM snooping from a VPLS service, the **no pim-snooping** command must be used.

Default

disabled (ipv6-multicast-disable)

mode

Syntax

mode mode

Context

config>service>vpls>pim-snooping

Description

This command sets the PIM snooping mode to proxy or plain snooping.

Parameters

mode	
	specifies PIM snooping mode
Values	snoop, proxy
Default	proxy

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

load-balancing

Syntax

```
load-balancing
```

Context

```
config>service>vpls
```

Description

This command accesses the context to configure load balancing.

I4-load-balancing

Syntax

[no] **I4-load-balancing**

Context

config>service>vpls>load-balancing

Description

This command enables or disables Layer 4 load balancing for the VPLS instance. 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.

Adding the Layer 4 source and destination port fields to the hashing algorithm generates a higher degree of randomness and a more even distribution of packets across the available LAG ports.

You can add additional fields to generate more randomness and more equal distribution of packets with the [teid-load-balancing](#) command.

Hashing based on the **I4-load-balancing** and **teid-load-balancing** commands and hashing based on the **per-service-hashing** command are mutually exclusive.

The **no** form of the command disables Layer 4 load balancing.

Default

no I4-load-balancing

per-service-hashing

Syntax

[no] **per-service-hashing**

Context

config>service>vpls>load-balancing

Description

This command enables or disables hashing based on the service ID. The result of the hashing calculation is used to determine the distribution of packets.

Hashing based on the **per-service-hashing** command and hashing based on the **I4-load-balancing** and **teid-load-balancing** commands are mutually exclusive.



Note: If **per-service-hashing** is not enabled, a 4-byte hash value is appended to internal overhead for VPLS multicast traffic at ingress. The egress internal hash value is discarded at egress before scheduling. Therefore, shaping rates at access and network ingress and for fabric policies may need to be adjusted accordingly. In addition, the 4-byte internal hash value may be included in any affected statistics counters.

The **no** form of the command disables per-service hashing.

Default

no per-service-hashing

spi-load-balancing

Syntax

[no] spi-load-balancing

Context

config>service>vpls>load-balancing

Description

This command enables use of the SPI in hashing for ESP/AH encrypted IPv4 or IPv6 traffic at the interface level.

The **no** form of this command disables SPI hashing.

Default

no spi-load-balancing

teid-load-balancing

Syntax

[no] teid-load-balancing

Context

config>service>vpls>load-balancing

Description

This command enables or disables TEID load balancing for the VPLS instance. 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.

Hashing based on the **teid-load-balancing** and **l4-load-balancing** commands and hashing based on the **per-service-hashing** command are mutually exclusive.

The **no** form of the command disables TEID load balancing.

Default

no teid-load-balancing

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 configured 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 44: SAP ID configurations](#) 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

The following table shows MTU values for specific VC types.

Table 67: Service service-MTU field descriptions

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

5.22.2.1.3 VPLS SAP commands

sap

Syntax

sap *sap-id* [**split-horizon-group** *group-name*] [**eth-ring** *ring-index*] [**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 split-horizon group can be associated with an Ethernet ring to prevent loops in cases where an Ethernet virtual ring is misconfigured on the main ring. Each A and B path in the major ring is configured in the group and associated with the sub-ring control instance in the VPLS service.

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 or on Ethernet ports. 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 44: SAP ID configurations](#) for a full list of SAP IDs.

group-name

specifies the name of the split horizon group to which the SAP belongs

ring-index

specifies the ring index of the Ethernet ring

create

keyword used to create a SAP instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context

cflowd

Syntax

[no] cflowd

Context

config>service>vpls>sap

Description

This command enables cflowd to collect traffic flow samples through a SAP for analysis. When cflowd is enabled on a VPLS service SAP, the Ethernet traffic can be sampled and processed by the system's cflowd engine and exported to cflowd version 10 collectors with the **l2-ip** template enabled.

When cflowd is enabled at the SAP level, all packets forwarded by the interface are subject to analysis according to the cflowd configuration.

For Layer 2 services, only ingress sampling is supported.

Default

no cflowd

discard-unknown-source

Syntax

[no] discard-unknown-source

Context

config>service>vpls>sap

config>service>vpls>spoke-sdp

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 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 "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 "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

5.22.2.1.4 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 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 "ATM QoS Traffic Descriptor Profiles" in the 7705 SAR Quality of Service Guide for information on the **atm-td-profile traffic-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)

5.22.2.1.5 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 (such as #, \$, or spaces), 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>sap>pim-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
config>service>vpls>spoke-sdp>pim-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, MLD, or PIM join message that would exceed the configured number of groups, the request is ignored. The **max-num-groups** command does not apply to mesh SDPs.

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 512

max-num-grp-sources**Syntax**

max-num-grp-sources *max-num-grp-sources*

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

max-num-grp-sources

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

ipv6-address

the IPv6 address used for the MLD query

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

ipv6-address

the source IPv6 source address in transmitted MLD reports

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 IPv4 address must be unique for each static group.

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.

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

src-ipv6-address
specifies the IPv6 unicast address

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

5.22.2.1.6 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 Nokia 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 Nokia 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 tiebreaking 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

5.22.2.1.7 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) Ethernet adapter card, the scheduler mode must be set to 16-priority. On a third-generation (Gen-3) Ethernet adapter card, the scheduler mode is always 4-priority. For information on adapter card generations, see the "Evolution of Ethernet Adapter Cards, Modules, and Platforms" section in the 7705 SAR 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**.

If the specified SAP is a LAG SAP, *agg-rate* and *cir-rate* can be configured regardless of the scheduler mode setting on Gen-2 or Gen-3 hardware. If the active port is on a Gen-3 card or platform, *agg-rate* and *cir-rate* are applicable. If the active port is on a Gen-2 card or platform, *agg-rate* and *cir-rate* apply when the scheduler mode is set to 16-priority. For details on the behavior of a mix-and-match LAG SAP, see the "LAG Support on Third-Generation Ethernet Adapter Cards, Ports, and Platforms" and "Network LAG Traffic Management" sections in the 7705 SAR Interface Configuration Guide.

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). Filters can be applied to VPLS and routed VPLS SAPs and SDPs.

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.

If an IP interface is attached to a VPLS service context (routed VPLS), the VPLS SAP or SDP configured IP or MAC 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, using the **v4-routed-override-filter** or **v6-routed-override-filter** command in the IES or VPRN service contexts.

If filter override is configured, the IP or MAC filter configured on the SAP or SDP applies to non-routed packets. If filter override is not configured, the IP or MAC filter configured on the SAP or SDP applies to both routed and non-routed 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.

The following table 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-dot1p 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.

Table 68: 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

Port/SAP type	Match-qinq-dot1p setting ¹	Existing packet tags	P-bits used for match
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.

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.

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).

The following table 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 69: 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
Dot1q/ 0 ²	n/a	Outer tag
QinQ/ X.Y ¹	False	Two outer tags ³
	True	Outer tag ³
QinQ/ X.* ¹	True or False	Outer tag

Egress port type/SAP type	qinq-mark-top-only state	Egress P-bits marked or re-marked
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.

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, which only support 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.

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

5.22.2.1.8 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

5.22.2.1.9 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 (toward 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 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 (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 *interface-name*, separated by "|". If no keyword is configured, 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 the *vlan-id* and dot1p bits in addition to the **ascii-tuple** information. The format is supported on dot1q and qinq encapsulated ports only. 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

5.22.2.1.10 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 destination 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-mps*, 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-mps*, 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 resets 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

entropy-label

Syntax

[no] **entropy-label**

Context

config>service>vpls>mesh-sdp

config>service>vpls>spoke-sdp

Description

This command enables or disables the use of entropy labels for mesh or spoke SDPs.

If **entropy-label** is enabled, the entropy label and entropy label indicator (ELI) are inserted in packets for which at least one LSP in the stack for the far end of the tunnel used by the service has advertised entropy label capability.

If the tunnel is an RSVP-TE type, **entropy-label** can also be controlled by disabling **entropy-label-capability** under the **config>router>rsvp** context at the far-end LER.

When the **entropy-label** and **entropy-label-capability** commands are both enabled, the entropy label value inserted at the iLER is always based on the service ID.

Default

no entropy-label

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

5.22.2.2 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:

- residential SHG
- DHCP
- mVPLS
- mac-subnet-length
- GRE SDP (cannot be bound to the VPLS)



Note: The DHCP relay functionality under **config>service>ies>if>dhcp** or **config>service>ies>if>ipv6>dhcp6-relay** 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).

5.22.2.3 VPLS security configuration commands

fw-bypass-policy

Syntax

fw-bypass-policy {*bypass-id* | *name*}
no fw-bypass-policy

Context

config>service>vpls

Description

This command associates a bypass policy with this VPLS instance. The bypass policy must already be configured in the **config>security>bypass** context. All protocols, or protocols associated with specific source or destination ports, defined by the bypass policy bypass the firewall lookup table and are permitted across the zone associated with this VPLS instance without being firewalled.

Default

no fw-bypass-policy

Parameters

bypass-id
the firewall bypass ID number
Values 1 to 65535

name
the name of the firewall bypass policy
Values 1 to 32 characters

zone

Syntax

zone {*zone-id* | *name*} [**create**]
no zone {*zone-id* | *name*}

Context

config>service>vpls

Description

This command creates a security zone within a VPLS context. Each zone must have a unique ID. When a zone is created with a name, the system automatically assigns it the first available zone ID value. A zone cannot be configured on a VPLS service with EVPN.

The **no** form of this command deletes the zone. When a zone is deleted, all configuration parameters for the zone are also deleted.

Default

no zone

Parameters

zone-id

the zone ID number.

Values 1 to 65534

name

the name of the zone.

Values 1 to 32 characters (must start with a letter). If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

create

keyword required when first creating the security zone. When the zone is created, you can enter the context without the **create** keyword.

abort

Syntax

abort

Context

config>service>vpls>zone

Description

This command discards changes made to a security feature.

Default

n/a

begin

Syntax

begin

Context

config>service>vpls>zone

Description

This command enters the mode to create or edit security features.

Default

n/a

commit

Syntax

commit

Context

config>service>vpls>zone

Description

This command saves changes made to security features.

Default

n/a

inbound

Syntax

inbound

Context

config>service>vpls>zone

Description

This command enables the context to configure limit parameters for inbound firewall sessions.

Default

n/a

outbound**Syntax**

outbound

Context

config>service>vpls>zone

Description

This command enables the context to configure limit parameters for outbound firewall sessions.

Default

n/a

limit**Syntax**

limit

Context

config>service>vpls>zone>inbound

config>service>vpls>zone>outbound

Description

This command enables the context to configure limits on concurrent sessions for inbound or outbound firewall sessions.

Default

n/a

concurrent-sessions**Syntax**

concurrent-sessions {tcp | udp | icmp | other} *sessions*

no concurrent-sessions {tcp | udp | icmp | other}

Context

config>service>vpls>zone>inbound>limit

```
config>service>vpls>zone>outbound>limit
```

Description

This command configures the maximum number of concurrent firewall sessions that can be established per zone, for the specified protocol, 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 specified direction and protocol
Values	1 to 16383

log

Syntax

```
log {log-id | name}  
no log
```

Context

```
config>service>vpls>zone
```

Description

This command applies a security log to the specified zone. The security log must already be configured in the **config>security>logging** context.

The **no** form of this command removes logging for the zone.

Default

n/a

Parameters

<i>log-id</i>	the identifier for the log
Values	1 to 100
<i>name</i>	the name of the log
Values	1 to 32 characters

mesh-sdp

Syntax

[no] mesh-sdp sdp-id:vc-id

Context

config>service>vpls>zone

Description

This command assigns a mesh SDP to the security zone.
The **no** form of the command removes the mesh SDP from the zone.

Default

n/a

Parameters

<i>sdp-id</i>	the SDP identifier
Values	1 to 17407
<i>vc-id</i>	identifies the virtual circuit
Values	1 to 4294967295

name

Syntax

name name
no name

Context

config>service>vpls>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. If the zone name was already configured with the [zone](#) command, this command renames the zone.

Default

n/a

Parameters

<i>name</i>	the name of the zone
Values	1 to 32 characters (must start with a letter). If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

policy

Syntax

policy {*policy-id* | *name*}

no policy

Context

config>service>vpls>zone

Description

This command specifies the policy to be used by the security zone to build its matching criteria for incoming packets. The policy must already be configured in the **config>security** context.

The **no** form of this command deletes the specified policy.

Default

n/a

Parameters

<i>policy-id</i>	the number of the referenced policy
Values	1 to 65535
<i>name</i>	the name of the referenced policy

sap

Syntax

[no] **sap** *sap-id*

Context

config>service>vpls>zone

Description

This command assigns a SAP to the security zone.
The **no** form of this command removes the SAP from the zone.

Default

n/a

Parameters

sap-id
specifies the physical port identifier portion of the SAP definition. See [Table 44: SAP ID configurations](#) for a full list of SAP IDs.

spoke-sdp

Syntax

[no] **spoke-sdp** *sdp-id:vc-id*

Context

config>service>vpls>zone

Description

This command assigns a spoke SDP to the security zone.
The **no** form of this command removes the spoke SDP from the zone.

Default

n/a

Parameters

sdp-id
uniquely identifies the SDP
Values 1 to 17407

vc-id
identifies the virtual circuit

Values 1 to 4294967295

5.22.2.4 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 70: Service egress labels field descriptions](#) describes the fields.

Output example

```
*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
=====

```

Table 70: Service egress labels field descriptions

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 71: Service FDB-info field descriptions](#) describes the fields.

Output example

```

*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 71: Service FDB-info field descriptions

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 to protect the system against undetected loops or duplicate MACs.

Label	Description
	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 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
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

Label	Description
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 72: Service FDB-MAC field descriptions](#) describes the fields.

Output example

```
*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 72: Service FDB-MAC field descriptions

Label	Description
ServId	The service ID number
MAC	The specified MAC address
Source-Identifier	The location where the MAC is defined
Type/Age	Static: FDB entries created by management
	Learned: dynamic entries created by the learning process
	OAM: entries created by the OAM process
	H: host, the entry added by the system for a static configured subscriber host
	D or DHCP: DHCP-installed MAC Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease
	P: indicates the MAC is protected by the MAC protection feature
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](#))

macsec

display MACsec related information about the service ([macsec](#))

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 73: Service service-ID \(all\) field descriptions](#) describes the fields.

Output example

```
*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
MTU              : 1514
SAP Count        : 1
Snd Flush on Fail: Disabled
Propagate MacFlush: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None
Oper State       : Down
Def. Mesh VC Id  : 5001
SDP Bind Count   : 1
Host Conn Verify : Disabled

-----
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
Split Horiz Grp  : shg5001
VC Type         : Ether
Admin Path MTU   : 0
Far End         : 10.10.10.10
Type            : Spoke
VC Tag          : n/a
Oper Path MTU   : 0
Delivery        : MPLS
Admin State      : Up
Acct. Pol       : None
Ingress Label    : 0
Ing mac Fltr     : n/a
Ing ip Fltr      : n/a
Ing ipv6 Fltr    : n/a
Admin ControlWord : Not Preferred
Last Status Change: 10/26/2010 20:13:08
Last Mgmt Change : 10/26/2010 20:13:09
Endpoint        : endpoint5000
PW Status Sig    : Enabled
Class Fwding State : Down
Oper State       : Down
Collect Stats    : Disabled
Egress Label     : 0
Egr mac Fltr     : n/a
Egr ip Fltr      : n/a
Egr ipv6 Fltr    : n/a
Oper ControlWord : False
Signaling        : TLDP
Force Vlan-Vc    : Disabled
Precedence       : 4
```

```

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 Scheduler Mode	: 4-priority	Egr Scheduler Mode	: 4-priority
Ing Agg Rate Limit	: 999000	Egr Agg Rate Limit	: max
Ing Agg cir	: 333000	Egr Agg cir	: 0
Ing Shaper Group	: test_sgl	Egr Shaper Group	: default
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	BPDUs Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: N/A
Last BPDUs 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 (Ingress)		
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

Forwarding Engine Stats (Egress)		
Dropped	: 0	n/a


```

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

Ingress Queue 3 (Profile)
Off. ColorIn     : 0          0
Off. ColorOut    : 0          0
Off. Uncolor     : 0          0
Dro. ColorOut    : 0          0
Dro. ColorIn/Uncolor : 0      0
For. InProf      : 0          0
For. OutProf     : 0          0
-----
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
Primary Factor   : 3
Mac Move Rate    : 2
Mac Move Retries : 3
Table Size      : 250
Learned Count    : 0
OAM-learned Count : 0
Host-learned Count : 0
Remote Age      : 900
High Watermark   : 95%
Mac Learning     : Enabled
Mac Aging        : Enabled
Mac Subnet Len   : 48
Incl PPP Circ-Id : no

Mac Move        : Disabled
Secondary Factor : 2
Mac Move Timeout : 10

Total Count     : 0
Static Count    : 0
DHCP-learned Count : 0

Local Age       : 300
Low Watermark   : 90%
Discard Unknown : Disabled
Relearn Only    : False

```

IGMP Snooping Base info

```

Admin State : Down
Querier     : No querier found

```

Sap/Sdp Id	Oper State	MRtr Port	Pim Port	Send Queries	Max Grps	Max Srcs	MVR From-VPLS	Num 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 Id	Oper State	MRtr Port	Send Queries	Max Groups	Num Groups	MVR From-VPLS	Num Groups
sap:1/2/4:1/100	Down	No	Disabled	No Limit	Local	Local	0
sdp:5001:100	Down	No	Disabled	No Limit	N/A	N/A	0

DHCP Summary, service 5001

Sap/Sdp	Snoop	Used/Provided	Arp Reply Agent	Info Option	Admin 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 73: Service service-ID (all) field descriptions

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

Label	Description
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
Service Destination Points (SDPs) (see Table 90: Service ID (SDP) field descriptions)	
Service Access Points (see Table 89: Service ID (SAP) field descriptions)	
IGMP Snooping Base info (see Table 80: Service service-ID IGMP and MLD snooping (base) field descriptions)	
MLD Snooping Base info (see Table 80: Service service-ID IGMP and MLD snooping (base) field descriptions)	
VPLS Spanning Tree Information (see Table 92: Service ID (STP) field descriptions)	
DHCP Summary (see Table 76: Service service-ID (DHCP summary) field descriptions)	
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

Label	Description
Service Endpoints (see Table 77: Service service-ID (endpoint) field descriptions)	

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 74: Service service-ID \(base\) field descriptions](#) describes the fields.

Output example

```
*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
MTU              : 1514
SAP Count        : 1
Snd Flush on Fail : Disabled
Propagate MacFlush: Disabled
Def. Gateway IP   : None
Def. Gateway MAC  : None
Oper State       : Down
Def. Mesh VC Id  : 5001
SDP Bind Count   : 1
Host Conn Verify : Disabled

-----
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 74: Service service-ID (base) field descriptions

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
Service Access and Destination Points	
Identifier	Specifies the service access (SAP) and destination (SDP) points

Label	Description
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 44: SAP ID configurations](#) for a full list of SAP IDs.
- sdp-id

the SDP identifier

Values1 to 17407
- vc-id

the virtual circuit ID on the SDP ID for which to display information

Values1 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

Output

The following output is an example of DHCP statistics information for a SAP, an SDP, and an interface, and [Table 75: Service service-ID \(DHCP statistics\) field descriptions](#) describes the fields.

Output example

```
*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 75: Service service-ID (DHCP statistics) field descriptions

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)

Label	Description
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 76: Service service-ID \(DHCP summary\) field descriptions](#) describes the fields.

Output example

```
*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 76: Service service-ID (DHCP summary) field descriptions

Label	Description
Sap/Sdp	The SAP or SDP identifier of the router interface
Snoop	Specifies whether 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 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 77: Service service-ID \(endpoint\) field descriptions](#) describes the fields.

Output example

```

*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 77: Service service-ID (endpoint) field descriptions

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 ignore standby signaling is configured True: standby signaling is ignored False: standby signaling is not ignored
Suppress Standby Signaling	Specifies whether suppress standby signaling is configured True: standby signaling is suppressed False: standby signaling is not suppressed
Block On Mesh Fail	Specifies whether 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

Label	Description
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
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 44: SAP ID configurations](#) for a full list of SAP IDs.

sdp-id

specifies the SDP ID for the display

Values 1 to 17407

ieee-address

specifies the 48-bit MAC address

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 78: Service service-ID \(FDB\) field descriptions](#) describes the fields.

Output example

```
*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 78: Service service-ID (FDB) field descriptions

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	<p>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 to protect the system against undetected loops or duplicate MACs</p> <p>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</p>

Label	Description
Mac Move Timeout	<p>The time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled.</p> <p>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 to avoid thrashing.</p>
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.
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

Label	Description
	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 79: Service service-ID \(IGMP and MLD snooping\) field descriptions](#) describes the fields.

Output example (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      : 10.12.1.1 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       : 10.12.1.1
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
-----
10.12.1.1        1/9/7:34          0d 00:00:38      63487s          3
-----
Number of mrouter: 1
-----
IGMP Snooping Proxy-reporting DB
-----
Group Address    Mode      Up Time          Num Sources
-----
10.0.0.1         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              Stdbby
-----
10.0.0.1         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              Stdbby
-----
-----
Number of groups: 0
-----
IGMP Snooping SDP 21:1 Port-DB
-----
Group Address    Mode      Type      From-VPLS  Up Time          Expires          Num Src
-----
10.0.0.1         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
-----
10.0.0.1         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			
10.12.0.12	10.0.0.1			
10.12.0.13	10.0.0.1			
10.12.0.14	10.0.0.1			
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 Id	Oper State	From VPLS	Num Local 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 Id	Oper State	From VPLS	Num MVR Groups
No MVR SAPs found.				

```
=====
*A:7705:Dut-C#
=====
```

Output example (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 Id	Oper State	MRtr Port	Send Queries	Max Num Groups	MVR From-VPLS	Num 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 Src	MC Stdbby
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
                  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
                  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 Id	Oper State	From VPLS	Num Local 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 Id	Oper State	From VPLS	Num MVR Groups
--------	------------	------------	-----------	----------------

No MVR SAPs found.

=====

*A:7705:Dut-C#

Table 79: Service service-ID (IGMP and MLD snooping) field descriptions

Label	Description
IGMP Snooping info for service <i>id</i> MLD Snooping info for service <i>id</i>	Shows the service <i>id</i> 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)

Label	Description
	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
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	

Label	Description
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
IGMP Snooping Proxy-reporting DB	
MLD Snooping Proxy-reporting DB	
Group Address	The IP multicast group address for which this entry contains information
Mode	<p>The type of membership reports received on the interface for the group: include or exclude</p> <p>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</p> <p>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</p>
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 Port-DB	
MLD Snooping SAP Port-DB	
Group Address	The IP multicast group address for which this entry contains information
Mode	<p>The type of membership reports received on the interface for the group: include or exclude</p> <p>Include – reception of packets sent to the specified multicast address is requested only from those IP source</p>

Label	Description
	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 SAP
MC Stdbby	The multicast standby state
Number of groups	The number of groups configured for this SAP
IGMP Snooping SDP Port-DB MLD Snooping SDP 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

Label	Description
	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
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	
IGMP Snooping Static Source Groups for SDP	
MLD Snooping Static Source Groups for SAP	
MLD Snooping Static Source Groups for SDP	
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

Label	Description
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
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

Label	Description
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)
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	
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

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.
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 80: Service service-ID IGMP and MLD snooping \(base\) field descriptions](#) describes the fields.

Output example

```
*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:1/10/6   Down  No    No    0    0    None Local  0
sdp:5007:5007 Down  Yes   No    0    0    None N/A   0
=====
*A:7705custDoc:Sar18>show>service>id# mld-snooping
```

```

-----
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 80: Service service-ID IGMP and MLD snooping (base) field descriptions

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 IPv4 address
grp-ipv6-address	displays the MLD snooping port database for the specified multicast group IPv6 address
sap-id	displays the IGMP or MLD snooping port database for the specified SAP. See Table 44: SAP ID configurations 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. <div>Values 1 to 17407</div>
vc-id	the virtual circuit ID on the SDP ID for which to display information <div>Values 1 to 4294967295</div>

Output

The following outputs are examples of port database information for IGMP snooping and MLD snooping displays for the specified VPLS service, and [Table 81: Service service-ID IGMP and MLD snooping \(port-DB\) field descriptions](#) describes the fields.

Output example (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
-----
10.0.0.2         exclude static local      29d 01:28:14  never      0
10.0.0.3         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 10.0.0.2
-----
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 10.0.0.3
-----
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
-----
10.12.12.12     29d 01:27:38  never         static        Fwd
10.12.13.13     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
-----
10.0.0.1         include dynamic N/A      0d 00:10:48  0s          2
-----
Number of groups: 1
=====
```



```
*A:7705:Dut-C#
```

Output example (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
      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
Vl 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 81: Service service-ID IGMP and MLD snooping (port-DB) field descriptions

Label	Description
Group Address	The IP multicast group address for which this entry contains information
Mode	<p>The type of membership reports received on the interface for the group: include or exclude</p> <p>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</p> <p>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</p>
Type	<p>Indicates how this group entry was learned: dynamic or static</p> <p>Dynamic – this group entry was learned by IGMP or MLD</p> <p>Static – this group entry was statically configured</p>
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
MC Stdbby	<p>(SAP only)</p> <p>Indicates the multicast standby state</p>
Number of groups	Indicates the number of groups configured for this SAP or SDP

Label	Description
IGMP Group MLD Group	The IPv4 or IPv6 multicast group address for which this entry contains information
Compat Mode	<p>The IGMP or MLD compatibility mode.</p> <p>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.</p> <p>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).</p> <p>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.</p>
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
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 IPv4 address
- grp-ipv6-address*
displays the MLD snooping proxy reporting database for the specified multicast group IPv6 address

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 82: Service service-ID IGMP and MLD snooping \(proxy-DB\) field descriptions](#) describes the fields.

Output example

```
*A:7705:Dut-C# show service id 1 igmp-snooping proxy-db detail

=====
IGMP Snooping Proxy-reporting DB for service 1
=====
-----
IGMP Group 10.0.0.1
-----
Up Time : 0d 00:03:21           Mode : exclude
-----
Source Address  Up Time
-----
No sources.
-----
Number of groups: 1
=====
```

```
*A:7705:Dut-C#
```

```
*A:7705:Sal18>show>service>id>igmp-snooping# proxy-db detail
```

```
=====
IGMP Snooping Proxy-reporting DB for service 5007
=====
```

```
-----
IGMP Group 10.0.0.2
```

```
-----
Up Time : 28d 01:13:36           Mode : exclude
```

```
-----
Source Address  Up Time
```

```
-----
No sources.
```

```
-----
IGMP Group 10.0.0.3
```

```
-----
Up Time : 28d 01:12:56           Mode : include
```

```
-----
Source Address  Up Time
```

```
-----
10.12.12.12     28d 01:12:56
```

```
10.12.13.13     28d 01:07:06
```

```
-----
Number of groups: 2
```

```
*A:7705:Sal18>show>service>id>igmp-snooping# proxy-db group 10.0.0.3
```

```
=====
IGMP Snooping Proxy-reporting DB for service 5007
=====
```

```
-----
IGMP Group 10.0.0.3
```

```
-----
Up Time : 28d 01:25:31           Mode : include
```

```
-----
Source Address  Up Time
```

```
-----
10.12.12.12     28d 01:25:31
```

```
10.12.13.13     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 82: Service service-ID IGMP and MLD snooping (proxy-DB) field descriptions

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 83: Service service-ID IGMP and MLD snooping \(querier\) field descriptions](#) describes the fields.

Output example

```
*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 83: Service service-ID IGMP and MLD snooping (querier) field descriptions

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 44: SAP ID configurations](#) 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 84: Service service-ID IGMP and MLD snooping \(static\) field descriptions](#) describes the fields.

Output example

```
*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
-----
*               10.0.0.2
10.12.12.12     10.0.0.3
10.12.13.13     10.0.0.3
-----
Static (*,G)/(S,G) entries: 3
-----
IGMP Snooping Static Source Groups for SDP 5007:5007
-----
Source          Group
-----
10.12.0.12      10.0.0.1
10.12.0.13      10.0.0.1
10.12.0.14      10.0.0.1
-----
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

*	10.0.0.2
10.12.12.12	10.0.0.3
10.12.13.13	10.0.0.3

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

10.10.0.12	10.0.0.1
10.10.0.13	10.0.0.1
10.12.0.14	10.0.0.1

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 84: Service service-ID IGMP and MLD snooping (static) field descriptions

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 44: SAP ID configurations](#) 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 85: Service service-ID IGMP and MLD snooping \(statistics\) field descriptions](#) describes the fields.

Output example

```
*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 85: Service service-ID IGMP and MLD snooping (statistics) field descriptions

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

Label	Description
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 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

Label	Description
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
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 86: Service service-ID \(labels\) field descriptions](#) describes the fields.

Output example

```
*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 86: Service service-ID (labels) field descriptions

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 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 87: Service service-ID \(MAC move\) field descriptions](#) describes the fields.

Output example

```
*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 87: Service service-ID (MAC move) field descriptions

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	<p>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 to protect the system against undetected loops or duplicate MACs.</p> <p>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.</p>
Mac Move Timeout	<p>The time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled.</p> <p>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 to avoid thrashing.</p>
Mac Move Retries	The number of times retries are performed for re-enabling the SAP or SDP
SAP Mac Move Information:	

Label	Description
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: Service AdminDown, PortOperDown, L2OperDown.
Time to RetryReset	<p>The time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled.</p> <p>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 to avoid thrashing.</p>
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: Svc AdminDown, SdpOperDown, NoIngVCLLabel, No EgrVCLLabel, PathMTUTooSmall.
Time to RetryReset	<p>The time, in seconds, to wait before a SDP that has been disabled after exceeding the maximum relearn rate is re-enabled.</p> <p>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 to avoid thrashing.</p>
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

Label	Description
Blockable Level	Specifies the level at which MAC move is blockable on the SDP (primary, secondary, or tertiary)

macsec

Syntax

macsec

Context

show>service>id

Description

This command displays MACsec security information for the specified service.

Output

The following output is an example of MACsec information, and [Table 88: Service-ID MACsec field descriptions](#) describes the fields.

Output example

```
*A:ALU-12# show service id 1 macsec
=====
MACsec (Summary), Service 1
=====
SAP          MACsec    MACsec    Encap    CA tags   CA-name
              port      sub-port  match    in-clear
-----
1/1/3        1/1/3      1         all      0         cal
=====
*A:ALU-12#
```

Table 88: Service-ID MACsec field descriptions

Label	Description
SAP	The service SAP
MACsec port	The port enabled for MACsec
MACsec sub-port	The subport enabled for MACsec
Encap match	The traffic encapsulation type to match: all traffic, untagged-only traffic, single-tag or dot1q traffic, double-tag or QinQ traffic
CA tags in-clear	The number of tags in clear text for this CA
CA-name	The name of the MACsec connectivity association for this SAP

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 44: SAP ID configurations](#) 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 89: Service ID \(SAP\) field descriptions](#) describes the fields.

Output example

```
*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
Ing Shaper Group   : default            Egr Shaper Group  : default
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
Calling-Station-Id : n/a                Tot Static Hosts  : 0
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	BPDUs Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: N/A
Last BPDUs 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 (Ingress)			
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	

```

Forwarding Engine Stats (Egress)
Dropped          : 0          n/a

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

Ingress Queue 3 (Profile)
Off. ColorIn : 0 0
Off. ColorOut : 0 0
Off. Uncolor : 0 0
Dro. ColorOut : 0 0
Dro. ColorIn/Uncolor : 0 0
For. InProf : 0 0
For. OutProf : 0 0
-----
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
Shared Q plcy      : n/a
=====
Egress qos-policy : 1
Multipoint shared : Disabled
=====

```

Table 89: Service ID (SAP) field descriptions

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: Service AdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, Parent IfAdminDown, NoSapIpipeCelpAddr, TodResource Unavail, TodMssResourceUnavail, SapParam Mismatch, CemSapNoEcidOrMacAddr, StandByFor McRing, ServiceMTUTooSmall, SapIngressNamed PoolMismatch, SapEgressNamedPoolMismatch, No SapEpipeRingNode.
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

Label	Description
Ingr Mac Fltr-Id	Not applicable
Egr Mac Fltr-Id	Not applicable
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
Ing Shaper Group	The ingress shaper group for the SAP
Egr Shaper Group	The egress shaper group for the SAP
ARP Reply Agent	Specifies whether 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

Label	Description
Collect Stats	Specifies whether accounting statistics are collected on the SAP
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 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.	

Label	Description
Sap Statistics	
Last Cleared Time	The date and time that a clear command was issued on statistics
Forwarding Engine Stats (Ingress)	
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
Forwarding Engine Stats (Egress)	
Dropped	The number of packets or octets dropped by the forwarding engine
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

Label	Description
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 <i>n</i>	The index of the ingress QoS queue of this SAP, where <i>n</i> 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 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
Off. ColorIn	Indicates the number of packets or octets colored as in-profile for the SAP (offered)
Off. ColorOut	Indicates the number of packets or octets colored as out-of-profile for the SAP (offered)
Off. Uncolor	Indicates the number of packets or octets that are unprofiled for the SAP (offered)
Dro. ColorOut	Indicates the number of packets or octets colored as out-of-profile that were dropped for the SAP
Dro. ColorIn/Uncolor	Indicates the number of packets or octets that were colored as in-profile or unprofiled that were dropped for the SAP
Egress Queue <i>n</i>	The index of the egress QoS queue of the SAP, where <i>n</i> 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

Label	Description
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)

sdp

Syntax

sdp [*sdp-id*[:*vc-id*]] [**detail**]

sdp far-end {*ip-address* | *ipv6-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 | ipv6-address

displays information for the SDP having this far-end IPv4 or IPv6 address
- detail

adds details to SDP information

Output

The following output is an example of service SDP information, and [Table 90: Service ID \(SDP\) field descriptions](#) describes the fields.

Output example

```
*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 Port Id: 0
Designated Bridge   : N/A

Fwd Transitions    : 0                  Bad BPDUs rcvd     : 0
Cfg BPDUs rcvd     : 0                  Cfg BPDUs tx       : 0
TCN BPDUs rcvd     : 0                  TCN BPDUs tx       : 0
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. 0thEncPkt*: 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

=====

Table 90: Service ID (SDP) field descriptions

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

Label	Description
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 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

Label	Description
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: Service AdminDown, SvcAdminDown, SdpOperDown, No IngVCLabel 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
MAC Pinning	Specifies whether MAC pinning is enabled in this SDP
Ignore Standby Sig	Specifies whether ignore standby signaling is configured True: standby signaling is ignored False: standby signaling is not ignored
Block On Mesh Fail	Specifies whether 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	

Label	Description
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
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	<p>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.</p> <p>If the SDP type is GRE, the following message displays: SDP delivery mechanism is not MPLS</p>
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 91: Service ID \(split horizon group\) field descriptions](#) describes the fields.

Output example

```
*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 91: Service ID (split horizon group) field descriptions

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 92: Service ID \(STP\) field descriptions](#) describes the fields.

Output example

```
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	BPDUs Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: N/A
Last BPDUs 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	BPDUs Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: N/A
Last BPDUs 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 92: Service ID (STP) field descriptions

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

Label	Description
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
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

Label	Description
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.
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	

Label	Description
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
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 toward 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 93: Service ingress-label field descriptions](#) describes the fields.

Output example

```
*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 93: Service ingress-label field descriptions

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

pim-snooping

Syntax
pim-snooping

Context
show>service>id

Description
This command enables the context to display PIM snooping information.

group

Syntax
group [grp-ip-address] [source ip-address] [type {starg | sg}] [detail] [family]

Context
show>service>id>pim-snooping

Description
This command displays the multicast group information.

Parameters

- grp-ip-address

displays information for the multicast group address (IPv4 or IPv6)
- ip-address

displays information for the source address or the RP address (IPv4 or IPv6)
- starg

displays information for (*,G) entries
- sg

displays information for (S,G) entries
- detail

displays detailed information
- family

displays either IPv4 or IPv6 information
- Values

ipv4 or ipv6

Output

The following output is an example of PIM snooping multicast group information, and [Table 94: Service VPLS PIM snooping group field descriptions](#) describes the fields.

Output example

```
*A:ALU# show service id 1 pim-snooping group

=====
PIM Snooping Groups ipv4
=====
Group Address      Source Address      Type      Incoming
                   Intf               (S,G)     SAP:1/1/2      Num
                   Oifs
-----
10.252.0.1         10.0.0.2            (S,G)     SAP:1/1/2      2
-----
Groups : 1
=====
*A:PE#
```

Table 94: Service VPLS PIM snooping group field descriptions

Label	Description
PIM Snooping Groups	
Group Address	The IP address of the multicast group
Source Address	The IP address of the multicast source
Type	The type of source-group pair: (S,G) or (*,G)
Incoming Intf	The interface identifier for the incoming multicast stream

Label	Description
Num Oifs	The number of outgoing interfaces
Groups	The number of groups displayed

neighbor

Syntax

neighbor [{sap *sap-id* | sdp *sdp-id:vc-id*} [**address** *ip-address*]] [**detail**] [*family*]

Context

show>service>id>pim-snooping

Description

This command displays PIM neighbor information.

Parameters

- sap-id*
displays information for the PIM neighbor associated with the specified SAP
- sdp-id:vc-id*
displays information for the PIM neighbor associated with the specified SDP
- ip-address*
displays information for the PIM neighbor with the specified IP address (IPv4 or IPv6)
- detail**
displays detailed PIM neighbor information
- family*
displays either IPv4 or IPv6 information for the specified neighbor
- Values** ipv4 or ipv6

Output

The following output is an example of PIM snooping neighbor information, and [Table 95: Service VPLS PIM snooping neighbor field descriptions](#) describes the fields.

Output example

```
*A:ALU# show service id 1 pim-snooping neighbor
=====
PIM Snooping Neighbors ipv4
=====
Sap/Sdp Id      Nbr DR Prty    Up Time        Expiry Time    Hold Time
Nbr Address
-----
SAP:1/1/1      1              0d 00:06:15    0d 00:01:40    105
10.0.1.1
```

```
SAP:1/1/2          1          0d 00:06:15   0d 00:01:30   105
 10.0.1.2
-----
Neighbors : 2
=====
*A:PE#
```

Table 95: Service VPLS PIM snooping neighbor field descriptions

Label	Description
PIM Snooping Neighbors	
Sap/Sdp Id Nbr Address	The SAP or SDP identifier and the IP address of the neighbor
Nbr DR Prty	The neighbor designated router parity
Up Time	The length of time that the multicast connection has been up
Expiry Time	The length of time remaining before the multicast connection goes down
Hold Time	The length of time that PIM snooping checks for the PIM state
Neighbors	The number of neighbors for the specified instance of PIM snooping

port

Syntax

port [**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [**group** [*grp-ip-address*]] [**detail**] [*family*]

Context

show>service>id>pim-snooping

Description

This command displays PIM port information.

Parameters

- sap-id*
displays port information for the specified SAP
- sdp-id:vc-id*
displays port information for the specified SDP
- group**
displays port information for multicast groups

- grp-ip-address*
displays port information for the specified multicast group address (IPv4 or IPv6)
- detail**
displays detailed port information
- family*
displays either IPv4 or IPv6 information for the specified port
- Values** ipv4 or ipv6

Output

The following output is an example of PIM snooping port information, and [Table 96: Service VPLS PIM snooping port field descriptions](#) describes the fields.

Output example

```
*A:ALU# show service id 1 pim-snooping port
=====
PIM Snooping Ports ipv4
=====
Sap/Sdp Id                Opr
-----
SAP:1/1/1                 Up
SAP:1/1/2                 Up
=====
*A:ALU#
```

Table 96: Service VPLS PIM snooping port field descriptions

Label	Description
PIM Snooping Ports	
Sap/Sdp Id	The SAP or SDP identifier for the port
Opr	The operational state of the port: Up or Down

statistics

Syntax

statistics [**sap** *sap-id*] [**sdp** *sdp-id:vc-id*] [*family*]

Context

show>service>id>pim-snooping

Description

This command displays PIM statistics information.

Parameters

- sap-id

displays the statistics associated with the specified SAP
- sdp-id:vc-id

displays the statistics associated with the specified SDP
- family

displays either IPv4 or IPv6 statistics
- Values

ipv4 or ipv6

Output

The following output is an example of PIM snooping statistics information, and [Table 97: Service VPLS PIM snooping statistics field descriptions](#) describes the fields.

Output example

```
*A:ALU# show service id 1 pim-snooping statistics
=====
PIM Snooping Statistics ipv4
=====
Message Type      Received      Transmitted   Rx Errors
-----
Hello             36            -             0
Join Prune        8             8             0
Total Packets     44            8
-----
General Statistics
-----
Rx Neighbor Unknown      : 0
Rx Bad Checksum Discard  : 0
Rx Bad Encoding           : 0
Rx Bad Version Discard   : 0
Join Policy Drops        : 0
-----
Source Group Statistics
-----
(S,G)                : 1
(*,G)                : 0
=====
```

Table 97: Service VPLS PIM snooping statistics field descriptions

Label	Description
Message Type	
Hello	The number of hello messages received, transmitted, and received with errors
Join Prune	The number of join-prune messages transmitted, received, and received with errors
Total Packets	The total number of hello and join-prune messages transmitted, received, and received with errors

Label	Description
General Statistics	
Rx Neighbor Unknown	The number of packets received from an unknown neighbor
Rx Bad Checksum Discard	The number of packets received and discarded because of a bad checksum
Rx Bad Encoding	The number of packets received with bad encoding
Rx Bad Version Discard	The number of packets received and discarded because of a bad version
Join Policy Drops	The number of join messages dropped due to the join policy actions
Source Group Statistics	
(S,G)	The number of (S,G)s in the multicast
(* ,G)	The number of (* ,G)s in the multicast

status

Syntax

status [family]

Context

show>service>id>pim-snooping

Description

This command displays PIM status information.

Parameters

family
displays either IPv4 or IPv6 status information

Values ipv4 or ipv6

Output

The following output is an example of PIM snooping status information, and [Table 98: Service VPLS PIM snooping status field descriptions](#) describes the fields.

Output example

```
*A:ALU# show service id 1 pim-snooping status
=====
PIM Snooping Status ipv4
```

```

=====
Admin State           : Up
Oper State            : Up
Mode Admin            : Proxy
Mode Oper             : Proxy
Hold Time             : 90
Designated Router     : 10.0.1.2
J/P Tracking          : Inactive
Up Time               : 0d 00:08:43
Group Policy          : None
=====
*A:ALU#

```

Table 98: Service VPLS PIM snooping status field descriptions

Label	Description
Admin State	The administrative state of PIM snooping for the specified service
Oper State	The operational state of PIM snooping for the specified service
Mode Admin	The administrative mode: snoop or proxy
Mode Oper	The operational mode: snoop or proxy
Hold Time	The length of time that PIM snooping checks for the PIM state
Designated Router	The IP address of the designated router
J/P Tracking	The join-prune status
Up Time	The length of time that the multicast connection has been up
Group Policy	The group policy name

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}
sap-using [ingress | egress] shaper-group shaper-group-name

```

Context

show>service

Description

This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters

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 44: SAP ID configurations](#) 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

shaper-group

specifies the shaper group for which to display matching SAPs

Output

The following output is an example of information about SAPs matching the specified properties, and [Table 99: Service SAP-using field descriptions](#) describes the fields.

Output example

```
*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 99: Service SAP-using field descriptions

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
Shaper Policy	Identifies the shaper policy that the shaper group belongs to
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

5.22.2.5 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]* | [*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 44: SAP ID configurations](#) 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 44: SAP ID configurations](#) 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 44: SAP ID configurations](#) 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 IPv4 address

grp-ipv6-address

clears MLD snooping statistics matching the specified group IPv6 address

src-ip-address

clears IGMP snooping statistics matching the specified source IPv4 address

src-ipv6-address

clears MLD snooping statistics matching the specified source IPv6 address

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 44: SAP ID configurations](#) 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

pim-snooping

Syntax

pim-snooping

Context

clear>service>id

Description

This command enables the context to clear PIM snooping information.

database

Syntax

database [[**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [**group** *grp-ip-address*] [**source** *src-ip-address*]] [*family*]

Context

clear>service>id>pim-snooping

Description

This command clears PIM snooping source group database information.

Parameters

sap-id

clears PIM snooping entries associated with the specified SAP

sdp-id:vc-id

clears PIM snooping entries associated with the specified SDP

grp-ip-address

clears PIM snooping information matching the specified group address

src-ip-address

clears PIM snooping information matching one particular source within the multicast group

family

clears either IPv4 or IPv6 information

Values ipv4 or ipv6

neighbor

Syntax

neighbor [*ip-address* | **sap** *sap-id* | **sdp** *sdp-id:vc-id*] [*family*]

Context

clear>service>id>pim-snooping

Description

This command clears PIM snooping neighbor information.

Parameters

ip-address

clears PIM snooping information for the neighbor with the specified IP address

sap-id

clears PIM snooping entries associated with the specified SAP

sdp-id:vc-id

clears PIM snooping entries associated with the specified SDP

family

clears either IPv4 or IPv6 information

Values ipv4 or ipv6

statistics

Syntax

statistics [**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [*family*]

Context

clear>service>id>pim-snooping

Description

This command clears PIM snooping statistics for the specified SAP or SDP.

Parameters

sap-id

clears PIM snooping statistics for the specified SAP

sdp-id:vc-id

clears PIM snooping statistics for the specified SDP

family

clears either IPv4 or IPv6 information

Values ipv4 or ipv6

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 44: SAP ID configurations](#) 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

5.22.2.6 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 44: SAP ID configurations](#) 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

pim-snooping

Syntax

[no] pim-snooping

Context

debug>service>id

Description

This command enables PIM-snooping debugging.

adjacency

Syntax

[no] adjacency

Context

debug>service>id>pim-snooping

Description

This command enables or disables debugging for PIM snooping adjacencies.

all

Syntax

all [**group** *grp-ip-address*] [**source** *src-ip-address*] [**detail**]

no all

Context

debug>service>id>pim-snooping

Description

This command enables or disables debugging for all the PIM snooping modules for the specified multicast address.

Parameters

grp-ip-address

debugs information associated with all PIM modules for the specified multicast group address

Values multicast group address (IPv4 or IPv6)

src-ip-address

debugs information associated with all PIM modules for the specified multicast source address

Values IPv4 or IPv6 address

detail

provides detailed debugging information on all PIM snooping modules

database

Syntax

database [**group** *grp-ip-address*] [**source** *src-ip-address*] [**detail**]

no database

Context

debug>service>id>pim-snooping

Description

This command enables or disables debugging for the PIM snooping database for the specified multicast address.

Parameters

grp-ip-address

debugs information associated with all PIM modules for the specified multicast group address

Values multicast group address (IPv4 or IPv6) or zero

src-ip-address

debugs information associated with the specified database for the specified multicast source address

Values IPv4 or IPv6 address

detail

provides detailed debugging information on the PIM snooping database

jp

Syntax

jp [**group** *grp-ip-address*] [**source** *src-ip-address*] [**detail**]

no jp

Context

debug>service>id>pim-snooping

Description

This command enables or disables debugging for the PIM snooping join-prune mechanism for the specified multicast address.

Parameters*grp-ip-address*

debugs information associated with the specified join-prune mechanism for the specified multicast group address

Values multicast group address (ipv4 or ipv6) or zero

src-ip-address

debugs information associated with the specified join-prune mechanism for the specified multicast source address

Values source IP address (IPv4 or IPv6)

detail

provides detailed debugging information on the join-prune mechanism

packet**Syntax**

packet [**hello** | **jp**] [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

no packet

Context

debug>service>id>pim-snooping

Description

This command enables or disables debugging for PIM snooping packets.

Parameters**hello | jp**

specifies the PIM snooping packet types

sap-id

debugs packets associated with the specified SAP

sdp-id:vc-id

debugs packets associated with the specified SDP

port**Syntax**

port [**sap** *sap-id* | **sdp** *sdp-id:vc-id*] [**detail**]

no port

Context

debug>service>id>pim-snooping

Description

This command enables or disables debugging for PIM snooping ports.

Parameters

sap-id

only debugs ports associated with the specified SAP

sdp-id:vc-id

only debugs ports associated with the specified SDP

detail

provides detailed debugging information on the SAP or SDP

red

Syntax

red [detail]

no red

Context

debug>service>id>pim-snooping

Description

This command enables or disables debugging for PIM snooping messages sent to the standby (redundant) CSM.

Parameters

detail

provides detailed debugging information

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 44: SAP ID configurations](#) 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 44: SAP ID configurations](#) for a full list of SAP IDs.

6 Internet enhanced service

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.

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).

Topics in this chapter include:

- [IES for in-band management](#)
- [IES for customer traffic](#)
- [Configuring IES with CLI](#)
- [IES command reference](#)

6.1 IES for in-band management

Topics in this section include:

- [Setting up connections between the NSP NFM-P 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 Shelf V2 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-M
- 7705 SAR-A
- 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.

6.1.1 Setting up connections between the NSP NFM-P 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 (NSP NFM-P).

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 NSP NFM-P.

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; see the 7705 SAR 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.

6.1.2 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*. Because 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.

6.1.3 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: Correct configuration of the traffic descriptor profiles is essential for correct 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.

6.1.4 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 about OAM, see the 7705 SAR 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 IGP and BGP protocols, including 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.

Because of 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 route to the peer is declared not active.



Note: Layer 2 AIS/RDI cells that are received on the IES SAP disable the IP interface. Link failures detected by BFD also disable the IP interface.

6.1.5 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. Layer 4 port attributes and the TEID attribute in the hashing algorithm can be configured with the **l4-load-balancing** and **teid-load-balancing** commands in the **config>service>ies>interface** context. Configuration of the **l4-load-balancing** command at the interface level overrides the

system-level settings for the specific interface. The **teid-load-balancing** command can only be configured at the interface level.

The system IP address can be included in or excluded from the hashing algorithm with the system-level **system-ip-load-balancing** command.

For more information about IP ECMP, see the 7705 SAR Router Configuration Guide, "Static routes, dynamic routes, and ECMP".


6.2 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](#)
- [VRRP on IES interfaces](#)
- [SAPs](#)
- [Spoke-SDP termination to IES](#)
- [Bandwidth optimization for low-speed links](#)
- [Hold up and hold down timers for IP interfaces](#)

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 the following figure. 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 instead of being encapsulated in a tunnel.

Figure 101: IES for customer access to the Internet



20720

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 V.35 port on the 12-port Serial Data Interface card, version 3 with speed set to 64 kb/s, 2048 kb/s, or any value from 128 kb/s to 1920 kb/s (every 128 kb/s)
- any T1/E1 port on the 7705 SAR-M
- any T1/E1 port on the 7705 SAR-A
- any T1/E1 port on the 4-port T1/E1 and RS-232 Combination module
- 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 on the 7705 SAR-M
- any Ethernet port on the 7705 SAR-A
- any Ethernet port on the 7705 SAR-Ax
- any Ethernet port on the 7705 SAR-Wx
- any Ethernet port on the 7705 SAR-H
- any Ethernet port on the 7705 SAR-Hc
- any Ethernet port on the 7705 SAR-X

Ports must be in access mode.

The encapsulation type for Ethernet ports must be null, dot1q, or qinq.

IES IPv6 SAPs are supported on the following cards, modules, and ports:

- 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-M
- any Ethernet port on the 7705 SAR-A
- any Ethernet port of the 7705 SAR-Ax
- any Ethernet port on the 7705 SAR-Wx
- the 7705 SAR-H
- any Ethernet port on the 7705 SAR-Hc
- any Ethernet port of the 7705 SAR-X

For more information about IPv6 addressing, see the 7705 SAR 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 instead of 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 about configuring fabric mode, see the 7705 SAR Quality of Service Guide, "Configurable ingress shaping to fabric (access and network)".

6.2.1 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 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 about DHCP and DHCPv6, see the 7705 SAR Router Configuration Guide, "DHCP and DHCPv6".

6.2.1.1 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. Because 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.

6.2.1.1.1 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 more 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

6.2.1.2 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 (see the 7705 SAR 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.

6.2.1.2.1 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*.

6.2.2 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.

6.2.3 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. IES interfaces support the provisioning of tunnel SAPs as part of IPsec provisioning. The *sap-id* for a public-side IPsec tunnel SAP is **tunnel-1.public:tag**.

For more information, see the [IPsec](#) chapter in this guide.

6.2.4 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, see the 7705 SAR Router Configuration Guide, "Configuring security parameters".

To enable NAT or firewall functionality for IES, security policy and profile parameters must be configured under the **config>security** context in the CLI, and a security zone must be configured under the **config>service>ies>zone** context.

A zone is created by adding at least one Layer 2 endpoint or Layer 3 interface to the zone configuration. The following table lists the supported interfaces and endpoints that can be added to zones under IES NAT or firewall.

Table 100: Security zone interfaces and endpoints for IES

CLI context	Interface/endpoint type	NAT	Firewall
IES	SAP	✓	✓
	Spoke-SDP termination	✓	✓
	IPsec public	✓	
	Routed VPLS	✓	✓

6.2.5 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 about proxy ARP, see the 7705 SAR Router Configuration Guide, "Proxy ARP".

6.2.6 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.

6.2.7 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.

6.2.8 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 IGP and BGP protocols, including 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.

Because of 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 route is declared “not active”, causing a reroute to an alternative path, if any.



Note: Link failures detected by BFD 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 about ICMP and ICMPv6, see the 7705 SAR Router Configuration Guide, “ICMP and ICMPv6”.

6.2.9 VRRP on IES interfaces

VRRP can be implemented on IES service interfaces to participate as part of a virtual router instance. This implementation prevents a single point of failure by ensuring access to the gateway address, which is configured on all IES service interfaces in the VRRP. VRRPv3 can also be implemented on IES service interfaces, including r-VPLS interfaces for IES.

The 7705 SAR supports VRRPv3 for IPv4 and IPv6 as described in RFC 5798. Within a VRRP router, the virtual routers in each of the IPv4 and IPv6 address families are in separate domains and do not overlap.



Note:

- VRRPv3 for IPv6 is not supported on a Layer 3 spoke-SDP termination.
- VRRP is not supported on an IPSec public interface.

For information about VRRP and VRRP IES service interface parameters, as well as the configuration parameters of VRRP policies, see the “VRRP” section in the 7705 SAR Router Configuration Guide. CLI command descriptions for VRRP policies are also specified in the 7705 SAR Router Configuration Guide.

For CLI command descriptions related to IES service interfaces, see [IES command reference](#).

6.2.10 SAPs

Topics in this section include:

- [Encapsulations](#)
- [Routing protocols](#)
- [QoS policies](#)
- [QinQ \(IES\)](#)
- [IP filter policies on an IES SAP](#)

6.2.10.1 Encapsulations

The following SAP encapsulations are supported on the 7705 SAR Internet enhanced service:

- Ethernet null
- Ethernet dot1q
- Ethernet qinq

- PPP/MLPPP/MC-MLPPP

6.2.10.2 Routing protocols

IES supports static routes on customer IP interfaces (that is, SAPs). These routes are redistributed into the global routing table of the 7705 SAR.

OSPFv2, RIP, and PIM routing protocols are supported on IES SAPs (that is, access IP interfaces). IES SAPs on V.35 ports on the 12-port Serial Data Interface card, version 3, support only OSPFv2 and static routing protocols.

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.

See the chapters on "OSPF" and "RIP" in the 7705 SAR Routing Protocols Guide for information about configuring these routing protocols.

IPv4 in IES supports PIM-SM and PIM-SSM. IPv6 in IES supports PIM-SSM. See the "IP multicast" chapter in the 7705 SAR Routing Protocols Guide for information about configuring these routing protocols.

6.2.10.3 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. See the 7705 SAR 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.

6.2.10.4 QinQ (IES)

IES supports QinQ functionality. For details, see [QinQ support](#).

6.2.10.5 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. See the 7705 SAR Router Configuration Guide, "Filter policies", for information about configuring IP filters.


6.2.11 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.

The following figure 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 102: SDP ID and VC label service identifiers (conceptual view of the service)



21510

The following figure 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 103: IES spoke-SDP termination



21511

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 about PPP/MLPPP ports, see the 7705 SAR Interface Configuration Guide, "Access, network, and hybrid ports"
- 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

The following figure shows an example of backhauling from a specific 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 104: Pseudowire-based backhaul (spoke-SDP termination at 7750 SR)



21522

6.2.12 Bandwidth optimization for low-speed links

The 7705 SAR can be used in deployments where the uplink bandwidth capacity and requirements are considerably less than if the router is used for fixed or mobile backhaul applications. For example, the

7705 SAR can be used to direct traffic from multiple individual homes for applications such as smart meter aggregation or relay connectivity. Connecting to end systems such as smart meters or relays requires uplink bandwidth capacity in terms of hundreds of kilobits per second, instead of hundreds of megabits per second.

One way to optimize operation in lower-bandwidth applications is to minimize head-of-line (HoL) blocking caused by large packets. HoL blocking occurs when transmission of a large non-mission-critical packet delays a mission-critical packet beyond acceptable limits. The propagation delay of large packets over a slow link is fairly significant. For example, the propagation delay when transmitting a 1500-byte packet over a 100 kb/s link is 120 ms. If a mission-critical packet is queued immediately after the first bit of a non-mission-critical 1500-byte packet begins transmission, the mission-critical packet must wait 120 ms before the uplink is available again.

To minimize HoL blocking, the 7705 SAR supports a lower MTU of 128 bytes (from the original 512-byte minimum) so that large IP packets are fragmented into 128-byte chunks. In the preceding example, transmitting a 128-byte packet over a 100 kb/s link only delays the next packet by 10.24 ms.

This lower MTU is supported on IES and VPRN interfaces (access interfaces) and on network interfaces. The IP MTU is derived from the port MTU, unless specifically configured with the **ip-mtu** command. This command is supported on access interfaces only.

The following must be considered when using a lower IP MTU:

- applicability – the lower IP MTU is only applicable for IP forwarded traffic and cannot be applied to pseudowire or VPLS traffic
- reassembly – the far-end/destination node must reassemble the packet before it can process the data, which may impact the performance of the end system or require different hardware to perform the reassembly
- extra overhead – each fragment must have an IPv4 header so that all fragments of the packet can be forwarded to the destination. Care must be taken to ensure that the extra IP overhead for each fragment does not offset the gain achieved by using the lower MTU. As an example, for a 128-byte packet, the IPv4 header, which is 20 bytes in length, constitutes approximately 15% of the total packet size.



Note:

- Lower IP MTU applies to IPv4 applications only. As per RFC 2640, IPv6 interfaces or dual-stack interfaces should not be configured to a value lower than 1280 bytes.
- Lower IP MTU is supported only on Ethernet encapsulated ports.
- Most routing and signaling protocols, such as OSPF, IS-IS, and RSVP-TE, cannot be supported with port MTUs lower than 512 bytes because of the protocol layer requirements and restrictions.
- Special care must be taken with routing protocols that use TCP, such as BGP and LDP. The minimum TCP MSS value supported on the 7705 SAR is 384 bytes; therefore, these protocols should only be enabled on links that can transport 384-byte IP packets without fragmentation. If there is a mismatch in TCP MSS in the network, this mismatch can potentially cause severe network performance issues because of the overhead caused by fragmentation and retransmissions, it can cause multi-vendor interoperability issues, and it can potentially cause the protocols to continuously flap.
- Not all OAM diagnostics are supported with lower port MTUs. Detailed information is provided in [OAM diagnostics restrictions with lower IP MTU](#).

6.2.12.1 OAM diagnostics restrictions with lower IP MTU

OAM tests require a minimum network port MTU in order to run; this value depends on the test. If the port MTU is set to a value lower than the minimum requirement, the test fails.

If the port MTU is set to a value that meets the minimum requirement, the packet size parameter can be configured for the test (for example, `oam sdp-ping 1 size 102`).

If the **size** parameter is not specified, the system builds the packet based on the default payload size. If the **size** parameter is configured and is greater than the default payload size, padding bytes are added to equal the configured value.

The packet size is dependent on the port MTU value; that is, if the minimum port MTU value is used, there are restrictions on the packet size. If the configured size is greater than the maximum value supported with the minimum port MTU, the test fails.

[Table 101: Port MTU requirements for OAM diagnostics \(GRE tunnels\)](#) and [Table 102: Port MTU requirements for OAM diagnostics \(LDP tunnels\)](#) list the minimum port MTU required for each OAM test and the maximum size of the OAM packet that can be configured when the minimum port MTU is used, based on SDP tunnel type.



Note: RSVP LSPs do not come up if the network port MTU value is lower than 302 bytes.

Table 101: Port MTU requirements for OAM diagnostics (GRE tunnels)

SDP type: GRE		
Test type	Minimum network port MTU requirement over Ethernet dot1q encapsulation (bytes)	OAM test size range (bytes)
sdp-ping	128	72 to 82
svc-ping	196	N/A ¹
vccv-ping	143	1 to 93
vccv-trace	143	1 to 93
vprn-ping	182	1 to 136
vprn-trace	302	1 to 256
mac-ping	188	1 to 142
mac-trace	240	1 to 194
cpe-ping	186	N/A ¹

Note:

1. Size is not configurable

Table 102: Port MTU requirements for OAM diagnostics (LDP tunnels)

SDP type: LDP		
Test type	Minimum network port MTU requirement over Ethernet dot1q encapsulation (bytes)	OAM test size range (bytes)
lsp-ping	128	1 to 106
lsp-trace	128	1 to 104
sdp-ping	128	72 to 102
svc-ping	176	N/A ¹
vccv-ping	128	1 to 98
vccv-trace	128	1 to 98
vprn-ping	182	1 to 156
vprn-trace	302	1 to 276
mac-ping	168	1 to 142
mac-trace	220	1 to 194
cpe-ping	166	N/A ¹

Note:

1. Size is not configurable

For information about OAM diagnostics, see the 7705 SAR OAM and Diagnostics Guide.

6.2.13 Hold up and hold down timers for IP interfaces

The 7705 SAR allows timers to be configured on an IES or VPRN IPv4 or IPv6 interface or on the base router to keep the IP interface in an operationally up or down state for a specified time beyond when it should be declared operationally up or down. The timers are configured at the IES service level using the **config>service>ies>interface>hold-time>up/down** commands. An **init-only** option enables the **down** delay to be applied only when the IP interface is first configured or after a system reboot. See [VPRN services](#) for information about how to configure the **hold-time** command on IES interfaces. See the 7705 SAR Router Configuration Guide for information about how to configure the **hold-time** command at the router level.

The configuration causes the system to delay sending notifications of any state change associated with the IP interface until the timer has expired.



Note: The **up** and **down** delay timers in the CLI are dynamic text fields; the fields are only displayed in the **show router interface detail** command output if they are configured. The field showing the time remaining is displayed only if the timer is actively counting down. If both **up** and

down timers are configured, the field displayed depends on the current operational state of the interface. For example, if the interface is operationally down, the configured hold **down** time is displayed.

6.3 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](#)

6.4 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 NSP NFM-P 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 NSP NFM-P.*
- Enable the service.



Note: *Remote address and static route configuration is beyond the scope of this document. For information, see the 7705 SAR Router Configuration Guide.

6.5 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 an example of the CLI display output.

Topics in this section include:

- [Creating an IES service](#)

- [Configuring IES interface parameters](#)
- [Configuring IES SAP parameters](#)
- [Configuring IES spoke SDP parameters](#)
- [Configuring VRRP](#)
- [Configuring a security zone within IES](#)
- [Configuring serial raw socket transport within IES](#)

6.5.1 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
...
```

6.5.2 Configuring IES interface parameters

Configure interface parameters for:

- [IES management service](#)
- [IES service](#)
- [IES IPv6 service](#)

6.5.2.1 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 10.1.1.1/8
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 10.1.1.1/8
    ip-mtu 1524
    no bfd
    exit
    no shutdown
...
-----
```

6.5.2.2 IES service

Use the following CLI syntax to configure interface parameters for the IES service.



Note: The IES 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.

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]
        unreachable
    ip-mtu octets
    ipcp
        dns ip-address [secondary ip-address]
        dns secondary ip-address
        peer-ip-address ip-address
    l4-load-balancing hashing-algorithm
    local proxy-arp
    loopback
    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
    teid-load-balancing
    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 192.168.0.0/16
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 192.168.0.0/16 broadcast host-ones
    no mac
    arp-timeout 14400
    no allow-directed-broadcasts
    icmp
        mask-reply
        unreachable 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 "proxyARPolicy"
local proxy-arp
remote proxy-arp
no shutdown...
-----

```

6.5.2.3 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]
      unreachable [number seconds]
    neighbor ipv6-address mac-address
    reachable-time seconds
    stale-time seconds

```

Example:

```

config>service# ies 9
config>service>ies$ interface "ies_interface"
config>service>ies>if$ ipv6
config>service>ies>if>ipv6$ address 2001:db8:0:1:1:1:1:1/24
config>service>ies>if>ipv6$ dhcp6-relay

```

```

config>service>ies>if>ipv6>dhcp6-relay$ server 2001:db8::1
config>service>ies>if>ipv6>dhcp6-relay$ option
config>service>ies>if>ipv6>dhcp6-relay>option$ interface-id ascii-tuple
config>service>ies>if>ipv6>dhcp6-relay>option$ exit
config>service>ies>if>ipv6$ icmp
config>service>ies>if>ipv6>icmp6$ packet-too-big 80 10
config>service>ies>if>ipv6>icmp6$ exit
config>service>ies>>if>ipv6# neighbor 2001:db8::2
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 2001:db8:0:1:1:1:1:1/24
    dhcp6-relay
        no description
        option
            interface-id ascii-tuple
            no remote-id
        server 2001:db8:0:1:1:1:1:1
    exit
    icmp6
        packet-too-big 80 10
        param-problem 100 10
        time-exceeded 100 10
        unreachable 100 10
    exit
    exit
    ...
    reachable-time 30
    stale-time 14400
    exit
...

```

6.5.3 Configuring IES SAP parameters

Configure IES SAP parameters for:

- [IES management SAP](#)
- [IES service SAP](#)

6.5.3.1 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#

```

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
-----

```

6.5.3.2 IES service SAP

Use the following CLI syntax to configure SAP parameters for the IES service.



Note: A SAP cannot be defined if the **loopback** command is enabled on the interface.

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
-----
```

6.5.4 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
-----

```

6.5.5 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 addresses 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 more information about VRRP CLI syntax and command descriptions for an IES service interface, see [IES command reference](#). For overview information about VRRP and VRRP IES interface parameters, see the "VRRP" chapter in the 7705 SAR Router Configuration Guide.

The following displays an IES interface VRRP owner configuration:

```

config>service>ies> info
#-----
...
  interface "vrrpowner"
    address 10.10.10.23/16
    vrrp 1 owner
      backup 10.10.10.24
      authentication-key "testabc"
    exit

```

```

    exit
    ...
#-----
config>service>ies#

config>service>ies>if># info
-----
...
    ipv6
        address 2001:db8:0:1:1:1:1:1/16
        vrrp 1 owner
            backup 2001:db8:0:1:1:1:1:2
        exit
    exit
    exit
    exit
...
-----

```

6.5.6 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-100-198.51.100.0/24
    exit
  exit
  exit
  policy 1 nat pool 1
  commit
  exit
  no-shutdown
-----
A:ALU-B>config>service>ies#

```

6.5.7 Configuring serial raw socket transport within IES

Configure an IP transport subservice within an IES service to enable the transport of serial data using raw sockets.

CLI syntax:

```

config>service
  ies service-id [customer customer-id] [create]
  ip-transport ipt-id [create]
    description description-string
    filter-unknown-host
    local-host ip-addr ip-addr port-num port-num] protocol {tcp|
udp}
    remote-host host-id [ip-addr ip-addr] [port-num port-num]
  [create]
    description description-string
    name host-name
    exit
    fc fc-name profile {in |out}
    shutdown
    tcp
    inactivity-timeout number
    max-retries seconds
    retry-interval seconds
    exit
  exit
  exit
exit

```

The following example displays an IP transport subservice configuration output.

```

A:ALU-B>config>service>ies# info
-----
configure
  service ies 20 create

```

```
ip-transport 1/2/4.1 create
description "ip-transport one"
filter-unknown-host
local-host ip-address 192.168.1.1 port-number 4000 protocol udp
    exit
remote-host 1 ip-address 192.168.1.7 port-number 4001 create
    exit
exit
no-shutdown
-----
A:ALU-B>config>service>ies#
```

6.6 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](#)

6.6.1 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
```

6.6.2 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
```

6.6.3 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
```

6.6.4 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 at 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
```

6.7 IES command reference

6.7.1 Command hierarchies

- Configuration commands
 - IES management configuration commands
 - IES service configuration commands
 - Routed VPLS commands
 - VRRP commands
 - IES security zone configuration commands
 - IES raw socket IP transport configuration commands
- Show commands
- Clear commands
- Debug commands

6.7.1.1 Configuration commands

6.7.1.1.1 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

```

6.7.1.1.2 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
    - cflowd-parameters
      - sampling {unicast | multicast} type {interface} [direction {ingress-only
| egress-only | both}]
      - no sampling {unicast | multicast}
    - 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
    - hold-time

```

[type np]

```

- down ip seconds [init-only]
- no down ip
- down ipv6 seconds [init-only]
- no down ipv6
- up ip seconds
- no up ip
- up ipv6 seconds
- no up ipv6
- icmp
- [no] mask-reply
- ttl-expired [number seconds]
- no ttl-expired
- unreachableables [number seconds]
- no unreachableables
- 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] [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
- 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]
- unreachableables [number seconds]
- no unreachableables
- link-local-address ipv6-address [preferred]
- no link-local-address
- [no] local-dhcp-server local-server-name
- neighbor ipv6-address mac-address
- no neighbor ipv6-address
- reachable-time seconds
- stale-time seconds
- tcp-mss value
- no tcp-mss
- l4-load-balancing hashing-algorithm
- 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 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] static-nat-inside
- tcp-mss value
- no tcp-mss
- [no] teid-load-balancing
- unnumbered {ip-int-name | ip-address}
- no unnumbered
- service-name service-name
- no service-name
- [no] shutdown

```

6.7.1.1.3 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 ip-filter-id
-   - no v4-routed-override-filter
-   - v6-routed-override-filter ipv6-filter-id
-   - no v6-routed-override-filter

```

6.7.1.1.4 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] [passive]
-   - no vrrp virtual-router-id
-   - [no] backup ipv6-address
-   - [no] bfd-enable service-id interface interface-name dst-ip ip-address
-   - [no] bfd-enable interface interface-name dst-ip 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] ntp-reply
-   - [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] [passive]
-   - no vrrp virtual-router-id
-   - authentication-key [authentication-key | hash-key] [hash | hash2]
-   - no authentication-key
-   - [no] backup ip-address
-   - [no] bfd-enable service-id interface interface-name dst-ip ip-address
-   - [no] bfd-enable interface interface-name dst-ip 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] ntp-reply
- [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

```

6.7.1.1.5 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

```

6.7.1.1.6 IES raw socket IP transport configuration commands

```

config
- service
- ies service-id [customer customer-id] [create]
- no ies service-id
- ip-transport ipt-id [create]
- no ip-transport ipt-id
  - description description-string
  - no description
  - dscp dscp-name
  - fc fc-name profile {in | out}
  - [no] filter-unknown-host
  - local-host ip-addr ip-addr port-num port-num protocol {tcp | udp}
  - no local-host
  - remote-host host-id [ip-addr ip-addr] [port-num port-num] [create]
  - no remote-host host-id
    - description description-string
    - no description
    - name host-name
    - no name
  - [no] shutdown
- tcp
  - inactivity-timeout seconds
  - max-retries number
  - retry-interval seconds

```

6.7.1.2 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]
- ip-transport ipt-id [detail | statistics]
- remote-host host-id [detail | statistics]
- macsec
- sap [sap-id] [detail]
- ingress-label start-label [end-label]
- ip-transport-using [ip-transport ipt-id]
- 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]

```

6.7.1.3 Clear commands

```
clear
- service
  - id service-id
    - dhcp
      - statistics [ip-int-name | ip-address]
    - dhcp6
      - statistics [ip-int-name | ip-address]
    - ip-transport ipt-id
      - remote-host host-id
      - statistics
    - statistics
```

6.7.1.4 Debug commands

```
debug
- service
  - id service-id
```

6.7.2 Command descriptions

- [IES generic configuration commands](#)
- [IES global configuration commands](#)
- [IES management configuration commands](#)
- [IES service configuration commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

6.7.2.1 IES generic configuration 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>if>sap

config>service>ies>ip-transport

config>service>ies>ip-transport>remote-host

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 (such as #, \$, or spaces), 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
config>service>ies>ip-transport
```

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 IP transport subservice

when an IP transport subservice within an IES service is shut down, all TCP/UDP packets received from remote hosts are dropped and any serial data received from the serial port is dropped. Any TCP connections that were up are closed and no new TCP connection requests are accepted.

It is not possible to make configuration changes to an IP transport subservice without performing a **shutdown** first.

The operational state of an IP transport subservice is relative to the operational state of the serial port for which the IP transport subservice is defined. When a serial port is shut down, the IP transport subservice associated with the serial port becomes operationally down.

When the **no shutdown** command is executed for an IP transport subservice, it becomes operationally up, serial data from the serial port is encapsulated in TCP/UDP packets destined for remote hosts, and TCP/UDP packets can be received by the local host, where raw serial data is then sent out the serial port.

6.7.2.2 IES global configuration 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

6.7.2.3 IES management configuration commands

- [IES management interface commands](#)
- [IES management SAP commands](#)

6.7.2.3.1 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). 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. 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 (such as #, \$, or spaces), 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/mask | *ip-address*

the IP address of the IP interface

netmask

the subnet mask in dotted-decimal notation

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.

cflowd-parameters

Syntax

cflowd-parameters

Context

config>service>ies>interface

Description

This command enables the context to configure cflowd parameters for the specified IP interface.

Cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement.

Default

n/a

sampling

Syntax

sampling {unicast | multicast} type {interface} [direction {ingress-only | egress-only | both}]
no sampling {unicast | multicast}

Context

config>service>ies>if>cflowd-parameters

Description

This command configures the cflowd sampling behavior to collect traffic flow samples through a router for analysis.

This command can be used to configure the sampling parameters for unicast and multicast traffic separately.

If cflowd sampling is enabled with no **direction** parameter specified, **ingress-only** sampling is enabled by default.

The **no** form of the command disables the configured type of traffic sampling on the interface.

Default

no sampling unicast no sampling multicast

Parameters

unicast

cflowd samples unicast traffic on the interface

multicast

cflowd samples multicast traffic on the interface

interface

specifies that all traffic entering or exiting the interface is subject to sampling. Interface is the only sampling type supported on the 7705 SAR and must be specified with this command.

direction

specifies the direction in which to collect traffic flow samples: **ingress-only**, **egress-only**, or **both**

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 128 to 9732

6.7.2.3.2 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. See the 7705 SAR Interface Configuration Guide for information about 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 44: SAP ID configurations](#) 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 about configuring IP filter IDs, see the 7705 SAR 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 specific 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 Shelf V2 or 7705 SAR-18 or on T1/E1 ports on a 7705 SAR-M
- configuring ATM-related configuration for ATM-based SAPs that exist on T1/E1 ASAP adapter cards on a 7705 SAR-8 Shelf V2 or 7705 SAR-18 or on T1/E1 ports on a 7705 SAR-M

If ATM functionality is not supported for a specific 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 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 about defining traffic descriptor profiles, see the 7705 SAR 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 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

6.7.2.4 IES service configuration 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 raw socket IP transport configuration commands](#)

6.7.2.4.1 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 (such as #, \$, or spaces), 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/mask | ip-address

the IP address or the IP address and mask of the IP interface

netmask

the subnet mask in dotted-decimal notation

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 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

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 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

hold-time

Syntax

hold-time

Context

```
config>service>ies>interface
```

Description

This command enables the CLI context to configure interface hold-up or hold-down timers.

Default

n/a

down

Syntax

down ip seconds [init-only]

no down ip

down ipv6 seconds [init-only]

no down ipv6

Context

```
config>service>ies>if>hold-time
```

Description

This command enables a delay in the activation of the IPv4 or IPv6 interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface up, unless the **init-only** option is configured. If the **init-only** option is first configured, the delay is only applied when the IP interface is first configured or after a system reboot.

The **no** form of this command disables the delay in the activation of the IPv4 or IPv6 interface. Removing the configuration during an active delay period stops the delay period immediately.

Default

n/a

Parameters

ip

specifies that the configured **down** delay is applied to an IPv4 interface

ipv6

specifies that the configured **down** delay is applied to an IPv6 interface

seconds

specifies the time delay, in seconds, before the interface is activated

Values 1 to 1200

init-only

specifies that the configured **down** delay is applied only when the interface is configured or after a reboot

up

Syntax

up ip seconds

no up ip

up ipv6 seconds

no up ipv6

Context

config>service>ies>if>hold-time

Description

This command enables a delay in the deactivation of the IPv4 or IPv6 interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface down.

The **no** form of this command disables the delay in the deactivation of the IPv4 or IPv6 interface. Removing the configuration during an active delay period stops the delay period immediately.

Default

n/a

Parameters**ip**specifies that the configured **up** delay applies to an IPv4 interface**ipv6**specifies that the configured **up** delay applies to an IPv6 interface**seconds**

specifies the time delay, in seconds, before the interface is deactivated

Values 1 to 1200**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 default value is derived from the port MTU. The **no** form of the command returns the default value.

Default

no ip-mtu – uses the value derived from the port MTU

Parameters

octets

the MTU for the interface

Values 128 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 addresses 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

load-balancing

Syntax

load-balancing

Context

config>service>ies>interface

Description

This command enables the context to configure load balancing hashing options on the interface. The options enabled at the interface level overwrite parallel system-level configurations.

Default

n/a

l4-load-balancing

Syntax

l4-load-balancing *hashing-algorithm*
no l4-load-balancing

Context

config>service>ies>interface>load-balancing

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 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

<i>hashing-algorithm</i>	
specifies that Layer 4 source and destination port fields are included in or excluded from the hashing calculation	
Values	includeL4 : include Layer 4 source and destination port fields in the hashing calculation for TCP/UDP packets excludeL4 : exclude Layer 4 source and destination port fields in the hashing calculation for TCP/UDP packets
Default	the system configuration setting (under the config>system context)

spi-load-balancing

Syntax

[no] spi-load-balancing

Context

config>service>ies>interface>load-balancing

Description

This command enables SPI hashing for ESP/AH encrypted IPv4 or IPv6 traffic at the interface level.

The **no** form of this command disables SPI hashing.

Default

no spi-load-balancing

teid-load-balancing

Syntax

[no] teid-load-balancing

Context

config>service>ies>interface>load-balancing

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

local-dhcp-server

Syntax

[no] local-dhcp-server *local-server-name*

Context

config>service>ies>interface

config>service>ies>if>ipv6

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

loopback

Syntax

[no] **loopback**

Context

config>service>ies>interface

Description

This command specifies that the interface is a loopback interface that has no associated physical interface. If this command is enabled, a SAP cannot be defined on the interface.

Default

no loopback

mac

Syntax

mac *ieee-address*

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

a 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 and cannot be all zeros. 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 (see the 7705 SAR 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 (such as #, \$, or spaces), 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/mask | ip-address

the IP address or the IP address and mask length of the IP interface

netmask

the subnet mask in dotted-decimal notation

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 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 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

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.

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.

static-nat-inside**Syntax**

[no] **static-nat-inside**

Context

config>service>ies>interface

Description

This command configures an interface as an inside (private) interface.

By default, all interfaces are outside (public) interfaces. The **no** form of this command returns the interface to the default setting.

Default

no static-nat-inside

tcp-mss

Syntax

tcp-mss *value*

no tcp-mss

Context

config>service>ies>interface

config>service>ies>if>ipv6

Description

This command configures the maximum segment size (MSS) in a TCP SYN or SYN-ACK packet during the establishment of a TCP connection. A **tcp-mss** value can be specified on an ingress interface, egress interface, or both. When configured on two interfaces, the smaller of the two values is used. If the TCP SYN packet has no TCP MSS field, the 7705 SAR assigns it the MSS value configured on the interface and recalculates the IP checksum. If the TCP SYN or SYN-ACK packet has an MSS field and the value is greater than the value configured on the interface, the 7705 SAR overwrites the packet MSS value with the lower value. If the MSS value is less than the value configured on the interface, the packet MSS value does not change. See the 7705 SAR Router Configuration Guide, "TCP MSS Configuration and Adjustment", for more information.

This command is supported on interfaces with IPv4 and IPv6 traffic, and a different MSS value can be configured for the IPv4 and IPv6 interfaces. This command is not supported on IPsec public interfaces in IES.

Default

no tcp-mss

Parameters

value

the MSS, in bytes, to be used in a TCP SYN or SYN-ACK packet

Values 384 to 9732

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

6.7.2.4.2 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] [preferred]
```

```
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

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.

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 address may be used as the source or destination address of packets sent from or to the interface.

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 (such as #, \$, or spaces), 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

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

link-local-address

Syntax

link-local-address *ipv6-address* [**preferred**]

no link-local-address

Context

config>service>ies>if>ipv6

Description

This command configures the IPv6 link-local address.

The **no** form of the command removes the configured link-local address, and the router automatically generates a default link-local address.

Removing a manually configured link-local address may impact routing protocols that have a dependency on that address.

Default

n/a

Parameters

ipv6-address

the IPv6 address

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 address may be 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>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

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

6.7.2.4.3 IES service VRRP commands

vrrp

Syntax

vrrp *virtual-router-id* [**owner**] [**passive**]

no vrrp *virtual-router-id*

Context

config>service>ies>interface

config>service>ies>if>ipv6

Description

This command creates or edits a virtual router ID (VRID) 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

owner

keyword used to identify this virtual router instance as owning the virtual router IP addresses. If the **owner** keyword is not specified at the time of VRID creation, the **vrrp backup** command must be used to define the virtual router IP addresses. The **owner** keyword is not required when entering the VRID for editing purposes. When created as **owner**, a VRID on an IP interface cannot have the **owner** parameter removed. The VRID must be deleted, and then recreated without the **owner** keyword, to remove ownership.

passive

keyword used to identify this virtual router instance as **passive**, owning the virtual router IP addresses. A **passive** VRID does not send or receive VRRP advertisement messages and is always in either the master state (if the interface is operationally up), or the initialize state (if the interface is operationally down). The **passive** keyword is not required when entering the VRID for editing purposes. When a VRID on an IP interface is created as

passive, the parameter cannot be removed from the VRID. The VRID must be deleted, and then recreated without the **passive** keyword, to remove the parameter.

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 is 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 alphanumeric 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 11 characters in length (encrypted) for a hash key or up to 110 characters for a hash2 key. If spaces are used in the string, the entire string must be enclosed 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 that 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 that 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*

[no] **backup** *ipv6-address*

Context

config>service>ies>if>vrrp

config>service>ies>if>ipv6>vrrp

Description

This command configures virtual router IP addresses for backup.

Default

n/a

Parameters

ip-address

specifies the destination IPv4 address for backup

ipv6-address

specifies the destination IPv6 address for backup

bfd-enable

Syntax

[no] **bfd-enable** *service-id* **interface** *interface-name* **dst-ip** *ip-address*

[no] **bfd-enable** **interface** *interface-name* **dst-ip** *ip-address*

Context

config>service>ies>if>vrrp

config>service>ies>if>ipv6>vrrp

Description

This command assigns a BFD session that provides a heartbeat mechanism for a VRRP instance. Only one BFD session can be assigned to a VRRP instance, but multiple VRRP instances can use the same BFD session.

BFD controls the state of the associated interface. By enabling BFD on a 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 specified in this command.

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 2147483690 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 *mac-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-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.

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 the backup router 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. See the "Master Down Interval" in the "VRRP" chapter of the 7705 SAR Router Configuration Guide for details.

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 is silently discarded.

Configuring the message interval value can be done in three ways: using only the milliseconds value, using only the seconds value, or using a combination of the two values. The following table shows the ranges for each way of configuring the message interval.

Table 103: Message interval configuration ranges

Configuration	IPv4	IPv6
Using milliseconds value only	100 to 900 ms	10 to 990 ms
Using seconds value only	1 to 255 s	1 to 40 s
Using combination milliseconds and seconds values	1 s 100 ms to 255 s 900 ms (1.1 s to 255.9 s)	1 s 10 ms to 40s 990 ms (1.01 s to 40.99 s)
Default setting	1 s	1 s

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.

Default

1 s

Parameters

seconds

the time interval, in seconds, between sending advertisement messages.

Values IPv4: 1 to 255
IPv6: 1 to 40

milliseconds

the time interval, in milliseconds, between sending advertisement messages. This parameter is not supported on non-redundant chassis.

Values IPv4: 100 to 900
IPv6: 10 to 990

ntp-reply

Syntax

[no] ntp-reply

Context

config>service>ies>if>vrrp
config>service>ies>if>ipv6>vrrp

Description

This command enables the reception of and response to Network Time Protocol (NTP) requests directed at the VRRP virtual IP address. This behaviour only applies to the router currently acting as the master VRRP.

The **no** form of this command disables NTP requests from being processed.

Default

no ntp-reply

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 parent 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.

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). VRRP policies are defined under the **config>vrrp>policy** context. For details, see the "VRRP" chapter in the 7705 SAR Router Configuration Guide.

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 effect 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 always preempts all other virtual routers when it is available.

Non-owner backup virtual router instances 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 only allows 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 is in-use priority for the virtual router instance.

Values	1 to 254
Default	100

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 parent 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 for 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 parent 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

6.7.2.4.4 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 does 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. See the 7705 SAR 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 are 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](#).

If the IES interface has been configured as a loopback interface with the [loopback](#) command, a SAP cannot be defined on the interface.

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

specifies the physical port identifier portion of the SAP definition. See [Table 44: SAP ID configurations](#) 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* as 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 does 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) Ethernet adapter card, the scheduler mode must be set to 16-priority. On a third-generation (Gen-3) Ethernet adapter card, the scheduler mode is always 4-priority. For information on adapter card generations, see the "Evolution of Ethernet Adapter Cards, Modules, and Platforms" section in the 7705 SAR 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**.

If the specified SAP is a LAG SAP, *agg-rate* and *cir-rate* is configured regardless of the scheduler mode setting on Gen-2 or Gen-3 hardware. If the active port is on a Gen-3 card or platform, *agg-rate* and *cir-rate* are applicable. If the active port is on a Gen-2 card or platform, *agg-rate* and *cir-rate* apply when the scheduler mode is set to 16-priority. For details on the behavior of a mix-and-match LAG SAP, see the "LAG Support on Third-Generation Ethernet Adapter Cards, Ports, and Platforms" and "Network LAG Traffic Management" sections in the 7705 SAR Interface Configuration Guide.

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 fails and an error message is 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 about configuring IP filter IDs, see the 7705 SAR 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.

The following table 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-dot1p 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.

Table 104: 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

Port/SAP type	Match-qinq-dot1p setting ¹	Existing packet tags	P-bits used for match
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.

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.

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).

The following table 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 are 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 105: 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
Dot1q/ 0 ²	n/a	Outer tag
QinQ/ X.Y ¹	False	Two outer tags ³
	True	Outer tag ³
QinQ/ X.* ¹	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 is 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 has its dot1p bits set to the inner tag's dot1p bits.

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 before associating the policy with a SAP. If the *policy-id* does not exist, an error is 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 is associated with an IES SAP at one time. Attempts to associate a second QoS policy of a specified 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, which only support 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.

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

6.7.2.4.5 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 destination 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 fails and an error message is 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

6.7.2.4.6 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**). After 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 specified 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 *ip-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

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)

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)

6.7.2.4.7 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. When 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 applies a security log to the specified zone. The security log must already be configured in the **config>security>logging** 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

the name of 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 (such as #, \$, or spaces), 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 are 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 are 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

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 (such as #, \$, or spaces), 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

6.7.2.4.8 IES raw socket IP transport configuration commands

ip-transport

Syntax

ip-transport *ipt-id* [create]

no ip-transport *ipt-id*

Context

config>service>ies

Description

This command creates an IP transport subservice within an IES service. An IP transport subservice is used to transmit serial raw socket data to and from a local host and remote host.

All IP transport subservices must be explicitly created using the **create** keyword. An IP transport subservice is owned by the service within which it is created. An IP transport subservice can only be associated with a single service. The **create** keyword is not needed when editing parameters for an existing IP transport subservice. An IP transport subservice must be first shut down before changes can be made to the configured parameters.

The **no** form of this command deletes the IP transport subservice with the specified *ipt-id*. When an IP transport subservice is deleted, all configured parameters for the IP transport subservice are also deleted.

Default

no ip-transport

Parameters

ipt-id

the IP transport subservice physical port identifier. The *ipt-id* must reference an RS-232 serial port that has been configured as a **socket** and has its encapsulation type set to **raw**. See the 7705 SAR Interface Configuration Guide, "Serial commands", for more information.

Values value in the format *slot/mda/port.channel*

create

creates this IP transport subservice

dscp

Syntax

dscp *dscp-name*

Context

config>service>ies>ip-transport

Description

This command configures the DSCP name used to mark the DSCP field in IP transport packets originating from this node.

Raw socket traffic redirection to a specific queue is enabled by the **fc** command.

Default

ef

Parameters

dscp-name
the DSCP name used to mark the DSCP field in IP transport packets

Table 106: Valid DSCP names

dscp-name
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

Syntax

fc [*fc-name*] **profile** {in | out}

Context

config>service>ies>ip-transport

Description

This command configures the forwarding class and profile marking for IP transport packets originating from this node.

Default

ef for fc, in for profile

Parameters

fc-name
the forwarding class name to use for the IP transport packets

Values be, l2, af, l1, h2, ef, h1, nc

profile {in| out}
the profile marking for the IP transport packets, either in or out

filter-unknown-host

Syntax

[no] **filter-unknown-host**

Context

```
config>service>ies>ip-transport
```

Description

This command filters connections from unknown hosts. An unknown host is any host that is not configured as a remote host.

The **no** form of this command disables the filter.

Default

no filter-unknown-host

local-host

Syntax

local-host *ip-addr ip-addr* **port-num** *port-num* **protocol** {**tcp** | **udp**}

no local-host

Context

```
config>service>ies>ip-transport
```

Description

This command creates the local host within the IP transport subservice.

The local host is required to accept TCP/UDP sessions initiated from far-end remote hosts, and for the node to initiate sessions toward the far-end remote hosts.

The local host must be created before a remote host is created.

The **no** form of this command deletes the local host.

Default

no local-host

Parameters

ip-addr

the IP address that is used for this local host. The IP address must be the same as a loopback or local interface IP address that is already configured within this service.

port-num

the port number that is used by remote hosts to establish TCP/UDP sessions to this local host

Values 1026 to 49150

protocol {**tcp** | **udp**}

the protocol type that is used for all sessions to and from this local host, either tcp or udp

remote-host

Syntax

remote-host *host-id* **ip-addr** *ip-addr* **port-num** *port-num* [**create**]

no remote-host *host-id*

Context

config>service>ies>ip-transport

Description

This command creates a remote host within the IP transport subservice. Multiple remote hosts may be created in order to send serial raw socket IP transport data to multiple destinations. The **create** keyword must be used for each remote host that is created.

The **no** form of this command deletes the remote host.

Default

no remote-host

Parameters

host-id

the remote host identifier

Values 1 to 2147483647 or a name string up to 64 characters

ip-addr

the IP address that is used to reach the remote host in order to route IP transport packets to that remote host

port-num

the destination port number that is used to reach the serial port socket on the remote host

Values 1 to 65535

create

creates this remote host

name

Syntax

name *host-name*

no name

Context

config>service>ies>ip-transport>remote-host

Description

This command configures a unique name for this remote host.

The **no** form of this command deletes the remote host name.

Default

n/a

Parameters

host-name

a unique name for this remote host, up to 64 characters long

tcp

Syntax

tcp

Context

config>service>ies>ip-transport

Description

This command enables the context to configure TCP parameters within this IP transport subservice.

Default

n/a

inactivity-timeout

Syntax

inactivity-timeout *seconds*

Context

config>service>ies>ip-transport>tcp

Description

This command specifies how long to wait before disconnecting a TCP connection because of traffic inactivity over the connection.

Default

30 s

Parameters

seconds

how long to wait, in seconds, before disconnecting a TCP connection

Values 1 to 65535

max-retries

Syntax

max-retries *number*

Context

config>service>ies>ip-transport>tcp

Description

This command specifies the number of times that a remote host, acting as a client, tries to establish a TCP connection after the initial attempt fails.

Default

5

Parameters

number

the number of attempts to establish a TCP connection after the initial attempt fails

Values 0 to 10

retry-interval

Syntax

retry-interval *seconds*

Context

config>service>ies>ip-transport>tcp

Description

This command specifies how long to wait before each TCP **max-retries** attempt.

Default

5 s

Parameters

seconds
how long to wait, in seconds, before each TCP **max-retries** attempt

Values 1 to 300

6.7.2.5 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
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 107: Service customer field descriptions](#) describes the fields.

Output example

```
A:ALU-2# show service customer 1
=====
Customer 1
=====
Customer-ID       : 1
Contact           : Tech Support
Description       : Default customer
Phone             : (613) 555-1122
=====
```

Table 107: Service customer field descriptions

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 108: Service egress field descriptions](#) describes the fields.

Output example

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
-----
*A:ALU-12#
```

Table 108: Service egress field descriptions

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 109: Service ID all field descriptions](#) describes the fields.

Output example (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 (Ingress)
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

Ingress Queue 3 (Profile)
Off. ColorIn     : 0          0
Off. ColorOut    : 0          0
Off. Uncolor     : 0          0
Dro. ColorOut    : 0          0
Dro. ColorIn/Uncolor : 0      0
For. InProf      : 0          0
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 109: Service ID all field descriptions

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

Label	Description
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, PortOper Down, and so on.
Multi Svc Site	Indicates the multiservice site that the SAP is a member of
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
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)
Ingr IPv6 Fltr-Id	Specifies the ingress IPv6 filter policy ID assigned to the SAP

Label	Description
Egr IPv6 Fltr-Id	Specifies the egress IPv6 filter policy ID assigned to the SAP
tod-suite	n/a
qinq-pbit-marking	Indicates the qinq P-bit marking for the SAP: both or top
Ing Scheduler Mode	Indicates the ingress scheduler mode for the SAP
Egr Scheduler Mode	Indicates the egress scheduler mode for the SAP
Ing Agg Rate Limit	Indicates the PIR rate limit in the access ingress direction for the aggregate of the SAP queues
Egr Agg Rate Limit	Indicates the PIR rate limit in the access egress direction for the aggregate of the SAP queues
Ing Agg cir	Indicates the CIR rate limit in the access ingress direction for the aggregate of the SAP queues
Egr Agg cir	Indicates the CIR rate limit in the access egress direction for the aggregate of the SAP queues
Ing Shaper Group	Indicates the ingress shaper group for the SAP
Egr Shaper Group	Indicates the egress shaper group for the SAP
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 (Ingress)	
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

Label	Description
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
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 (Priority)	Index of the ingress QoS queue of this SAP, where n is the index number
Off. Combined	Combined total number of high-priority and low-priority packets or octets offered to the forwarding engine
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

Label	Description
For. OutProf	Number of out-of-profile packets or octets (rate above CIR) forwarded
Ingress Queue <i>n</i> (Profile)	Index of the ingress QoS queue of this SAP, where <i>n</i> is the index number
Off. ColorIn	Number of packets or octets colored as in-profile for the SAP (offered)
Off. ColorOut	Number of packets or octets colored as out-of-profile for the SAP (offered)
Off. Uncolor	Number of packets or octets that are unprofiled for the SAP (offered)
Dro. ColorOut	Number of packets or octets colored as out-of-profile that were dropped for the SAP
Dro. ColorIn/Uncolor	Number of packets or octets that were colored as in-profile or unprofiled that were dropped for the SAP
For. InProf	Number of forwarded packets or octets colored as in-profile (FC profile set to "in" or "no profile" and rate less than or equal to CIR)
For. OutProf	Number of forwarded packets or octets that were colored as out-of-profile (FC profile set to "out" or "no profile" and rate above CIR)
Egress Queue <i>n</i>	Index of the egress QoS queue of the SAP, where <i>n</i> 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
Egress TD Profile	Profile ID of the traffic descriptor applied to the egress SAP
Alarm Cell Handling	Indicates that OAM cells are being processed

Label	Description
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

Label	Description
Local Proxy ARP	Indicates whether local proxy ARP is enabled or disabled
Policies	Specifies the policy statements applied to proxy ARP
ICMP Details	
Unreachables	Maximum number of ICMP destination unreachable messages that the IP interface issues 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 issues 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 44: SAP ID configurations](#) 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 110: Service ID ARP field descriptions](#) describes the fields.

Output example

```
*A:ALU-2# show service id 4 arp
=====
ARP Table
=====
IP Address      MAC Address      Type    Expiry    Interface    SAP
-----
10.2.3.3        Other    00h00m00s to Internet    n/a
=====
*A:ALU-2#
```

Table 110: Service ID ARP field descriptions

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 111: Service ID base field descriptions](#) describes the fields.

Output example

```
*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 111: Service ID base field descriptions

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

Label	Description
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

Output

The following output is an example of service ID DHCP statistics information, and [Table 112: Service ID DHCP statistics field descriptions](#) describes the fields.

Output example

```
*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 112: Service ID DHCP statistics field descriptions

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
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 113: Service ID DHCP summary field descriptions](#) describes the fields.

Output example

```
*A:ALU-2 show service id 4 dhcp summary
=====
DHCP Summary, service 4
=====
Interface Name      Arp    Used/
SapId/Sdp           Populate Provided      Info  Admin
-----
to Internet         No      0/0                Keep   Down
-----
Interfaces: 1
=====
*A:ALU-2
```

Table 113: Service ID DHCP summary field descriptions

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

Label	Description
	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
displays the interface information associated with the specified IPv4 or IPv6 address

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 114: Service ID interface field descriptions](#) describes the fields.

Output example

```
*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
10.2.3.3/24                n/a
-----
Interfaces : 1
=====
*A:ALU-2
```

Table 114: Service ID interface field descriptions

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

ip-transport

Syntax

ip-transport *ipt-id* [**detail** | **statistics**]

Context

show>service>id

Description

This command displays information for a specified IP transport subservice within this IES service. If no IP transport subservice is specified, summary information is displayed for all IP transport subservices associated with the IES service.

Parameters

- ipt-id

the physical port associated with the IP transport subservice, in the format *slot/mda/port.channel*
- detail

displays detailed information for the specified IP transport subservice
- statistics

displays statistical information for the specified IP transport subservice

Output

The following output is an example of IP transport subservice summary information for a specified service, and [Table 115: Service IP transport summary field descriptions](#) describes the fields.

Output example

```
*A:ALU-12# show service id 100 ip-transport
=====
IP Transport (Summary), Service 100
=====
IptId      LocalIP      LocalPort Proto RemHost DSCP FC FltrUnkn Adm  Opr
-----
1/2/4.1    192.168.1.1  3000      tcp    2       ef   ef disabled Up   Up
-----
Entries found: 1
=====
*A:ALU-12#
```

Table 115: Service IP transport summary field descriptions

Label	Description
IP Transport (Summary), Service x	
IptId	The IP transport subservice physical port identifier
LocalIP	The IP address (IPv4) that is used for the local host
LocalPort	The port number that is used by remote hosts to establish TCP/UDP sessions to the local host
Proto	The protocol type that is used for all sessions to and from the local host (either TCP or UDP)
RemHost	The number of remote hosts associated with the IP transport subservice
DSCP	The DSCP name used to mark the DSCP field in IP transport packets
FC	The FC name used for IP transport packets
FltrUnkn	Indicates whether the filter-unknown-host command is enabled or disabled on the IP transport subservice

Label	Description
Adm	The administrative state of the IP transport subservice
Opr	The operational state of the IP transport subservice
Entries found:	The number of IP transport subservices associated with this service

The following output is an example of detailed information for a specified IP transport subservice within a specified service, and [Table 116: Service IP transport detailed field descriptions](#) describes the fields.

Output example

```
*A:7705:Dut-C# show service id 100 ip-transport 1/2/4.1 detail
=====
IP Transport
=====
Service Id       : 100 (IES)
IP Transport Id  : 1/2/4.1
Description      : (Not Specified)
Admin State     : Up
Oper State      : Up
Oper Flags      : (Not Specified)
Local IP Address : 192.168.1.1
Local IP Protocol : tcp
DSCP            : ef
FC              : ef
TCP Inact Timeout : 30
TCP Max Retries  : 5
TCP Retry Interval : 5
Num Remote Hosts : 0
Last Mgmt Change : 12/07/2016 16:48:22
Last Oper Change : 12/07/2016 16:48:22
-----
IP Transport Accumulated Statistics
-----
Known Remote Hosts
Packets sent           : 44
Characters sent        : 66000
Packets received       : 67
Characters received    : 51114
Connections            : 2
  To                   : 2
  From                 : 0
Connection retries     : 20
Connection failures    : 2
Currently connected    : 0
Unknown Remote Hosts
Packets sent           : 119
Characters sent        : 178500
Packets received       : 153
Characters received    : 116039
Successful connections from : 2
Rejected due to unknown host filter : 37
Rejected due to out of resources : 0
Inactivity timeouts    : 0
Last RemIp:RemPort    : 192.168.1.7:4001
Currently connected    : 0
Dropped packets due to no remote hosts : 27
=====
*A:7705:Dut-C#
```

Table 116: Service IP transport detailed field descriptions

Label	Description
IP Transport	
Service Id	The ID that identifies the service (the service type is shown in brackets)
IP Transport Id	The physical port identifier for this IP transport subservice
Description	The description associated with this IP transport subservice
Admin State	The administrative state of this IP transport subservice
Oper State	The operational state of this IP transport subservice
Oper Flags	The operational flags associated with this IP transport subservice
Local IP Address	The IP address (IPv4) that is used for the local host
Local Port Number	The port number that is used by remote hosts to establish TCP/UDP sessions to the local host
Local IP Protocol	The protocol type that is used for all sessions to/from the local host (either TCP or UDP)
DSCP	The DSCP name used to mark the DSCP field in IP transport packets
Filter Unknown Host	Indicates whether the filter-unknown-host command is enabled or disabled for this IP transport subservice
FC	The FC name used for IP transport packets
Profile	The profile marking for the IP transport packets (in or out)
TCP Inact Timeout	The configured inactivity timeout value for TCP connections
TCP Max Retries	The configured maximum retry value for TCP connections
TCP Retry Interval	The configured retry interval value for TCP connections
Num Remote Hosts	The number of remote hosts associated with this IP transport subservice
Last Mgmt Change	The date and time of the most recent management-initiated change to this IP transport subservice
Last Oper Change	The date and time of the most recent operational status change for this IP transport subservice
IP Transport Accumulated Statistics	

Label	Description
Known Remote Hosts	
Packets sent	The number of packets sent to the host
Characters sent	The number of data characters sent to the host
Packets received	The number of packets received from the host
Characters received	The number of data characters received from the host
Connections To From	The number of connections to and from the host
Connection retries	The number of connection retries to the host
Connection failures	The number of connection failures to the host
Currently connected	The number of hosts currently connected
Unknown Remote Hosts	
Packets sent	The number of packets sent to the host
Characters sent	The number of data characters sent to the host
Packets received	The number of packets received from the host
Characters received	The number of data characters received from the host
Successful connections from	The number of successful connections from the host
Rejected due to unknown host filter	The number of rejected connection attempts from the host due to the filter-unknown-host command being enabled
Rejected due to out of resource	The number of connection attempts from the host that were rejected due to the unavailability of resources
Inactivity timeouts	The number of connections from the host that timed out due to inactivity
Last RemIp:RemPort	The IP address (IPv4) and port number used by the host for the last connection
Currently connected	The number of hosts that are currently connected
Dropped packets due to no remote hosts	The number of packets dropped due to no hosts being connected

remote-host

Syntax

remote-host *host-id* [**detail** | **statistics**]

Context

show>service>id>ip-transport

Description

This command displays information for a specified remote host within this IP transport subservice within this service. If no remote host is specified, summary information is displayed for all remote hosts within this IP transport subservice.

Parameters

host-id

the remote host identifier

Values 1 to 2147483647 or a name string up to 64 characters long

detail

displays detailed information for a specified remote host

statistics

displays summary information for a specified remote host

Output

The following output is an example of IP transport subservice remote host summary information when no remote host is specified, and [Table 117: IP transport subservice remote host summary field descriptions](#) describes the fields.

Output example

```
*A:7705:Dut-C# show service id 100 ip-transport 1/6/4.1 remote-host
=====
IPT Remote Host (Summary), Service 100 IPT 1/6/4.1
=====
RemId      RemIp:RemPort      Rcvd Chars  Sent Chars  Drop Chars  State
          Rcvd Pkts  Sent Pkts  Drop Pkts  Up Time
-----
1          192.168.1.1:3000   2555        2044        0           connected
                    5           4           0           00h01m21s
(unknown)  192.168.1.7:4000   0           2044        5110        connected
                    0           4           10          00h00m42s
-----
Number of known remote hosts: 1
Number of unknown remote hosts: 1
Total entries found: 2
=====
*A:7705:Dut-C#
```


Table 117: IP transport subservice remote host summary field descriptions

Label	Description
IP Remote Host (Summary), Service x IPT x/x/x.x	
RemId	The remote host identifier
RemIp:RemPort	The IP address (IPv4) and port number used by the remote host
Rcvd Chars	The number of data characters received from the remote host
Sent Chars	The number of data characters sent to the remote host
Drop Chars	The number of data characters destined for the remote host that were dropped
State	The operational state of the packet transport session connection to the remote host
Rcvd Pkts	The number of packets received from the remote host
Sent Pkts	The number of packets sent to the remote host
Drop Pkts	The number of packets destined for the remote host that were dropped
Up Time	The amount of time that the remote host has been connected
Number of known remote hosts	The number of known remote hosts associated with the IP transport subservice
Number of unknown remote hosts	The number of unknown remote hosts associated with the IP transport subservice
Total entries found	The total number of hosts associated with the IP-Transport subservice

The following output is an example of IP transport subservice detailed information for a specified remote host, and [Table 118: IP transport subservice remote host detailed field descriptions](#) describes the fields.

Output example

```
*A:7705:Dut-C# show service id 100 ip-transport 1/2/4.1 remote-host 1 detail
=====
IPT Remote Host
=====
Service Id       : 100 (IES)
IP Transport Id  : 1/2/4.1
Remote Host Id   : 1
Name             : (Not Specified)
Description      : (Not Specified)
IP Address       : 192.168.1.6          Port Number      : 4000
Last Mgmt Change : 12/07/2016 16:48:44  Up Time          : 00h01m44s
Session State    : connected
```

```

Last Connect      : successful
-----
IPT Remote Host Statistics
-----
Sent Pkts       : 134          Sent Chars       : 201000
Dropped Pkts    : 0           Dropped Chars    : 0
Rcvd Pkts       : 267         Rcvd Chars       : 201000
Session information
  Connections    : 2
  To             : 1
  From           : 1
  Connection retries : 0
  Connection failures : 0
  Closed by far end : 1
  Inactivity timeouts : 0
=====
*A:7705:Dut-C#

```

Table 118: IP transport subservice remote host detailed field descriptions

Label	Description
IP Remote Host	
Service Id	The ID that identifies the service (the service type is shown in brackets)
IP Transport Id	The physical port identifier for the IP transport subservice
Remote host Id	The host identifier associated with this remote host
Name	The name associated with this remote host
Description	The description associated with this remote host
IP Address	The IP address associated with this remote host
Port Number	The port number associated with this remote host
Last Mgmt Change	The date and time of the most recent management-initiated change to this remote host
Session State	The operational state of the packet transport session to this host
Up Time	The amount of time that this remote host has been connected
Last Connect	Indicates whether the last connection attempt to this remote host was successful or unsuccessful
IP Remote Host Statistics	
Sent Pkts	The number of packets sent to this remote host
Sent Chars	The number of data characters sent to this remote host

Label	Description
Dropped Pkts	The number of packets destined for this remote host that were dropped
Dropped Chars	The number of data characters destined for this remote host that were dropped
Rcvd Pkts	The number of packets received from this remote host
Rcvd Chars	The number of data characters received from this remote host
Session information	
Connections To From	The number of connections to and from the host
Connection retries	The number of connection retries to the host
Connection failures	The number of connection failures to this host
Closed by far end	The number of connections closed by the far end
Inactivity timeouts	The number of connections that were timed out due to inactivity

macsec

Syntax

macsec

Context

show>service>id

Description

This command displays MACsec security information for the specified service.

Output

The following output is an example of MACsec information, and [Table 119: Service-ID MACsec field descriptions](#) describes the fields.

Output example

```
*A:ALU-12# show service id 1 macsec
=====
MACsec (Summary), Service 1
=====
SAP          MACsec    MACsec    Encap    CA tags   CA-name
              port      sub-port  match    in-clear
```

1/1/3	1/1/3	1	all	0	ca1
*A:ALU-12#					

Table 119: Service-ID MACsec field descriptions

Label	Description
SAP	The service SAP
MACsec port	The port enabled for MACsec
MACsec sub-port	The subport enabled for MACsec
Encap match	The traffic encapsulation type to match: all traffic, untagged-only traffic, single-tag or dot1q traffic, double-tag or QinQ traffic
CA tags in-clear	The number of tags in clear text for this CA
CA-name	The name of the MACsec connectivity association for this SAP

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 is displayed. See [Table 44: SAP ID configurations](#) 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 57: Service-ID SAP field descriptions](#) in [VLL services command reference](#) for field descriptions.

Output example

*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
Ing Shaper Group  : n/a                     Egr Shaper Group  : n/a
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 (Ingress)		
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 120: Service ingress label field descriptions](#) describes the fields.

Output example

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: Service ingress label field descriptions

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

ip-transport-using

Syntax

ip-transport-using [**ip-transport** *ipt-id*]

Context

show>service

Description

This command displays IP transport subservice information for a specified port. If no port is specified, the command displays a summary of all IP transport subservices defined for the IES service.

Parameters

ipt-id

the physical port associated with the IP transport subservice, in the format *slot/mda/port.channel*

Output

The following output is an example of **ip-transport-using** information, and [Table 121: IP transport-using field descriptions](#) describes the fields.

Output example

```
*A:ALU-48# show service ip-transport-using
=====
IP Transports
=====
IptId          SvcId      Type  Adm  Opr
-----
1/2/4.1        100        IES   Up   Up
-----
Entries found: 1
-----
*A:ALU-48#
```

Table 121: IP transport-using field descriptions

Label	Description
IP Transports	
IptId	The IP transport subservice physical port identifier
SvcId	The service identifier
Type	The type of service
Adm	The administrative state of the IP transport subservice
Opr	The operational state of the IP transport subservice
Entries found	The number of IP transport subservices using this service

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}

sap-using [ingress | egress] shaper-group *shaper-group-name*

Context

show>service

Description

This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The 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 44: SAP ID configurations](#) for a full list of SAP IDs.

ip-address

displays the interface information associated with the specified IPv4 or IPv6 address

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

shaper-group

specifies the shaper group for which to display matching SAPs

Output

The following output is an example of service SAP-using information, and [Table 122: Service SAP-using field descriptions](#) describes the fields.

Output example

```
*A:ALU-48# show service sap-using
```

```
=====
```

```
Service Access Points
```

```
=====
```

PortId	SvcId	Ing. QoS	Ing. Fltr	Egr. QoS	Egr. Fltr	Adm	Opr
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. QoS	Ing. Fltr	Egr. QoS	Egr. Fltr	Adm	Opr
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#

```

```

*A:7705custDoc:Sar18>show>service# sap-using ingress shaper-group test_sg1
=====
Service Access Points Using Ingress Shaper Group "test_sg1"
=====
PortId          SvcId      Scheduler  Shaper Policy  Opr
                Mode
-----
1/2/1           30         4-priority test_shaper_policy  Down
-----
Number of SAPs : 1
=====
*A:Sar18 Dut-B>config>service>epipe>sap>ingress#

```

Table 122: Service SAP-using field descriptions

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
Shaper Policy	Identifies the shaper policy that the shaper group belongs to
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 specific 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 123: Service service-using field descriptions](#) describes the fields.

Output example

```
*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 123: Service service-using field descriptions

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

6.7.2.6 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

the IP interface name

ip-address

the IPv4 or IPv6 address

ip-transport

Syntax

ip-transport *ipt-id*

Context

clear>service>id

Description

This command clears configured information pertaining to a specified IP transport subservice.

If no port identifier is specified, information is cleared for all IP transport subservices.

Parameters

ipt-id

the IP transport subservice physical port identifier, in the format *slot/mda/port.channel*

remote-host

Syntax

remote-host *host-id*

Context

clear>service>id>ip-transport

Description

This command clears configured information pertaining to a specified remote host assigned to this IP transport subservice.

Parameters

host-id

the remote host identifier

Values 1 to 2147483647 or a name string up to 64 characters long

statistics

Syntax

statistics

Context

clear>service>id>ip-transport

clear>service>id>ip-transport>remote-host

Description

This command clears statistics-related information pertaining to all configured IP transport subservices or to all configured remote hosts for a specified IP transport subservice.

6.7.2.7 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

7 VPRN services

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](#)

7.1 VPRN service overview

Topics in this section include:

- [Routing prerequisites](#)
- [BGP support](#)
- [IPSec support](#)
- [Security zones and VPRN](#)
- [Static one-to-one NAT and VPRN](#)
- [Unicast and multicast address translation](#)
- [Route distinguishers](#)
- [Route target constraint](#)
- [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. Because 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 correct 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.

The following figure 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 105: Virtual private routed network



20949

VPRN is supported on the following:

- any DS1/E1 port on the 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card
- the 16-port T1/E1 ASAP Adapter card
- the 32-port T1/E1 ASAP Adapter card
- the Packet Microwave Adapter card
- any V.35 port on the 12-port Serial Data Interface card, version 3
- any T1/E1 port on the 7705 SAR-M
- any T1/E1 port on the 7705 SAR-A
- any T1/E1 port on the 4-port T1/E1 and RS-232 Combination module
- 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 port on the 7705 SAR-A
- any port on the 7705 SAR-X
- any Ethernet port on the 7705 SAR-M
- any Ethernet port on the 7705 SAR-Ax
- any Ethernet port on the 7705 SAR-Wx
- any Ethernet or T1/E1 port on the 7705 SAR-H
- any Ethernet port on the 7705 SAR-Hc

Ports must be in access mode.

7.1.1 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 be used

7.1.2 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 about BGP route tunnels, see the 7705 SAR 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.

VPRN BGP is configured through the **config>service>vpn>bgp** context. Global BGP is configured through the **config>router>bgp** context.

7.1.2.1 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 /IPv6 and VPN-IPv4/VPN-IPv6 routes. The supported VPRN scenarios are described in the following table.

Table 124: 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
IPv6 (ingress PE)	VPN-IPv6 route with next hop A resolved by a GRE, LDP, RSVP, or BGP tunnel	VPN-IPv6 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
MPLS (egress PE)	IPv6 route with next hop A resolved by an IPv6 route	VPN-IPv6 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, see the 7705 SAR Routing Protocols Guide, "BGP FRR with prefix-independent convergence".

7.1.2.2 BGP next-hop resolution and peer tracking

The 7705 SAR can attach a route policy to the BGP next-hop resolution process and can allow a route policy to be associated with the optional BGP peer-tracking function. These two features are supported for VPRN BGP service.

BGP next-hop resolution determines the best matching route (or tunnel) for the BGP next-hop address and uses information about this resolving route when running the best-path selection algorithm and programming the forwarding table. Attaching a policy to BGP next-hop resolution provides additional control over which IP routes in the routing table can become resolving routes. Similar flexibility and control is available for BGP peer tracking, which is an optional feature that allows a session with a BGP neighbor to be taken down if there is no IP route to the neighbor address or if the best matching IP route is rejected by the policy.

Use the following CLI syntax to configure next-hop resolution and peer-tracking policies:

CLI syntax:

```
config>service>vprn>bgp
  next-hop-resolution
    policy policy-name
    no policy
  peer-tracking-policy policy-name
  no peer-tracking-policy
```

For details, see the "Route policies for BGP next-hop resolution and peer tracking" section in the 7705 SAR Router Configuration Guide.

7.1.3 IPSec support

The 7705 SAR supports IPSec and IPSec tunnels, where VPRN or IES 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 public-side IPSec tunnel SAP is **tunnel-1.public:tag**. The *sap-id* for a private-side IPSec tunnel SAP is **tunnel-1.private:tag**.

For more information, see the [IPSec](#) chapter in this guide.

7.1.4 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, see the 7705 SAR 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 2 endpoints or 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. Layer 3 zones support both NAT and firewall security policies. Layer 2 zones support only firewalls. 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>epipe>zone**
- **config>service>vpls>zone**
- **config>service>vprn>zone**

Layer 2 and Layer 3 firewalls share system resources; that is, they share the maximum number of policies, profiles, and session ID space supported by the system.

A zone is created by adding at least one Layer 2 endpoint or Layer 3 interface to the zone configuration. Multiple zones can be created within each Layer 3 service or within the router context. Layer 2 services support only one zone. Layer 2 endpoints or Layer 3 interfaces from different services cannot be grouped into a single common zone. The following table lists the supported interfaces and endpoints that can be added to zones in each CLI context for NAT or firewall.

Table 125: Security zone interfaces and endpoints per context

CLI context	Interface/endpoint type	NAT	Firewall
Router	Layer 3	✓	✓
Epipe	SAP		✓
	Spoke-SDP termination		✓
VPLS	SAP		✓
	Spoke-SDP termination		✓
	Mesh SDP		✓
	EVPN		
VPRN	SAP	✓	✓
	Spoke-SDP termination	✓	✓
	IPSec private	✓	✓
	IPSec public	✓	
	Routed VPLS	✓	✓



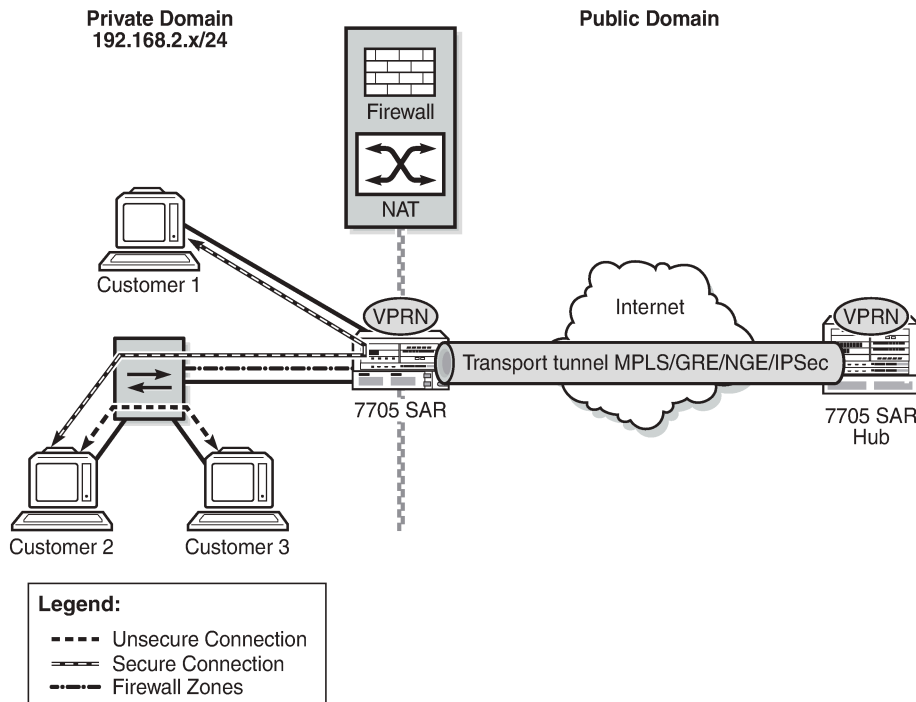
Note:

- NAT and firewalls are not supported on V.35 ports on the 12-port Serial Data Interface card
- A group of endpoints used for pseudowire redundancy cannot be added to a zone configured under an Epipe.

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.

The following figure 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 106: Firewall protection for the network core



25129

A security zone can also be created with spoke SDPs or autobind tunnels that have been configured for a VPRN service using MP-BGP. For autobind tunnels, the **auto-bind-tunnel resolution-filter** type can be set to gre, ldp, rsvp, sr-isis, sr-ospf, or sr-te. When a zone contains spoke SDPs or autobind MPLS tunnels, it cannot contain any other type of interface. A zone configured with spoke SDPs or autobind 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. After 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 autobind tunnel to another within the same zone is not firewalled.

7.1.5 Static one-to-one NAT and VPRN

With static one-to-one NAT, NAT is performed on packets traveling from an inside (private) interface to an outside (public) interface or from an outside interface to an inside interface. Static one to-one NAT can be applied to a single IP address or a subnet of IP addresses and is performed on the IP header of a packet, not on the UDP/TCP port.

Mapping statements, or entries, can be configured to map an IP address range to a specific IP address. The direction of the NAT mapping entry dictates whether NAT is performed on a packet source IP address

or subnet or on a packet destination IP address or subnet. The 7705 SAR supports inside mapping entries that map an inside IP address range to an outside IP address range sequentially.

With an inside mapping entry, the following points apply:

- Packets that originate from an inside interface and are destined for an inside interface are forwarded without any NAT being applied.
- If there is a matching one-to-one NAT mapping entry, packets that originate from an inside interface and are destined for an outside interface undergo static one-to-one NAT where NAT changes the source IP address of the packet IP header. The packet is forwarded whether or not a NAT mapping entry is found unless the **drop-packets-without-nat-entry** command is enabled. When a mapping entry is not found and the **drop-packets-without-nat-entry** command is enabled, the packet is not forwarded.
- If there is a matching one-to-one NAT mapping entry, packets that originate from an outside interface and are destined for an inside interface undergo static one-to-one NAT where NAT changes the destination IP address of the packet IP header. The packet is forwarded whether or not a NAT mapping entry is found unless the **drop-packets-without-nat-entry** command is enabled. When a mapping entry is not found and the **drop-packets-without-nat-entry** command is enabled, the packet is not forwarded.
- Packets that originate from an outside interface and are destined for an outside interface are forwarded without any NAT being applied.

Static one-to-one NAT is supported in the GRT and in VPRNs. For more information about static one-to-one NAT, see the 7705 SAR Router Configuration Guide, "Static one-to-one NAT".

For VPRNs, one-to-one NAT can be configured between an inside interface and a outside MP-BGP MPLS transport tunnel interface. Policies should be used to not leak the inside interface/IP address via MP-BGP to the peer. Policies should be used to leak the NAT routes to the MP-BGP peer.

The following table lists the types of outside and inside interfaces that are supported in a VPRN for one-to-one NAT.

Table 126: VPRN interfaces supported for static one-to-one NAT

VPRN interface type	Outside	Inside
SAP interface	✓	✓
R-VPLS interface	✓	✓
Layer 3 spoke SDP interface	✓	✓
IPSec private interface	✓	✓
Autobind GRE/MPLS (MP-BGP), where MPLS includes segment routing, LDP, and RSVP	✓	

7.1.6 Unicast and multicast address translation

The 7705 SAR supports unicast-to-multicast address translation and multicast-to-multicast address translation.

For unicast-to-multicast translation, the 7705 SAR translates the destination IP address of the unicast flow to a multicast group. For multicast-to-multicast translation, the 7705 SAR acts as a host to upstream (S,G)s and performs address translation to the downstream (S,G).

Unicast and multicast address translation is supported on the following adapter cards and platforms:

- on the 7705 SAR-8 Shelf V2 and the 7705 SAR-18:
 - 2-port 10GigE (Ethernet) Adapter card
 - 6-port Ethernet 10Gbps Adapter card
 - 8-port Gigabit Ethernet Adapter card, version 3
 - 10-port 1GigE/1-port 10GigE X-Adapter card, version 2 (supported on the 7705 SAR-18 only)
- 7705 SAR-Ax
- 7705 SAR-H
- 7705 SAR-Hc
- 7705 SAR-Wx
- 7705 SAR-X


Unicast and multicast address translation is supported in the GRT and in VPNs. For VPNs, IPv4 addressing on SAP-to-SAP connections is supported.

For more information about unicast-to-multicast address translation and multicast-to-multicast address translation, see the 7705 SAR Routing Protocols Guide, "Unicast and multicast address translation".

7.1.7 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 implementation supports the three (3) Type-Value combinations, as defined in RFC 2547bis. The following figure illustrates the RD structure.

Figure 107: Route distinguisher structure



20950

The three Type-Value combinations supported are described in the following table.

Table 127: 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 field	Value field	Notes
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

7.1.7.1 PE-to-CE route exchange

Routing information between the provider edge (PE) and customer edge (CE) can be exchanged by the following methods:

- EBGp (for IPv4 and IPv6 address families)
- OSPF
- OSPFv3
- RIP
- static routes

Each protocol provides controls to limit the number of routes learned from each CE router.

7.1.7.1.1 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.

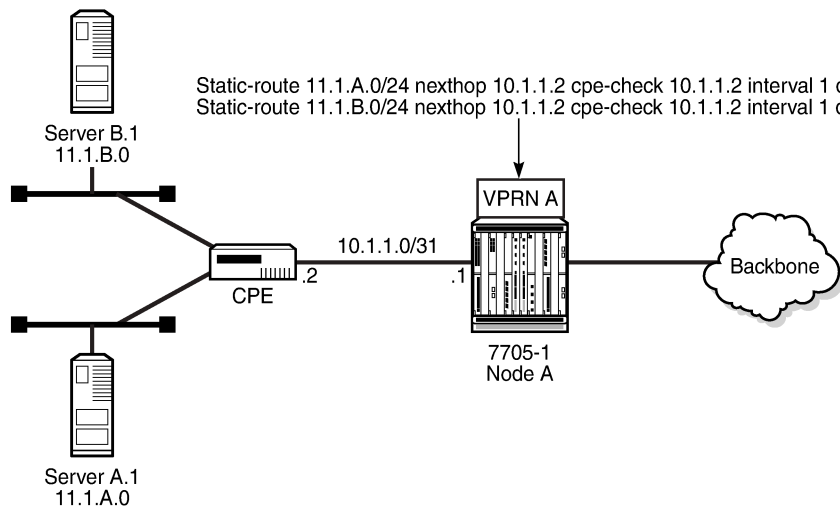
A route belonging to a VPRN must use the protocol owner, VPN-IPv4, to denote that it is a VPRN route. This can be used within the route policy match criteria.

7.1.7.1.2 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.

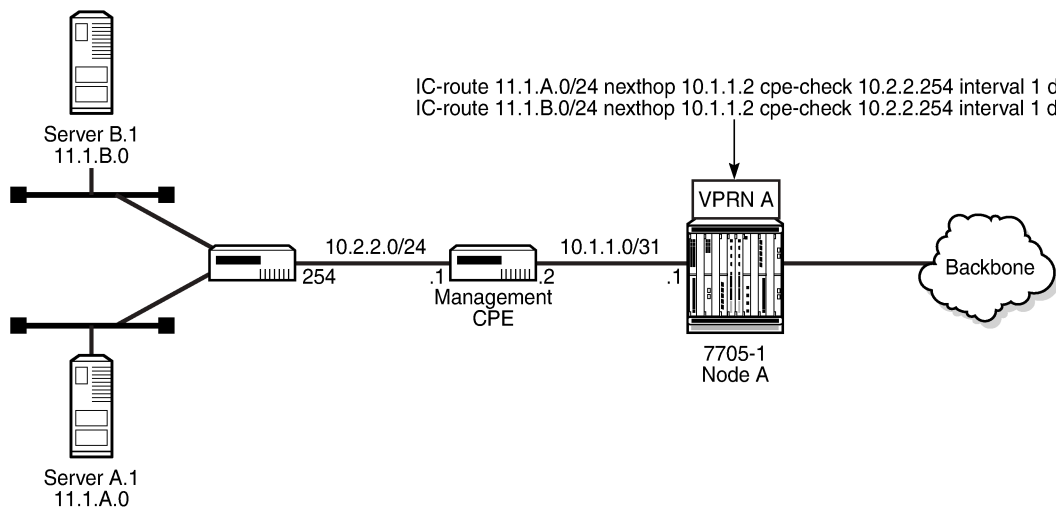
The following figures illustrate the use of CPE connectivity check in directly connected and multiple-hop connected routes.

Figure 108: Directly connected IP target



20951

Figure 109: Multiple hops to IP target



20952

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.

7.1.8 Route target constraint

Route target constraint (RTC) is a mechanism that allows a router to advertise route target membership information to its BGP peers to indicate interest in receiving only VPN routes tagged with specific route target extended communities. Upon receiving RTC route information, peers restrict the advertised VPN routes to only those requested, minimizing control plane load in terms of protocol traffic and potentially reducing RIB memory usage.

The route target membership information is carried using MP-BGP, using an AFI value of 1 and SAFI value of 132. The NLRI of an RTC route encodes an Origin AS and a route target extended community using prefix type encoding with host bits after the prefix-length set to zero.

In order for two routers to exchange route target membership NLRI, they must advertise the corresponding AFI and SAFI to each other during capability negotiation. The use of MP-BGP means route target membership NLRI are propagated, loop-free, within an autonomous system and between autonomous systems, using well-known BGP route selection and advertisement rules.

Route target constrained route distribution and outbound route filtering (ORF) both allow routers to advertise which route target extended communities they want to receive in VPN routes from peers. RTC, however, is more widely supported, is simpler to configure, and its distribution scope is not limited to a direct peer.

7.1.8.1 Configuring the route target address family

RTC is supported only by the base router BGP instance. When the **family** command in the BGP router group or neighbor CLI context includes the **route-target** keyword, the RTC capability is negotiated with the associated set of EBGP and IBGP peers.

ORF and RTC are mutually exclusive for a particular BGP session. The CLI does not attempt to block the configuration of both ORF and RTC, but if both capabilities are enabled for a session, the ORF capability will not be included in the OPEN message sent to the peer.

7.1.8.2 Originating RTC routes

When the base router has one or more RTC peers (BGP peers with which the RTC capability has been successfully negotiated), one RTC route is created for each route target extended community imported into a locally configured Layer 3 VPN service. These imported route targets are configured in the following contexts:

- **config>service>vprn**
- **config>service>vprn>mvpn**

By default, RTC routes are automatically advertised to all RTC peers without the need for an export policy to explicitly accept them. Each RTC route has a prefix, a prefix length, and path attributes. The prefix value is the concatenation of the origin AS (a 4-byte value representing the 2-octet or 4-octet AS of the originating router, as configured using the **config>router>autonomous-system** command) and 0 or 16 to 64 bits of a route target extended community encoded in one of the following formats: 2-octet AS specific extended community, IPv4 address specific extended community, or 4-octet AS specific extended community.

A router may be configured to send the default RTC route to a group or neighbor with the **default-route-target** CLI command. The default RTC route is a special route that has a prefix length of zero. Sending

the default RTC route to a peer conveys a request to receive all VPN routes from that peer, whether they match the route target extended community. The default RTC route is typically advertised by a route reflector to PE clients. Advertising the default RTC route to a peer does not suppress other more specific RTC routes from being sent to that peer. A received default RTC route is never propagated to other routers.

7.1.8.3 Receiving and readvertising RTC routes

All received RTC routes that are considered valid are stored in the RIB-In. RTC routes are considered invalid and treated as withdrawn if the prefix length is configured to be any of the following:

- 1 to 31
- 33 to 47
- 48 to 96 and the 16 most-significant bits are not 0x0002, 0x0102, or 0x0202

If multiple RTC routes are received for the same prefix value (same NLRI), then standard BGP best-path selection procedures are used to determine the best route. The propagation of the best path installs RIB-Out filter rules as it travels from one router to the next, and this process creates an optimal VPN route distribution tree rooted at the source of the RTC route.

The best RTC route per prefix is readadvertised to RTC peers based on the following rules:

- The best path for a default RTC route (prefix length 0, origin AS only with prefix length 32, or origin AS plus 16 bits of a route target type with prefix-length 48) is never propagated to another peer.
- A PE with only IBGP RTC peers that is not a route reflector or an ASBR does not readvertise the best RTC route to any RTC peer, due to standard IBGP split-horizon rules.
- A route reflector that receives its best RTC route for a prefix from a client peer readadvertises that route (subject to export policies) to all its client and non-client IBGP peers (including the originator), per standard route reflector operation. When the route is readadvertised to client peers, the route reflector sets the ORIGINATOR_ID to its own router ID and modifies the NEXT_HOP to be its local address for the sessions (for example, system IP).
- A route reflector that receives its best RTC route for a prefix from a non-client peer readadvertises that route (subject to export policies) to all its client peers, per standard route reflector operation. If the route reflector has a non-best path for the prefix from any of its clients, it advertises the best of the client-advertised paths to all non-client peers.
- An ASBR that is not a PE or a route reflector, that receives its best RTC route for a prefix from an IBGP peer, readadvertises that route (subject to export policies) to its EBGP peers. The NEXT_HOP and AS_PATH of the re-advertised route are modified per standard BGP rules. No aggregation of RTC routes is supported.
- An ASBR that is not a PE or a route reflector, that receives its best RTC route for a prefix from an EBGP peer, readadvertises that route (subject to export policies) to its EBGP and IBGP peers. The NEXT_HOP and AS_PATH of the re-advertised route are modified per standard BGP rules. No aggregation of RTC routes is supported.



Note: These advertisement rules do not handle hierarchical route reflector topologies properly. This is a limitation of the current RTC standard.

7.1.8.4 Using RTC routes

In general, the best VPN route for every prefix or NLRI in the RIB is sent to every peer supporting the VPN address family. Export policies may be used to prevent some prefixes or NLRIs from being advertised to specific peers. These export policies may be configured statically, or created dynamically by using ORF or RTC with a peer. ORF and RTC are mutually exclusive for a session.

When RTC is configured on a session that also supports VPN address families using route targets (VPN-IPv4, VPN-IPv6, or MVPN-IPv4), the advertisement of the VPN routes is affected as follows:

- When the session comes up, the advertisement of the VPN routes is delayed for a short while to allow RTC routes to be received from the peer.
- After the initial delay, the received RTC routes are analyzed and acted upon. If $S1$ is the set of routes previously advertised to the peer and $S2$ is the set of routes that should be advertised based on the most recent received RTC routes, then:
 - the set of routes in $S1$ but not in $S2$ are withdrawn immediately (subject to the minimum route advertisement interval (MRAI))
 - the set of routes in $S2$ but not in $S1$ are advertised immediately (subject to the MRAI)
- If a default RTC route is received from a peer $P1$, the set of VPN routes that are advertised to $P1$ are routes that:
 - are eligible for advertisement to $P1$ per BGP route advertisement rules
 - have not been rejected by manually configured export policies
 - have not been advertised to the peer

This applies whether or not $P1$ advertised the best route for the default RTC prefix. A default RTC route is a route with any of the following:

- NLRI length = zero
- NLRI value = origin AS and NLRI length = 32
- NLRI value = {origin AS+0x0002 | origin AS+0x0102 | origin AS+0x0202} and NLRI length = 48
- If an RTC route for prefix A (origin-AS = $A1$, RT = $A2/n$, $n > 48$) is received from an IBGP peer $I1$ in autonomous system $A1$, the set of VPN routes that are advertised to $I1$ are the routes that:
 - are eligible for advertisement to $I1$ per BGP route advertisement rules
 - have not been rejected by manually configured export policies
 - carry at least one route target extended community with value $A2$ in the n most significant bits
 - have not been advertised to the peer

This applies whether or not $I1$ advertised the best route for A .

- If the best RTC route for a prefix A (origin-AS = $A1$, RT = $A2/n$, $n > 48$) is received from an IBGP peer $I1$ in autonomous system B , the set of VPN routes that are advertised to $I1$ are routes that:
 - are eligible for advertisement to $I1$ per BGP route advertisement rules
 - have not been rejected by manually configured export policies
 - carry at least one route target extended community with value $A2$ in the n most significant bits
 - have not been advertised to the peer

This applies only if $I1$ advertised the best route for A .

- If the best RTC route for a prefix A (origin-AS = A1, RT = A2/n, $n > 48$) is received from an EBGp peer E1, the set of VPN routes that are advertised to E1 are the routes that:
 - are eligible for advertisement to E1 per BGP route advertisement rules
 - have not been rejected by manually configured export policies
 - carry at least one route target extended community with value A2 in the n most significant bits
 - have not been advertised to the peer

This applies only if E1 advertised the best route for A.

7.1.9 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 a destination IP address in the packet header matches any system IP address configured under **grt-lookup** with a GRT **static-route-entry** set to the system IP address specified under **vprn>static-route-entry>grt**, the packet is extracted to the CSM for processing. If the **vprn>grt-lookup>enable-grt>allow-local-management** command is not enabled, the packet is routed using the 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 **allow-local-management** 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 configured under **vprn>static-route-entry>grt**. If the command **allow-local-management** is enabled under the VPRN instance, a packet arriving on a transport tunnel will be extracted to the CSM before hitting the blackhole route. In this case, the only effect of the blackhole route will be to advertise the system IP address to the far-end peer. If the command **allow-local-management** is not enabled, packet forwarding will be the default forwarding mode; that is, all packets destined for the system IP address will be blackholed because of the static route configuration.

For MP-BGP VPRNs, when the command **allow-local-management** 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 local IP address and destination IP address of either the far-end NSP NFM-P 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 FIB using the **export-grt** command. The route installed in the GRT FIB 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.

The following table lists the management protocols supported by IPv4 and IPv6 GRT in the reverse and forward directions.

Table 128: IPv4 and IPv6 GRT-supported management protocols


Protocol	Reverse direction (toward the 7705 SAR)	Forward direction (from the 7705 SAR)
FTP Passive and Active	✓ ¹	✓ ²
SFTP	✓ ¹	
NTP		✓ ³
RADIUS	✓	
SCP	✓ ¹	✓ ³
SNMP	✓	
SNMP Trap		✓
SSH	✓ ¹	✓ ³
TACACS+	✓	
Telnet	✓ ¹	✓ ³
TWAMP ⁴	✓ ⁵	✓ ⁶
TWAMP Light	✓ ⁵	✓ 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)

The following figure shows an example of IPv4 in-band management of the 7705 SAR and the switches behind it by the NSP NFM-P and a TACACS+ server. In the example, an IPSec tunnel is being used as the VPRN to transport the management traffic via a secure and encrypted medium over the public internet.

Figure 110: 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 **allow-local-management** 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 FIB using the **export-grt** command.
- The above configuration will add route 192.168.1.0/24 to the GRT FIB 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 be resolved by the VPRN IPsec tunnel.

The following figure shows an example of IPv6 in-band management of the 7705 SAR and the switches behind it by the NSP NFM-P and a TACACS+ server.

Figure 111: 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 **allow-local-management** 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, 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 FIB using the **export-grt** command.

The above configuration adds route fd00:1:129:1::/64 to the GRT FIB 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.

7.2 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](#)
- [Multicast VPN \(MVPN\)](#)

- [VPRN autobinding tunnels](#)
- [Spoke SDPs](#)
- [Spoke-SDP termination to VPRN](#)
- [IPv6 on virtual private edge router](#)
- [IPv6 over IPv4 LAN-to-LAN IPSec tunnels](#)
- [Bandwidth optimization for low-speed links](#)
- [Support for NTP](#)

7.2.1 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:

- Unnumbered interfaces (see [Unnumbered interfaces](#))
- DHCP options (see [DHCP and DHCPv6](#))
- Local DHCP server options (see [Local DHCP or DHCPv6 server](#))
- IPSec tunnel interfaces (see [IPSec support](#))
- IPCP options (see [IPCP](#))
- ICMP options (see [Troubleshooting and fault detection services](#))
- VRRP options (see [VRRP on VPRN interfaces](#))

Configuration options found on core IP interfaces not supported on VPRN IP interfaces are:

- NTP broadcast receipt

7.2.1.1 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.

7.2.1.2 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.

Because 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-A, 7705 SAR-Ax, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-M, 7705 SAR-Wx, and 7705 SAR-X. The 7705 SAR supports a maximum of 62 servers per node on the 7705 SAR-8 Shelf V2 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.

7.2.1.2.1 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 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 about DHCP and DHCPv6, see the 7705 SAR Router Configuration Guide, "DHCP and DHCPv6".

7.2.1.2.1.1 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. Because 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.

7.2.1.2.1.1.1 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 more 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

7.2.1.2.1.2 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 (see the 7705 SAR 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.

7.2.1.2.1.2.1 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*.

7.2.1.2.2 Local DHCP or 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.

7.2.1.2.2.1 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.

7.2.1.3 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.

7.2.1.4 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 IGP and BGP protocols, including 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 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.

7.2.1.5 VRRP on VPRN interfaces

VRRP can be implemented on VPRN service interfaces to participate as part of a virtual router instance. This implementation prevents a single point of failure by ensuring access to the gateway address, which is configured on all VPRN service interfaces in the VRRP. VRRPv3 can also be implemented on VPRN service interfaces, including r-VPLS interfaces for VPRN.

The 7705 SAR supports VRRPv3 for IPv4 and IPv6 as described in RFC 5798. Within a VRRP router, the virtual routers in each of the IPv4 and IPv6 address families are a domain unto themselves and do not overlap.



Note: VRRPv3 for IPv6 is not supported on a Layer 3 spoke-SDP termination.

For information about VRRP and VRRP VPRN service interface parameters, as well as the configuration parameters of VRRP policies, see the “VRRP” section in the 7705 SAR Router Configuration Guide. CLI command descriptions for VRRP policies are also given in the 7705 SAR Router Configuration Guide.

For CLI command descriptions related to VPRN service interfaces, see [VPRN services command reference](#).

7.2.1.6 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. Layer 4 port attributes and the TEID attribute in the hashing algorithm can be configured with the **l4-load-balancing** and **teid-load-balancing** commands in the **config>service>vprn>interface** context. Configuration of the **l4-load-balancing** command at the interface level overrides the

system-level settings for the specific interface. The **teid-load-balancing** command can only be configured at the interface level.

The system IP address can be included in or excluded from the hashing algorithm with the system-level **system-ip-load-balancing** command.

For more information about IP ECMP, see the 7705 SAR Router Configuration Guide, "Static routes, dynamic routes, and ECMP".

7.2.1.7 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 about proxy ARP, see the 7705 SAR Router Configuration Guide, "Proxy ARP".

7.2.1.8 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.

7.2.1.9 Hold up and hold down timers for IP interfaces

The 7705 SAR allows timers to be configured on a VPRN or IES IPv4 or IPv6 interface or on the base router to keep the IP interface in an operationally up or down state for a specified time beyond when it should be declared operationally up or down. The timers are configured at the VPRN service level using the **config>service>vprn>interface>hold-time>up/down** commands.

An **init-only** option enables the **down** delay to be applied only when the IP interface is first configured or after a system reboot. See [Internet enhanced service](#) for information about how to configure the **hold-time** command on IES interfaces. See the 7705 SAR Router Configuration Guide for information about how to configure the **hold-time** command at the router level.

The configuration causes the system to delay sending notifications of any state change associated with the IP interface until the timer has expired.



Note: The **up** and **down** delay timers in the CLI are dynamic text fields; the fields are only displayed in the **show router interface detail** command output if they are configured. The field showing the time remaining is displayed only if the timer is actively counting down. If both **up** and **down** timers are configured, the field displayed depends on the current operational state of the interface. For example, if the interface is operationally down, the configured hold **down** time is displayed.

7.2.2 SAPs

Topics in this section include:

- [Encapsulations](#)
- [QoS policies](#)
- [QinQ \(VPRN\)](#)
- [Filter policies on a VPRN SAP](#)

VPRN service also supports SAPs for IPsec tunnels (see [IPsec support](#)).

7.2.2.1 Encapsulations

The following SAP encapsulations are supported on the 7705 SAR VPRN service:

- Ethernet null
- Ethernet dot1q
- Ethernet qinq
- PPP
- MLPPP
- MC-MLPPP
- LAG



Note: When gathering statistics on VPRN SAPs and ports, the SAP ingress counters may be different from the port counters as the SAP takes into account encapsulation headers. For Ethernet-encapsulated SAPs, counters can be adjusted by configuring the packet byte offset, which adjusts the packet size that schedulers, shapers, and the SAP counters take into account by offsetting the configured number of bytes. For information about packet byte offset, see the 7705 SAR Quality of Service Guide, "Packet byte offset".

7.2.2.2 QoS policies

For each instance of VPRN service, QoS policies can be applied to the ingress and egress VPRN interface SAPs.

At VPRN access ingress, traffic can be classified as unicast or multicast traffic types. In a VPRN access ingress QoS policy, users can create queues that map to forwarding classes. For each forwarding class, traffic can be assigned to a queue that is configured to support unicast, multicast, or both. As shown in the

following example, for fc "af", both unicast and multicast traffic use queue 2, and for fc "l2", only multicast traffic uses queue 3.

```
configure qos sap-ingress qos 2 create
    queue 1 create
    exit
    queue 2 create
    exit
    queue 3 create
    exit
    fc "af" create
        queue 2
        multicast-queue 2
    exit
    fc "l2" create
        multicast-queue 3
    exit
```

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 must 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, see the 7705 SAR Quality of Service Guide.

7.2.2.2.1 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 Quality of Service Guide.

7.2.2.3 QinQ (VPRN)

VPRN supports QinQ functionality. For details, see [QinQ support](#).

7.2.2.4 Filter policies on a VPRN SAP

IPv4 and IPv6 filter policies can be applied to ingress and egress VPRN SAPs.

See the 7705 SAR Router Configuration Guide, "Filter policies", for information about configuring IP filters.

7.2.3 PE-to-CE routing protocols

The 7705 SAR supports the following PE-to-CE routing protocols for VPRN service:

- EBGp (for IPv4 and IPv6 address families)
- OSPF
- OSPFv3 (for IPv6 address families)
- RIP
- static routes

EBGP is supported within both the router context and VPRN service context. Both OSPF and OSPFv3 are supported within the router context as well as within the VPRN service context; however, there are some minor differences in the command sets depending on the context.

7.2.3.1 Using OSPF or OSPFv3 in IP VPNs

Using OSPF or OSPFv3 as a PE-to-CE routing protocol allows the version of 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/OSPFv3 learned routes as OSPF/OSPFv3 externals**

This feature uses OSPF or OSPFv3 as the protocol between the PE and CE routers; however, instead of transporting the OSPF/OSPFv3 LSA information across the IP-VPN, the OSPF/OSPFv3 routes are "imported" into MP-BGP as AS externals. As a result, other OSPF- or OSPFv3-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/OSPFv3 instance**

This occurs if the OSPF or OSPFv3 route type (in the OSPF/OSPFv3 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/OSPFv3 domain ID BGP extended community attribute carried with the VPN route.

- **sham links**

A sham link is a logical PE-to-PE unnumbered point-to-point interface that rides over the PE-to-PE transport tunnel. A sham link can be associated with any area and can appear as an intra-area link to CE routers attached to different PEs in a VPN.

Sham links are not supported on OSPFv3.

- **import policies**

By default, OSPF imports all the routes advertised via LSAs. Import policies allow routes that match a certain criteria, such as neighbor IP addresses, to be rejected. Users must use caution when applying import policies, since not using certain routes may result in network stability issues.

Import policies are supported within the VPRN context and the base router context. Import policies are not supported on OSPFv3.

7.2.3.1.1 DN bit

When a type 3 LSA is sent from a PE router to a CE router, the DN bit in the LSA options field is set. This ensures that if any CE router sends this type 3 LSA to a PE router, the PE router will not redistribute it further.

When a PE router needs to distribute to a CE router a route that comes from a site outside the CE router's OSPF/OSPFv3 domain, the PE router presents itself as an autonomous system boundary router (ASBR) and distributes the route in a type 5 LSA. The DN bit must be set in these LSAs to ensure that they will be ignored by any other PE routers that receive them.

DN bit loop avoidance is also supported.

7.2.3.2 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 about TTL security, see the 7705 SAR Routing Protocols Guide, "TTL security".

7.2.4 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 required. When the **auto-bind-tunnel** command is used, all services traverse the same LSPs and do not allow alternate tunneling mechanisms (such as GRE) or the ability to configure sets of LSPs with bandwidth reservations for specific customers, as is available with explicit SDPs for the service.


7.2.5 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.

7.2.6 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. The following figure shows an example of RIP metric propagation in a VPRN across two autonomous systems.

Figure 112: 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.

7.2.7 Multicast VPN (MVPN)

The two main multicast VPN (MVPN) service implementations are the *draft-rosen-vpn-mcast* and the next-generation multicast VPN (NG-MVPN).

The 7705 SAR supports NG-MVPNs, which use BGP for customer-multicast (C-multicast) signaling.

The V.35 ports on the 12-port Serial Data Interface card, version 3 do not support multicast VPN.

The 7705 SAR conforms to the relevant sections of the following RFCs related to MVPNs:

- RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*
- RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root* (only as source router)
- RFC 6513, *Multicast in MPLS/BGP IP VPNs*

This section includes information about the following topics:

- [Multicast in IP-VPN applications](#)
- [MVPN building blocks](#)
- [Provider tunnel support](#)
- [Inter-AS and intra-AS solutions](#)
- [NG-MVPN non-segmented inter-AS solution](#)
- [Mrinfo and Mtrace](#)
- [Multicast-only fast reroute](#)

- [mLDP point-to-multipoint support](#)
- [mLDP fast upstream switchover](#)
- [Multicast Source Discovery Protocol](#)

7.2.7.1 Multicast in IP-VPN applications

This section focuses on IP-VPN multicast functionality. As a prerequisite for MVPN, readers should be familiar with the "IP multicast" material in the 7705 SAR Routing Protocols Guide, where multicast protocols (PIM, IGMP, and MLD) are described.

Applications for this feature include enterprise customers implementing a VPRN solution for their WAN networking needs, video delivery systems, and customer applications that include stock-ticker information as well as financial institutions for stock and other types of trading data.

Implementation of next-generation VPRN (NG-VPRN) requires the separation of the provider's core multicast domain from customer multicast domains, and the customer multicast domains from each other.

[Figure 113: Multicast in an IP-VPN application](#) shows an example of multicast in an IP-VPN application and shows the following domains:


- provider's domain
 - core routers (1 through 4)
 - edge routers (5 through 10)
- customers' IP-VPNs, each having their own multicast domain
 - VPN-1 (CE routers 12, 13, and 16)
 - VPN-2 (CE routers 11, 14, 15, 17, and 18)

In this VPRN multicast example, VPN-1 data generated by the customer behind router 16 is multicast by PE router 9 to PE routers 6 and 7 for delivery to CE routers 12 and 13, respectively. VPN-2 data generated by the customer behind router 15 is forwarded by PE router 8 to PE routers 5, 7, and 10 for delivery to CE routers 11 and 18, 14, and 17, respectively.

The demarcation points for these domains are in the PEs (routers 5 through 10). The PE routers participate in both the customer multicast domain and the provider multicast domain. The customer CEs are limited to a multicast adjacency with the multicast instance on the PE, where the PE multicast instance is specifically created to support that specific customer IP-VPN. As a result, customers are isolated from the provider core multicast domain and other customer multicast domains, while the provider core routers only participate in the provider multicast domain and are isolated from all customer multicast domains.

The PE for a customer's multicast domain becomes adjacent to the CE routers attached to that PE and to all other PEs that participate in the IP-VPN (customer) multicast domain. The adjacencies are set up by the PE that encapsulates the customer multicast control data and the multicast streams inside the provider's multicast packets. The encapsulated packets are forwarded only to the PE nodes that are attached to the same customer's edge routers as the originating stream and are part of the same customer VPRN. This process prunes the distribution of the multicast control and data traffic to the PEs that participate in the customer's multicast domain.

Figure 113: Multicast in an IP-VPN application



25930

7.2.7.2 MVPN building blocks

This section includes information about the following topics:

- [PMSI](#)
- [MVPN using BGP control plane](#)
- [Auto-discovery](#)
- [PE-CE multicast protocols and services](#)
- [PE-PE transmission of C-multicast routing using BGP](#)
- [PE-PE multicast protocols](#)
- [PE-PE multicast data transmission](#)

7.2.7.2.1 PMSI

A provider-multicast (P-multicast) service interface (PMSI), described in RFC 6513, refers to an abstract service in the service provider's core network that can take a packet from one PE, belonging to one MVPN, and deliver a copy of the packet to some or all of the other PEs supporting that MVPN.

The most common PMSI uses a multicast distribution tree (MDT). An MDT is a point-to-multipoint traffic path that is instantiated using forwarding table entries that support packet replication. For example, an MDT forwarding entry would specify the incoming interface—where the node expects to receive a packet flowing up or down the MDT—and the set of outgoing interfaces that each receive a copy of the packet. The MDT forwarding state can be set up using an IP multicast signaling protocol such as PIM, or an MPLS protocol such as multicast LDP (mLDP) or RSVP-TE.

The 7705 SAR supports mLDP PMSI only.

This section includes information about the following topics:

- [PMSI types](#)
- [Creating a PMSI](#)

7.2.7.2.1.1 PMSI types

There are two types of PMSIs: inclusive and selective.

An inclusive PMSI (I-PMSI) includes all of the PEs supporting an MVPN. A selective PMSI (S-PMSI) includes a subset of the PEs supporting an MVPN (that is, an S-PMSI is a subset of an I-PMSI). An MVPN can have more than one S-PMSI.

7.2.7.2.1.1.1 Inclusive PMSIs (I-PMSIs)

An MVPN has one I-PMSI. The I-PMSI carries MVPN-specific control information between the PEs of the MVPN. In the 7705 SAR implementation, by default, all C-multicast flows use the I-PMSI. This minimizes the number of PE router states in the service provider core, but wastes bandwidth because a C-multicast flow on an I-PMSI is delivered to all PEs in the MVPN, even when only a subset of the PEs have receivers for the flow. To reduce wasted bandwidth, a service provider can migrate the C-multicast flow from the I-PMSI to an S-PMSI that includes only the PEs with receivers interested in that (S,G) flow.

On a 7705 SAR, migration of a C-multicast-flow from I-PMSI to S-PMSI can be configured to be initiated automatically by the PE closest to the source of the C-multicast-flow, where the migration trigger is based on the data rate of the flow. Migration occurs when the data rate exceeds the configured threshold.

7.2.7.2.1.1.2 Selective PMSIs (S-PMSIs)

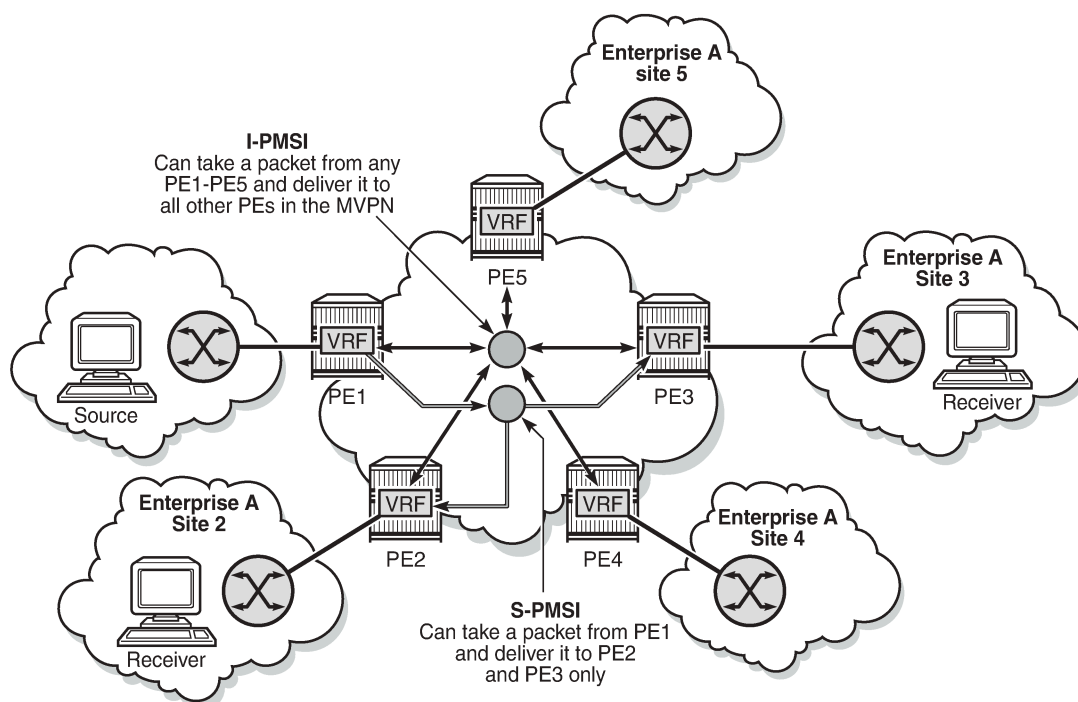
A selective PMSI is one that includes a subset of the PEs supporting an MVPN. Each MVPN can have zero or more S-PMSIs. As stated above, the transition from I-PMSI to S-PMSI is triggered when the data rate exceeds the user-configured threshold.

7.2.7.2.1.1.3 I-PMSI versus S-PMSI

The following figure illustrates the difference between an I-PMSI and an S-PMSI. In the figure, the arrowheads indicate send and receive capabilities supported by the PMSI. Two-way arrows imply sender and receiver transmissions, and one-way arrows imply sender-only or receiver-only transmission.

In the figure, all the VRFs that are part of the MVPN domain receive PDUs from the I-PMSI MDT entries, whether or not the VRFs are configured to receive the PDUs. When the traffic for an (S,G) exceeds the configured data rate threshold, the multicast tree for that (S,G) switches from an I-PMSI to an S-PMSI. Each (S,G) has its own S-PMSI tree built when the threshold for that (S,G) has been exceeded.

Figure 114: I-PMSI and S-PMSI



25927

7.2.7.2.1.2 Creating a PMSI

The 7705 SAR supports multicast LDP (mLDP) only as the mechanism to build a PMSI tunnel.


When the C-multicast protocol for a leaf node in an MVPN initiates a multicast request, it triggers mLDP to generate an LSP. The C-multicast protocol can be IGMP or PIM, which is configured under the VPRN service.

Multicast LDP tries to resolve the source address of (S,G) by looking up the source (S) in the routing table manager (RTM). If there is a resolution, mLDP generates a FEC toward the source (S). The mLDP FEC, as described in RFC 6388, contains the root node's system IP address and an opaque value. The opaque value contains a point-to-multipoint LSP ID, which uniquely identifies the point-to-multipoint LSP in the context of the root node. [Figure 115: P2MP FEC and MP LDP opaque value as per RFC 6388](#) illustrates a point-to-multipoint FEC element and an opaque value.

The P2MP ID is generated on the root node and is advertised to the leaf node via a BGP MVPN address family route update (see [Figure 116: BGP MVPN address family updates](#)).

In the following figure, the point-to-multipoint FEC element contains of the address of the root of the point-to-multipoint LSP and an opaque value. The opaque value consists of one or more LDP multiprotocol (MP) opaque value elements. The opaque value is unique within the context of the root node, and for the 7705 SAR it is the P2MP ID. The combination of "Root Node Address Type", "Root Node Address", and "Opaque Value" uniquely identifies a point-to-multipoint LSP within the MPLS network.

Figure 115: P2MP FEC and MP LDP opaque value as per RFC 6388




MP LDP Opaque Value RFC 6388

25928

The following figure shows:

- the BGP MVPN address family update (#1), which contains the unique P2MP ID
- the LDP FEC (#2), which is generated when the C-multicast IGMP or PIMv4 prompts the LDP to generate the mLDP FEC from the leaf to the root node
- the opaque value (basic type) (#3), which is encoded into the mLDP FEC and contains the P2MP ID

Figure 116: BGP MVPN address family updates



25929

7.2.7.2.2 MVPN using BGP control plane

To communicate auto-discovery routes and C-multicast signaling, add the **mvpn-ipv4** address family to the BGP address family configuration. For more information, see the “Configuring BGP address families” section in the 7705 SAR Routing Protocols Guide.

The 7705 SAR MVPN implementation is based on NG-VPN standards and supports the following features:

- auto-discovery
 - PE-PE transmission of C-multicast routing using BGP
 - IPv4
 - use of mLDP for S-PMSIs used as PMSIs
 - inter-AS with direct VRF connect (option A) and non-segmented mLDP option C
 - inter-AS as root router only (option B)
 - inter-AS/intra-AS root, ASBR/ABR, leaves and transit router for non-segmented option C using mLDP
- For non-segmented inter-AS/inter-area option C, an ASBR/ABR cannot be the root of the P2MP tree for mLDP.

7.2.7.2.3 Auto-discovery

Auto-discovery for multicast VPN refers to the process by which a PE with a multicast MVPN service dynamically learns about the other PEs supporting the same MVPN.

The basic auto-discovery function is to discover the identity of all other PEs in the MVPN. This information is essential for setting up the I-PMSI.

Advanced auto-discovery functions are:

- discovering the subsets of PEs in the MVPN that are interested in receiving a specific multicast flow
- discovering autonomous system border routers (ASBRs) in other ASs that are interested in receiving a specific multicast flow
- discovering C-multicast sources that are actively sending traffic across the service provider backbone (in a PMSI)
- discovering bindings between multicast flows and PMSIs

The MVPN standards define two different options for MVPN auto-discovery, BGP and PIM. The 7705 SAR only uses MP-BGP for auto-discovery and S-PMSI signaling.

With BGP auto-discovery, MVPN PEs advertise special auto-discovery routes to their peers using multiprotocol extensions to BGP.

Using BGP for auto-discovery does not imply that BGP must be used for C-multicast signaling, nor does it impose any restrictions about the technology used to tunnel MVPN packets between PEs in the service provider backbone.

7.2.7.2.3.1 MVPN membership auto-discovery using BGP

BGP-based auto-discovery is performed by referring to a multicast VPN address family (for example, **mvpn-ipv4**). Any PE that attaches to an MVPN must issue a BGP update message containing an NLRI in the address family, along with a specific set of attributes.

The PE router uses route targets to specify an MVPN route import and export policy. The route target may be the same target as the one used for the corresponding unicast VPN, or it may be a different target. For a specific MVPN, the PE router can specify separate import route targets for sender sites and receiver sites.

The route distinguisher (RD) that is used for the corresponding unicast VPN can also be used for the MVPN.

In addition, the bindings of C-trees to P-tunnels are discovered using BGP S-PMSI auto-discovery routes.

7.2.7.2.4 PE-CE multicast protocols and services

A PE with an MVPN service must learn about the networks and multicast receivers located beyond the CE devices at the MVPN customer site. Typically, IGMP or a PIM protocol is used by the CE to inform the PE that it (the CE) wants to receive a particular multicast flow because it has downstream receivers of that flow. The 7705 SAR supports IGMP (versions 1, 2, and 3), PIM sparse mode (SM), and PIM source-specific multicast (SSM) as the CE-to-PE protocol.

The use of PIM as the CE-to-PE protocol requires that the PE learn about networks beyond the CE so that the PE can appropriately select the correct upstream next hop for sending PIM join and prune

messages. The join and prune messages normally follow the reverse path of unicast data traffic and establish the required multicast forwarding state in the PE, CE, and other PIM routers at the customer site. The reachability of networks beyond the CE can be learned through a routing protocol such as OSPF, RIP, or BGP, or it can be configured statically when static routes or multiprotocol BGP are used between the CE and a 7705 SAR PE.

Layer 2 services, such as routed VPLS, can be used to snoop IGMP from the VPLS access interface and translate IGMP to PIM on the PE-CE Layer 3 interface.

The 7705 SAR supports IPv4 PE-CE protocols (for example, IGMP and PIM).

7.2.7.2.5 PE-PE transmission of C-multicast routing using BGP

MVPN C-multicast routing information is exchanged between PEs by using C-multicast routes that are carried by the MVPN NLRI.

7.2.7.2.6 PE-PE multicast protocols

When a PE gets a request for a multicast flow from a connected CE in an MVPN, it must convey that request to the PE closest to the source of the multicast traffic. In order for a PE to know that another PE is closer to a source, unicast routes must be exchanged between the PEs. This is done by exchanging VPN-IP routes using multiprotocol BGP (MP-BGP). The VPN-IP routes exchanged for this purpose may carry additional information as compared to VPN-IP routes used only for unicast routing. In particular, when NG-MVPN signaling is used, as per RFC 6513, a route that is a candidate for upstream multicast hop (UMH) selection carries two additional BGP extended communities: a source-AS extended community and a VRF route import extended community.

C-multicast signaling is the signaling of joins and prunes from a PE that is connected to a site with receivers of a multicast flow to another PE that is closest to a sender of the multicast flow. Similar to auto-discovery, the MVPN standards allow either PIM or BGP to be used for C-multicast signaling.

The 7705 SAR uses only BGP for C-multicast signaling and multicast route advertisement between PEs.

When BGP is used for C-multicast signaling, a PE announces its desire to join a source C-tree by announcing a special source-join BGP NLRI using BGP multiprotocol extensions. The source-join BGP NLRI has the same AFI and SAFI as the BGP auto-discovery routes (described in RFC 6513). When a PE wants to leave an inter-site source tree, it withdraws the source-join BGP NLRI that it had previously advertised. A PE directs a source-join BGP NLRI to a specific upstream PE—the one it determines to be closest to the source—by including the VRF route import extended community associated with that upstream PE; other PEs may receive the source-join BGP NLRI, but do not import and use it.

Using C-multicast signaling protocols with BGP means that each MVPN PE typically has a small number of BGP sessions (for example, two interior border gateway protocol (IBGP) sessions with two route reflectors in the local AS).

The use of BGP minimizes the control plane load, but may lead to slightly longer join and leave latencies than is the case for the faster recovery of lost BGP messages by the TCP layer underlying the BGP sessions. This is due to the route reflector propagating join and prune messages from downstream PEs to upstream PEs.

7.2.7.2.7 PE-PE multicast data transmission

A PMSI can be built on one or more point-to-point, point-to-multipoint, or multipoint-to-multipoint tunnels that carry customer multicast packets transparently through the service provider core network. The MVPN standards provide several technology options for PMSI tunnels:

- RSVP-TE LSP (point-to-point and point-to-multipoint)
- mLDP LSP (point-to-point, point-to-multipoint, and multipoint-to-multipoint)
- GRE tunnel (point-to-point and point-to-multipoint)

Only point-to-multipoint mLDP as a PMSI tunnel is supported.

The 7705 SAR platforms support the following transport options:

- I-PMSI – mLDP point-to-multipoint LSPs
- S-PMSI – mLDP point-to-multipoint LSPs

7.2.7.3 Provider tunnel support

The following provider tunnel features are supported:

- I-PMSI
- S-PMSI

Topics in this section include:

- [Point-to-multipoint I-PMSI and S-PMSI](#)
- [Point-to-multipoint LDP I-PMSI and S-PMSI](#)
- [Point-to-multipoint LSP S-PMSI](#)
- [MVPN sender-only and receiver-only](#)

7.2.7.3.1 Point-to-multipoint I-PMSI and S-PMSI

BGP C-multicast signaling must be enabled for an MVPN instance to use point-to-multipoint mLDP to create an I-PMSI or S-PMSI.

By default, all PE nodes participating in MVPN receive data traffic over an I-PMSI. Optionally, for efficient data traffic distribution, S-PMSIs can be used to send traffic to PE nodes that have at least one active receiver connected.

Only one unique multicast flow is supported over each mLDP point-to-multipoint LSP S-PMSI.

The number of S-PMSIs that can be initiated per MVPN instance is set by the **maximum-p2mp-spmsi** command. A point-to-multipoint LSP S-PMSI cannot be used for more than one (S,G) stream when the maximum number of S-PMSIs per MVPN is reached. Multicast flows that cannot switch to an S-PMSI remain on the I-PMSI.

7.2.7.3.2 Point-to-multipoint LDP I-PMSI and S-PMSI

A point-to-multipoint LDP LSP as an inclusive or selective provider tunnel is available with BGP NG-MVPN only. A point-to-multipoint LDP LSP is set up dynamically from leaf nodes upon auto-discovery of leaf PE

nodes that are participating in multicast VPN. Each LDP I-PMSI or S-PMSI LSP can be used with a single MVPN instance only.

The **multicast-traffic** command (under **config>router>ldp>interface-parameters>interface**) must be configured on a per-LDP interface basis to enable a point-to-multipoint LDP setup. Point-to-multipoint LDP must also be configured as an inclusive or selective provider tunnel on a per-MVPN basis. Use the **mldp** command (under **provider-tunnel>inclusive** or **>selective**) to dynamically initiate a point-to-multipoint LDP LSP to leaf PE nodes learned via NG-MVPN auto-discovery signaling. S-PMSI is for efficient data distribution and is optional.

7.2.7.3.3 Point-to-multipoint LSP S-PMSI

NG-MVPN allows the use of a point-to-multipoint LDP LSP as the S-PMSI. An S-PMSI is generated dynamically, based on the user-configured traffic bandwidth threshold for a number of multicast flows. Use the **data-threshold** command (under **provider-tunnel>selective**) to set the bandwidth threshold.

In MVPN, the root node PE discovers all the leaf PEs via I-PMSI auto-discovery routes. All multicast PDUs traverse through the I-PMSI until the configured threshold is reached on the root node. When the configured threshold is reached on the root node, the root node signals the desire to switch to an S-PMSI via BGP signaling of the S-PMSI auto-discovery NLRI.

Because of the way that LDP normally works, mLDP point-to-multipoint LSPs are set up (unsolicited) from the leaf PEs toward the root node PE. The leaf PE discovers the root node PE via auto-discovery routes (I-PMSI or S-PMSI). The tunnel identifier carried in the PMSI attribute is used as the point-to-multipoint FEC element.

The tunnel identifier consists of the root node PE address, along with a point-to-multipoint LSP ID. The generic LSP identifier value is automatically generated by the root node PE.

7.2.7.3.4 MVPN sender-only and receiver-only

The I-PMSI can be optimized by configuring PE nodes that function as a sender-only or receiver-only node. By default, PE nodes are both sender and receiver nodes (sender-receiver).

In MVPN, by default, if multiple PE nodes form a peering within a common MVPN instance, then each PE node originates a local multicast tree toward the other PE nodes in this MVPN instance. This behavior creates an I-PMSI mesh across all PE nodes in the MVPN. Typically, a VPN has many sites that host multicast receivers only, and has a few sites that host sources only or host both receivers and sources.

MVPN **sender-only** and **receiver-only** commands allow the optimization of control-plane and data-plane resources by preventing unnecessary I-PMSI mesh setups when a PE device hosts only multicast sources or only multicast receivers for an MVPN.

For PE nodes that host only multicast sources for a VPN, operators can configure the MVPN to block those PE nodes from joining I-PMSIs that belong to other PEs in the MVPN. For PE nodes that host only multicast receivers for a VPN, operators can block those PE nodes in order to set up a local I-PMSI to other PEs in this MVPN.

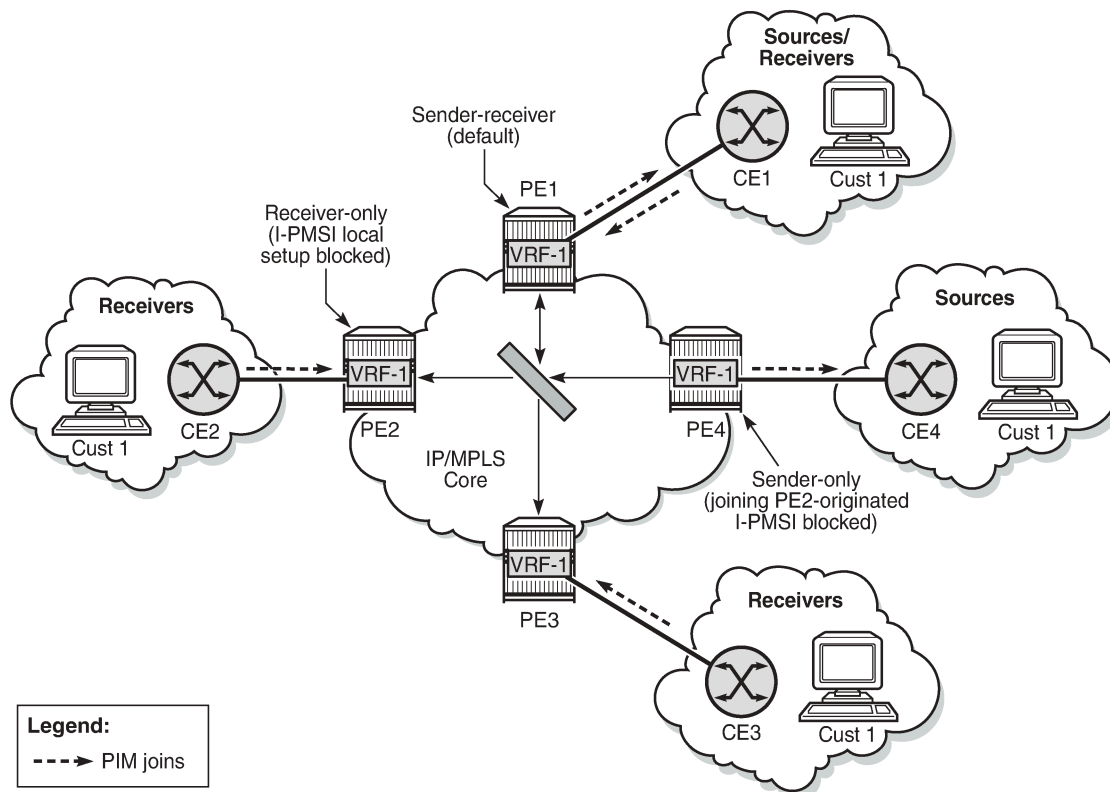
MVPN **sender-only** and **receiver-only** commands are supported with NG-MVPN using IPv4 LDP provider tunnels for both IPv4 and IPv6 customer multicast. The following figure shows a four-site MVPN with sender-only, receiver-only, and sender-receiver (default) sites.



Note: The physical location of the BSR and RP nodes when sender-only or receiver-only is enabled is important. Because the source DR sends unicast-encapsulated traffic toward the RP,

the RP needs to be at a sender-receiver or sender-only site, so that (*,G) traffic can be sent over the tunnel. The BSR needs to be deployed at the sender-receiver site. The BSR can be at a sender-only site if the RPs are at the same site. The BSR needs to receive packets from other candidate-BSR and candidate-RP nodes and also needs to send BSM packets to all other BSR and RP nodes.

Figure 117: I-PMSI sender-receiver, sender-only, and receiver-only: optimized I-PMSI mesh



25931

7.2.7.4 Inter-AS and intra-AS solutions

An MVPN service that spans more than one AS is called an inter-AS MVPN. As is the case with unicast-only IP VPN services, there are different approaches for supporting inter-AS MVPNs. Generally, the approaches belong to one of two categories:

- all P-tunnels and P-multicast trees start and end on PEs and ASBRs in the same AS
- P-tunnels and P-multicast trees extend across multiple ASs

In the first category, the P-tunnels and P-multicast trees start on PEs and ASBRs of an AS, and end on PEs and ASBRs in that same AS (extending no further). In this scenario, C-multicast traffic that must cross an AS boundary is handed off natively between the ASBRs on each side of the AS boundary.

From the perspective of each ASBR, the other ASBR is simply a collection of CEs, each reachable through separate logical connections (for example, VPRN SAPs). In this type of deployment, no auto-discovery signaling is required between the different ASs, and the exchange of C-multicast routes and C-multicast

signaling uses the same protocols and procedures as described in [PE-CE multicast protocols and services](#) for PE-CE interfaces.

In the second category, P-tunnels and P-multicast trees extend across the boundaries between different ASs. In this scenario, the PMSI extends end-to-end between the PEs of the MVPN, even when those PEs are in different ASs. ASBRs need to exchange auto-discovery information in order to determine whether:

- the neighbor AS has PEs with sites in the MVPN
- the neighbor AS is a transit node on the best path to a remote AS that has sites of the MVPN

If a P-multicast tree is used to transport the PMSI, there are two options for extending the P-tree across multiple ASs:

- **non-segmented inter-AS MDT**

The end-to-end P-tree is end-to-end between all the PEs supporting the MVPN, passing through ASBRs as necessary.

- **segmented inter-AS MDT**

The end-to-end P-tree is formed by stitching together a sub-tree from each AS. A sub-tree of an AS connects only the PEs and ASBRs of that AS. A point-to-point tunnel between ASBRs on each side of an AS boundary is typically used to stitch the sub-trees together.

Constructing and using a non-segmented inter-AS MDT is similar to constructing and using an intra-AS MDT, except that BGP auto-discovery messages are propagated by ASBRs across AS boundaries, where the BGP auto-discovery messages are I-PMSI intra-AS auto-discovery routes, despite the reference to intra-AS.

When segmented inter-AS tunnels are used for an NG-MVPN, the ASBRs configured to support that MVPN will originate inter-AS I-PMSI auto-discovery routes for that MVPN toward their external peers after having received intra-AS I-PMSI auto-discovery routes for the MVPN from one or more PEs in their own AS. The inter-AS I-PMSI auto-discovery messages are propagated through all ASs that support the MVPN (that is, through all ASs that have PEs or ASBRs for the MVPN).

When an ASBR receives an inter-AS I-PMSI auto-discovery route, and it is the best route for the NLRI, the ASBR sends a leaf auto-discovery route to the exterior Border Gateway Protocol (EBGP) peer that advertised the route. The leaf auto-discovery route is used to set up a point-to-point, one-hop MPLS LSP that stitches together the P-multicast trees of each AS.

The 7705 SAR supports inter-AS and intra-AS option A.


The 7705 SAR supports non-segmented inter-AS and intra-AS option C as an ABR or ASBR router or as a leaf router. The 7705 SAR can be part of non-segmented inter-AS and intra-AS with option B as a root node.

7.2.7.4.1 7705 SAR as source of non-segmented inter-AS or intra-AS network

The 7705 SAR can be the source of a non-segmented inter-AS or intra-AS network, as per RFC 6512 and RFC 6513.

The following figure shows inter-AS option B connectivity via non-segmented mLDP, where the 7705 SAR is acting as a root node. In inter-AS solutions, the leaf and ABR/ASBR nodes need recursive opaque FEC to route the mLDP FEC through the network.

Figure 118: Inter-AS option B: non-segmented solution



25933

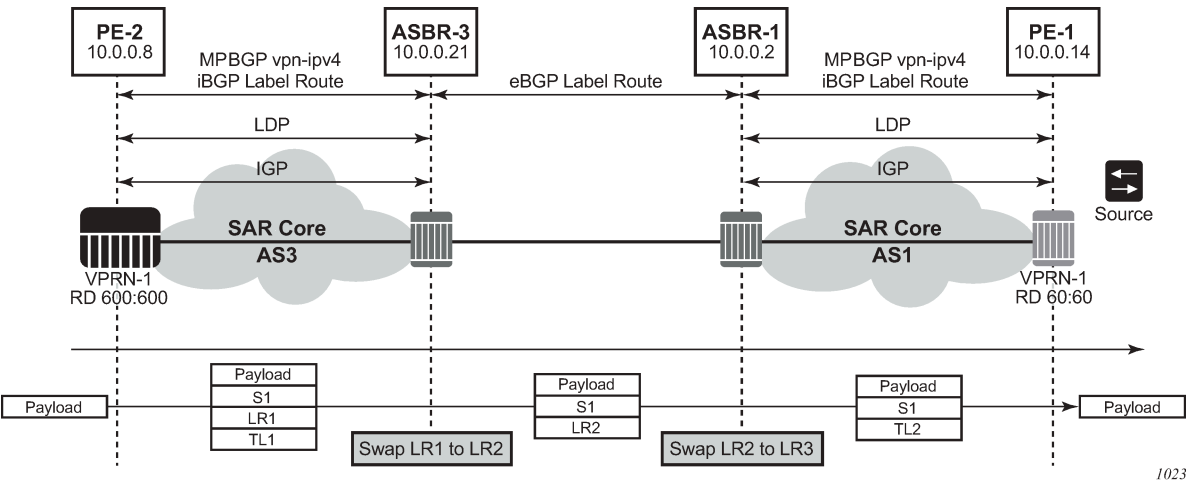
7.2.7.5 NG-MVPN non-segmented inter-AS solution

This feature allows multicast services to use segmented protocols and span them over multiple autonomous systems (ASs) in the same way as unicast services. Because IP VPN or GRT services span multiple IGP areas or multiple ASs, either for a network designed to deal with scale or as result of commercial acquisitions, operators may require inter-AS VPN (unicast) connectivity. For example, an inter-AS VPN can break the IGP, MPLS and BGP protocols into access segments and core segments, allowing higher scaling of protocols by segmenting them into their own islands. The 7705 SAR allows for a similar provisioning of multicast services and for spanning these services over multiple IGP areas or multiple ASs.

For unicast VPRNs, inter-AS or intra-AS option C breaks the IGP, BGP and MPLS protocols at ABR routers (for multiple IGP areas) and ASBR routers (for multiple ASs). At ABR and ASBR routers, a stitching mechanism of MPLS transport is required to allow transition from one segment to the next, as shown in the following figure.

In the figure, the 3107 BGP label route (LR) is stitched at ASBR1 and ASBR3. At ASBR1, the LR1 is stitched with LR2, and at ASBR3, the LR2 is stitched with TL2.

Figure 119: Unicast VPN option C with segmented MPLS



Previously, segmenting an LDP MPLS tunnel at ASBRs or ABRs was not possible with NG-MVPN. Therefore, RFC 6512 and 6513 used a non-segmented mechanism to transport the multicast data over P-tunnels end-to-end through ABR and ASBR routers. The signaling of LDP needed to be present and possible between two ABR routers or two ASBR routers in different ASs.

For unicast VPNs, it was usually preferred to only have EBGP between ASBR routers.

The 7705 SAR now has non-segmented intra-AS and inter-AS signaling for NG-MVPN. The non-segmented solution is possible for inter-ASs as option C.

7.2.7.5.1 Non-segmented inter-AS VPN option C support

The 7705 SAR supports the inter-AS option C VPN solution. option C uses recursive opaque type 7 as shown in the following table.

Table 129: Recursive opaque types

Opaque type	Opaque name	RFC	7705 SAR use
1	Basic Type	RFC 6388	VPRN Local AS
7	Recursive Opaque (Basic Type)	RFC 6512	Inter-AS option C MVPN over mLDP

In inter-AS option C, the PEs in two different ASs have their system IP addresses in the RTM, but the intermediate nodes in the remote AS do not have the system IP addresses of the PEs in their RTM. Therefore, for NG-MVPN, a recursive opaque value in mLDP FEC is needed to signal the LSP to the first ASBR in the local AS path.

For inter-AS option C, on a leaf PE, a route exists to reach the root PE system IP address. Because ASBRs can use BGP unicast routes, recursive FEC processing using BGP unicast routes (not VPN recursive FEC processing using PMSI routes) is required.

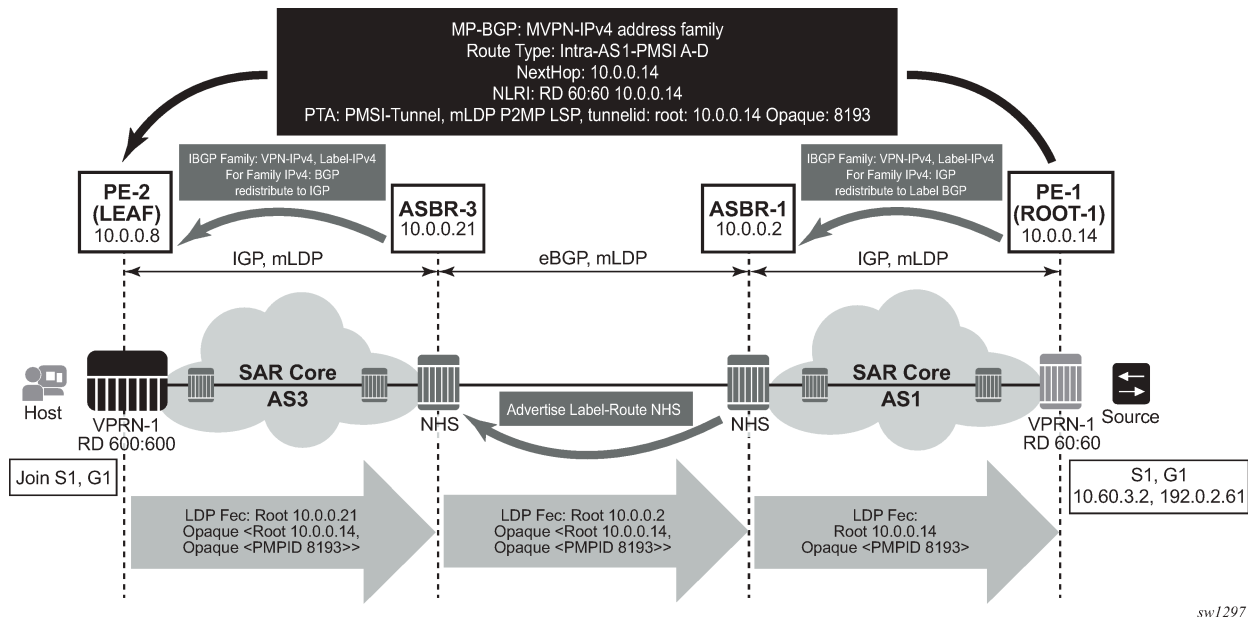
7.2.7.5.1.1 I-PMSI and S-PMSI establishment

I-PMSI and S-PMSI functionality follow RFC 6513 section 8.1.1 and RFC 6512 section 2. The VRR Route Import External community now encodes the VRF instance in the local administrator field.

Option C uses an outer opaque of type 7 and inter opaque of type 1.

The following figure shows the processing required for I-PMSI and S-PMSI inter-AS establishment.

Figure 120: Non-segmented mLDP PMSI establishment (option C)



sw/297

For non-segmented mLDP trees, A-D procedures follow those of the intra-AS model, with the exception that NO EXPORT Community must be excluded; LSP FEC includes mLDP recursive FEC (and not VPN recursive FEC).

For I-PMSI on inter-AS option C:

- A-D routes are not installed by ASBRs and next-hop information is not changed in MVPN A-D routes
- BGP labeled routes are used to provide inter-domain connectivity on remote ASBRs

On receipt of an intra-AS I-PMSI A-D route, PE2 resolves PE1's address (N-H in PMSI route) to a labeled BGP route with a next hop of ASBR3 because PE1 is not known via IGP. PE2 sources an mLDP FEC with a root node of ASBR3 and an opaque value, shown below, containing the information advertised by PE1 in the I-PMSI A-D route.

PE-2 LEAF FEC: {Root = ASBR3, Opaque Value: {Root: ROOT-1, Opaque Value: P2MP-ID xx}}

When the mLDP FEC arrives at ASBR3, it notes that it is the identified root node, and that the opaque value is a recursive opaque value. ASBR3 resolves the root node of the recursive FEC (ROOT-1) to a labeled BGP route with the next hop of ASBR1 because PE-1 is not known via IGP. ASBR3 creates a new mLDP FEC element with a root node of ASBR1 and an opaque value that is the received recursive opaque value.

ASBR3 FEC: {Root: ASBR1, Opaque Value: {Root: ROOT-1, Opaque Value: P2MP-ID xx}}

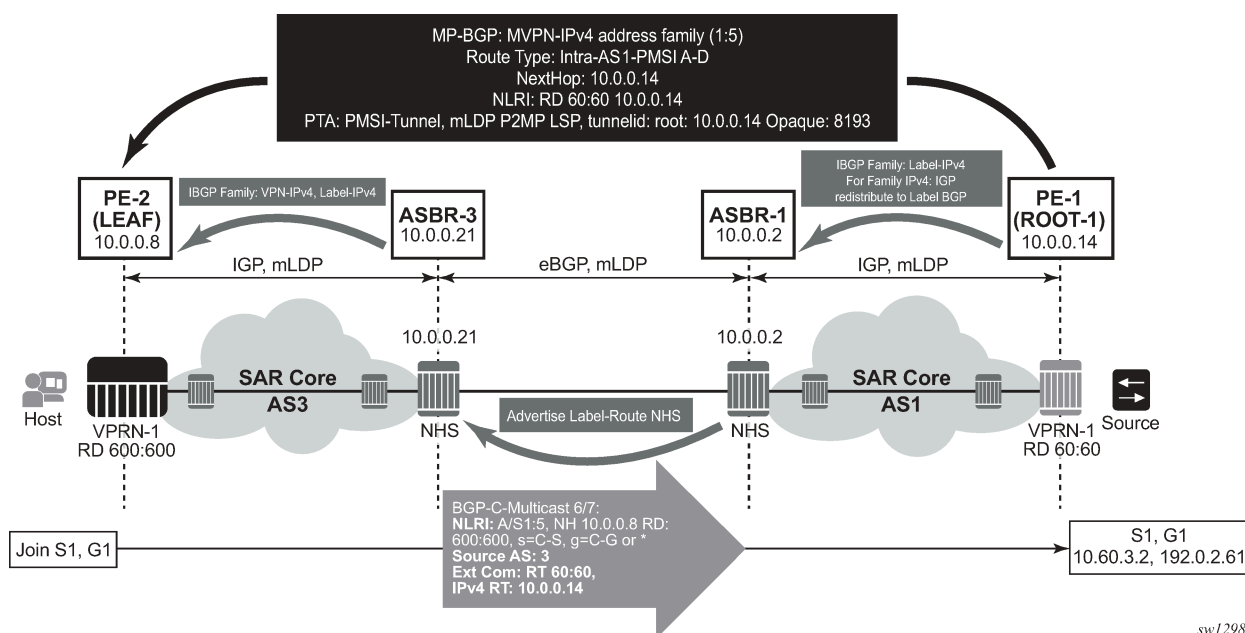
When the mLDP FEC arrives at ASBR1, it notes that it is the root node and that the opaque value is a recursive opaque value. As PE-1's address is known to ASBR1 via IGP, no further recursion is required. Regular processing begins, using the received opaque mLDP FEC information.

The functionality as described above for I-PMSI applies to S-PMSI and (C-*, C-*) S-PMSI.

7.2.7.5.1.1.1 C-multicast route processing

C-multicast route processing functionality follows RFC 6513 section 8.1.2 (BGP used for route exchange). The processing is similar to BGP unicast VPN route exchange. The following figure shows C-multicast route processing with non-segmented mLDP PMSI details.

Figure 121: Non-segmented mLDP C-multicast exchange (option C)



sw1298

7.2.7.5.1.1.2 LEAF node cavities



Caution: ASBRs do not currently support receiving a non-recursive opaque FEC (opaque type 1).

The LEAF (PE-2) must have the ROOT-1 system IP address installed in the RTM via BGP. If ROOT-1 is installed in the RTM via IGP, the LEAF will not generate the recursive opaque FEC and ASBR 3 will therefore not process the LDP FEC correctly.

7.2.7.5.2 Configuration example

No configuration is required for option C on ASBRs.

Policy is required for a root or leaf PE for removing the NO_EXPORT community from MVPN routes, which can be configured using an export policy on the PE.

The following is an example of configuring a policy on PEs to remove the NO_EXPORT community:

```
*A:Dut-A>config>router>policy-options# info
-----
community "no-export" members "no-export"
policy-statement "remNoExport"
  default-action accept
  community remove "no-export"
exit
exit
-----
*A:Dut-A>config>router>policy-options#
```

The following is an example of configuring the policy under BGP in a global, group, or peer context:

```
*A:Dut-A>config>router>bgp# info
-----
vpn-apply-export
export "remNoExport"
```

7.2.7.5.3 Inter-AS non-segmented mLDP

See the 7705 SAR MPLS Guide, "Inter-AS non-segmented mLDP" for information.

7.2.7.5.4 ECMP

See the 7705 SAR MPLS Guide, "ECMP support" under "Inter-AS non-segmented mLDP" for information about ECMP.

7.2.7.6 Mrinfo and Mtrace

When using **mrinfo** and **mtrace** in a Layer 3 VPN context, the configuration for the VPRN should have a loopback address configured that has the same address as the core VPRN instance's system address (that is, the BGP next hop).

For more information, see the "IP multicast debugging tools" section in the 7705 SAR OAM and Diagnostics Guide.

7.2.7.7 Multicast-only fast reroute

The 7705 SAR supports multicast-only fast reroute (MoFRR) in the context of GRT for mLDP. The multicast traffic is duplicated on a primary mLDP multicast tree and a secondary mLDP multicast tree.

For more information, see the "Multicast-only fast reroute (MoFRR)" section in the 7705 SAR Routing Protocols Guide.

7.2.7.8 mLDP point-to-multipoint support

The 7705 SAR supports mLDP point-to-multipoint traffic.

For more information, see the "LDP point-to-multipoint support" section in the 7705 SAR MPLS Guide.

7.2.7.9 mLDP fast upstream switchover

This feature allows a downstream LSR of an mLDP FEC to perform a fast switchover in order to source the traffic from another upstream LSR while IGP and LDP are converging due to a failure of the upstream LSR, where the upstream LSR is the primary next hop of the root LSR for the point-to-multipoint FEC.

For more information, see the "Multicast LDP fast upstream switchover" section in the 7705 SAR MPLS Guide.

7.2.7.10 Multicast Source Discovery Protocol

7705 SAR supports Multicast Source Discovery Protocol (MSDP) for MVPNs.

MSDP is a mechanism that allows rendezvous points (RPs) to share information about active sources. When RPs in remote domains hear about the active sources, they can pass on that information to the local receivers and multicast data can be forwarded between the domains. MSDP allows each domain to maintain an independent RP that does not rely on other domains, but it also enables RPs to forward traffic between domains. PIM-SM is used to forward the traffic between the multicast domains.

In addition to supporting MSDP on MVPNs in the VPRN service context, the 7705 SAR supports MSDP in the base router context. For information about MSDP, see the 7705 SAR Routing Protocols Guide.

In an MVPN, a PE node can act as an RP and run the MSDP functionality.

To interconnect multicast domains and to learn about source in other domains, MSDP peering is maintained between RP nodes. MSDP peering occurs over a TCP connection and control information is exchanged between peers to learn about multicast sources in other domains and to distribute information about multicast sources in the local domain.

When MSDP is configured in a service provider MVPN for an IP VPN customer, at least one of the PEs that are part of that MVPN becomes an MSDP peer to customer-instance RPs. MSDP groups are configured on PEs to limit source-active (SA) advertisements to routers within a group. As the PE RP learns about multicast sources within its domain via PIM-SM, it encapsulates the first data packet in an MSDP SA message and distributes it to all of its peer RP nodes. Based on the RPF check, each peer node sends the control message to other peers to distribute information about the active source. If there is an existing entry for the multicast group, the RP node joins the shortest path tree toward the source.

7.2.8 VPRN autobinding tunnels

The 7705 SAR supports autobinding for selecting tunnels in the tunnel table manager (TTM) in the following resolution contexts:

- resolution of RFC 3107 BGP label route prefix using tunnels to a BGP next hop
- resolution of a VPN-IPv4 or VPN-IPv6 prefix to a BGP next hop

The command to autobind tunnels is **config>service>vprn>auto-bind-tunnel**, which has **resolution** and **resolution-filter** options.

The user configures the **resolution** option to enable autobind resolution to tunnels in the TTM. If the **resolution** option is explicitly set to **disabled**, the autobinding to the tunnel is removed.

If **resolution** is set to **any**, any supported tunnel type in the resolution context will be selected following the TTM preference. The following tunnel types are selected in order of preference: RSVP, LDP, segment

routing, and GRE. The user can configure the preference of the segment routing tunnel type in the TTM for a specific IGP instance.

If **resolution** is set to **filter**, one or more explicit tunnel types are specified using the **resolution-filter** option, and only these specified tunnel types will be selected according to the TTM preference.



Note:

- If a VPRN is configured with **auto-bind-tunnel** using GRE and the BGP next hop of a VPN route matches a static blackhole route, all traffic matching that VPN route will be blackholed even if the static blackhole route is later removed. Similarly, if a static blackhole route is added after **auto-bind-tunnel** GRE has been enabled, the blackholing of traffic will not be performed optimally. In general, static blackhole routes that match VPN route next hops should be configured first, before the **auto-bind-tunnel** GRE command is applied.
- An SDP specified by **vprn>spoke-sdp** is always preferred over an autobind tunnel, regardless of the tunnel table manager (TTM) preference.

7.2.9 Spoke SDPs

For VPRN service, spoke SDPs can be used only for providing network connectivity between the PE routers.


7.2.10 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.

The following figure 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 122: SDP ID and VC label service identifiers (conceptual view of the service)



21510

The following figure 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 123: VPRN spoke-SDP termination



22726

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 about PPP/MLPPP ports, see the 7705 SAR Interface Configuration Guide, "Access, network, and hybrid ports".
- 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 (see the following figure): 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 124: Pseudowire-based backhaul (spoke-SDP termination at 7750 SR)



21512

- Scenario 2 (see the following figure): an MP-BGP-based solution can provide a fully routed scenario


Figure 125: VPRN in mobile backhaul application



21513

- Scenario 3 (see the following figure): 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 126: Spoke-SDP termination to VPRN



21514

7.2.11 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 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 or on the 12-port Serial Data Interface card, version 3 (v.35 ports). 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 128: IPv4 and IPv6 GRT-supported management protocols](#). The access control lists for 6VPE are shown in the following three tables.

Table 130: 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 131: 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 132: 6VPE access control list, r-VPLS override

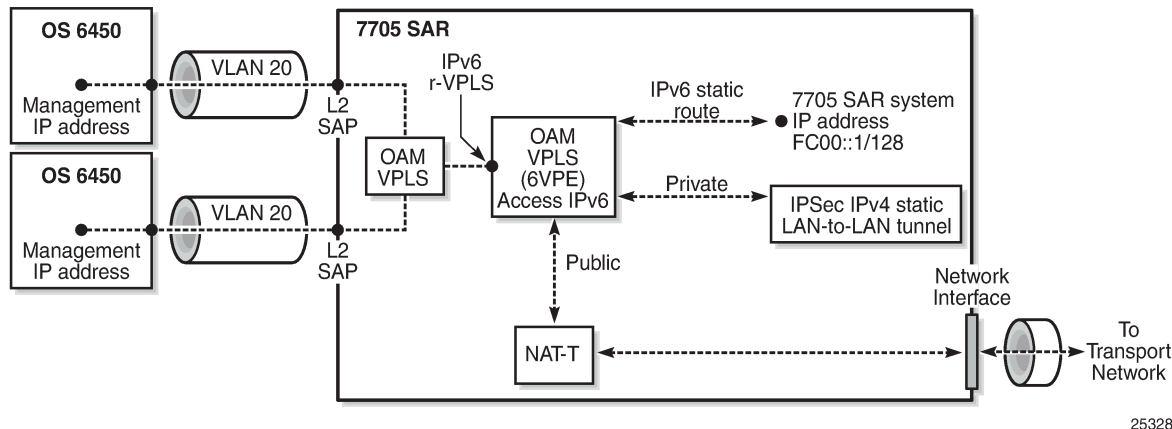
Service	Ingress override-v4	Ingress override-v6
IES	Yes	Yes
VPRN	Yes	Yes

7.2.12 IPv6 over IPv4 LAN-to-LAN IPsec tunnels

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 the following figure. 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 127: Access IPv6 traffic aggregation and encryption



7.2.13 Bandwidth optimization for low-speed links

The 7705 SAR can be used in deployments where the uplink bandwidth capacity and requirements are considerably less than if the router is used for fixed or mobile backhaul applications. For example, the 7705 SAR can be used to direct traffic from multiple individual homes for applications such as smart meter aggregation or relay connectivity. Connecting to end systems such as smart meters or relays requires uplink bandwidth capacity in terms of hundreds of kilobits per second instead of hundreds of megabits per second.

The 7705 SAR is optimized to operate in environments with megabits per second of uplink capacity for network operations. Therefore, many of the software timers are designed to ensure the fastest possible detection of failures, without considering bandwidth limitations. In deployments with very low bandwidth constraints, the system must also be optimized for effective operation of the routers without any interruption to mission-critical customer traffic. This can be achieved by:

- minimizing head-of-line (HoL) blocking by supporting a lower MTU
- redirecting self-generated traffic (SGT) to data queues (see the 7705 SAR Quality of Service Guide, "SGT redirection", for information)

One way to optimize operation in lower-bandwidth applications is to minimize HoL blocking caused by large packets. HoL blocking occurs when transmission of a large non-mission-critical packet delays a mission-critical packet beyond acceptable limits. The propagation delay of large packets over a slow link is fairly significant. For example, the propagation delay when transmitting a 1500-byte packet over a 100 kb/s link is 120 ms. If a mission-critical packet is queued immediately after the first bit of a non-mission-critical 1500-byte packet begins transmission, the mission-critical packet must wait 120 ms before the uplink is available again.

To minimize HoL blocking, the 7705 SAR now supports a lower MTU of 128 bytes (from the original 512-byte minimum) so that large IP packets can be fragmented into 128-byte chunks. In the preceding example, transmitting a 128-byte packet over a 100 kb/s link only delays the next packet by 10.24 ms.

This lower MTU is supported on IES and VPRN interfaces (access interfaces) and on network interfaces. The IP MTU is derived from the port MTU, unless specifically configured with the **ip-mtu** command. This command is supported on access interfaces only.

The following must be considered when using a lower IP MTU:

- applicability – the lower IP MTU is only applicable for IP forwarded traffic and cannot be applied to pseudowire or VPLS traffic
- reassembly – the far-end/destination node must reassemble the packet before it can process the data, which may impact the performance of the end system or may require different hardware to perform the reassembly
- extra overhead – each fragment must have an IPv4 header so that all fragments of the packet can be forwarded to the destination. Care must be taken to ensure that the extra IP overhead for each fragment does not offset the gain achieved by using the lower MTU. As an example, for a 128-byte packet, the IPv4 header, which is 20 bytes in length, constitutes approximately 15% of the total packet size.

**Note:**

- Lower IP MTU applies to IPv4 applications only. As per RFC 2640, IPv6 interfaces or dual-stack interfaces should not be configured to a value lower than 1280 bytes.
- Lower IP MTU is supported only on Ethernet encapsulated ports.
- Most routing and signaling protocols, such as OSPF, IS-IS, and RSVP-TE, cannot be supported with port MTUs lower than 512 bytes due to the protocol layer requirements and restrictions.
- Special care must be taken with routing protocols that use TCP, such as BGP and LDP. The minimum TCP MSS value supported on the 7705 SAR is 384 bytes; therefore, these protocols should only be enabled on links that can transport 384-byte IP packets without fragmentation. If there is a mismatch in TCP MSS in the network, this mismatch can potentially cause severe network performance issues due to the overhead caused by fragmentation and retransmissions, it can cause multi-vendor interoperability issues, and it can potentially cause the protocols to continuously flap.

- Not all OAM diagnostics are supported with lower port MTUs. Detailed information is provided in [OAM diagnostics restrictions with lower IP MTU](#).

7.2.13.1 OAM diagnostics restrictions with lower IP MTU

OAM tests require a minimum network port MTU to run; this value depends on the test. If the port MTU is set to a value lower than the minimum requirement, the test will fail.

If the port MTU is set to a value that meets the minimum requirement, the packet size parameter can be configured for the test (for example, **oam sdp-ping 1 size 102**).

If the **size** parameter is not specified, the system builds the packet based on the default payload size. If the **size** parameter is configured and is greater than the default payload size, padding bytes are added to equal the configured value.

The packet size is dependent on the port MTU value; that is, if the minimum port MTU value is used, there are restrictions on the packet size. If the configured size is greater than the maximum value supported with the minimum port MTU, the test will fail.

The following tables list the minimum port MTU required for each OAM test and the maximum size of the OAM packet that can be configured when the minimum port MTU is used, based on SDP tunnel type.



Note: RSVP LSPs will not come up if the network port MTU value is lower than 302 bytes.

Table 133: Port MTU requirements for OAM diagnostics (GRE tunnels)

SDP type: GRE		
Test type	Minimum network port MTU requirement over Ethernet dot1q encapsulation (bytes)	OAM test size range (bytes)
sdp-ping	128	72 to 82
svc-ping	196	N/A ¹
vccv-ping	143	1 to 93
vccv-trace	143	1 to 93
vprn-ping	182	1 to 136
vprn-trace	302	1 to 256
mac-ping	188	1 to 142
mac-trace	240	1 to 194
cpe-ping	186	N/A ¹

Note:

1. Size is not configurable

Table 134: Port MTU requirements for OAM diagnostics (LDP tunnels)

SDP type: LDP		
Test type	Minimum network port MTU requirement over Ethernet dot1q encapsulation (bytes)	OAM test size range (bytes)
lsp-ping	128	1 to 106
lsp-trace	128	1 to 104
sdp-ping	128	72 to 102
svc-ping	176	N/A ¹
vccv-ping	128	1 to 98
vccv-trace	128	1 to 98
vprn-ping	182	1 to 156
vprn-trace	302	1 to 276
mac-ping	168	1 to 142
mac-trace	220	1 to 194
cpe-ping	166	N/A ¹

Note:

1. Size is not configurable

For information about OAM diagnostics, see the 7705 SAR OAM and Diagnostics Guide.

7.2.14 Support for NTP

On the 7705 SAR, communication with external NTP clocks over VPRNs is supported for external NTP servers and peers and for external NTP clients.

Communication with external servers and peers is controlled using the same commands as those used in the base routing context; see the 7705 SAR Basic System Configuration Guide, "System time commands", for information. Communication with external clients is controlled using commands in the VPRN context. Support for external clients can be as a unicast or a broadcast service. In addition, authentication keys for external clients are configurable on a per-VPRN basis.

7.3 Configuring a VPRN service with CLI

This section provides information to configure virtual private routed network (VPRN) services using the CLI.

Topics in this section include:

- [Basic configuration](#)
- [Common configuration tasks](#)
- [Configuring VPRN components](#)
- [Service management tasks](#)

7.4 Basic configuration

The following fields require specific input (there are no defaults) to configure a basic VPRN service:

- customer ID (see [Configuring customer accounts](#))
- interface parameters
- spoke SDP parameters (at VPRN service level)

The following example displays a VPRN service configuration.

```
*A:ALU-1>config>service>vprn# info
-----
vrf-import "vrfImpPolCust1"
vrf-export "vrfExpPolCust1"
autonomous-system 10000
route-distinguisher 10001:1
auto-bind-tunnel
    resolution-filter
        ldp
    exit
    resolution filter
exit
vrf-target target:10001:1
interface "to-cel" create
    address 172.16.0.0/12
    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-entry 10.1.1.1/8
    next-hop 10.1.1.2
    no shutdown
    exit
exit
-----
*A:ALU-1>config>service>vprn#
```

7.5 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. Enabling BGP in the global routing context is needed in some configurations; for example, MVPN in PIM-SM.

To configure BGP for the router, see the 7705 SAR 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-tunnel** parameter. Alternatively, define a VPRN spoke SDP. When using the **vprn>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 PIM parameters (optional).
7. Define OSPF parameters (optional).
8. Create a VPRN interface.
9. Define SAP parameters on the VPRN interface:
 - select nodes and ports
 - optional – select QoS policies other than the default (configured in the **config>qos** context)
 - optional – select filter policies (configured in the **config>filter** context)
 - optional – select accounting policy (configured in the **config>log** context)
 - optional – configure DHCP features
10. Define BGP parameters for the VPRN (**config>service>vprn>bgp**) (optional).
11. Define RIP parameters (optional).
12. Define spoke SDP parameters on the VPRN interface.
13. Configure VRRP (optional).
14. Enable the service.

7.6 Configuring VPRN components

Topics in this section include:

- [Creating a VPRN service](#)
- [Configuring global VPRN parameters](#)
- [Configuring router interfaces](#)
- [Configuring static route entries for VPRN](#)
- [Configuring BGP for VPRN](#)
- [Configuring IPv6 parameters for VPRN BGP](#)
- [Configuring VPRN IPv6 neighbor discovery parameters](#)

- [Configuring OSPF or OSPFv3 for VPRN](#)
- [Configuring RIP for VPRN](#)
- [Configuring IGMP for VPRN](#)
- [Configuring PIM for VPRN](#)
- [Configuring MVPN 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 serial raw socket transport within a VPRN](#)
- [Configuring VPRN router advertisement](#)

7.6.1 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
```

The following example displays a VPRN service configuration.

```
*A:ALU-1>config>service# info
-----
...
    vprn 1 customer 1 create
        route-distinguisher 10001:0
        no shutdown
    exit
...
-----
*A:ALU-1>config>service>vprn#
```

7.6.2 Configuring global VPRN parameters

The autonomous system (AS) number for a VPRN service is configured in the **config>service>vprn** context. Local AS numbers can be set at the VPRN BGP global, group, and neighbor levels.

A spoke SDP can be bound to the VPRN service using the **auto-bind-tunnel** 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 10001:0
    spoke-sdp 2 create
    exit
    no shutdown
  exit
...
-----
*A:ALU-1>config>service#
```

7.6.3 Configuring router interfaces

See the 7705 SAR 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 10.0.0.0/8
    port 1/1/33
  exit
  interface "if2"
    address 10.0.0.1/8
    port 1/1/34
  exit
  interface "if3"
    address 10.0.0.2/8
    port 1/1/35
  exit
...
#-----
ALU48>config>router#
```

7.6.4 Configuring static route entries for VPRN

The 7705 SAR VPRN service supports static routes to next-hop addresses.

Only one next-hop IP address can be specified per IP interface for static routes.

Use the following CLI syntax to create a VPRN static route entry. Multiple types of static routes (**black-hole**, **grt**, **indirect**, **ipsec-tunnel**, and **next-hop**) can be applied to the same entry. Unless **no shutdown** is specified, the **static-route-entry** will be created in a **shutdown** state.

CLI syntax:

```
config>service>vprn>
static-route-entry {ip-prefix/prefix-length}
    black-hole {ip-int-name | ip-address | ipv6-address}
        description description-string
        metric metric
        preference preference
        prefix-list prefix-list-name [all | none]
        no shutdown
        tag tag
    grt
        description description-string
        metric metric
        preference preference
        no shutdown
    indirect ip-address
        cpe-check cpe-ip-address
            drop-count count
            interval seconds
            log
        description description-string
        metric metric
        preference preference
        prefix-list prefix-list-name {all | none}
        no shutdown
        tag tag
    ipsec-tunnel [ipsec-tunnel-name]
        description description-string
        metric metric
        preference preference
        no shutdown
        tag tag
    next-hop {ip-int-name | ip-address | ipv6-address}
        bfd-enable
        cpe-check cpe-ip-address
            drop-count count
            interval interval
            log
        description description-string
        metric metric
        preference preference
        prefix-list prefix-list-name [all | none]
        no shutdown
        tag tag
```

Example:

```
config>service>vprn# static-route-entry 10.5.5.5/8
static-route-entry# next-hop 10.1.1.2
next-hop# metric 1
next-hop# preference 5
next-hop# tag 20
next-hop# no shutdown
```

7.6.5 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 an example of a VPRN service with a configured autonomous system number, see [Configuring global VPRN parameters](#).

- a router ID

For an example of a VPRN service with a configured router ID, see [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

VPRN BGP is administratively enabled upon creation. Minimally, to enable VPRN BGP in a VPRN instance, you must associate an autonomous system number and router ID for the VPRN service, create a peer group, neighbor, and associate a peer AS number. There are no default VPRN BGP groups or neighbors. Each VPRN BGP group and neighbor must be explicitly configured.

All parameters configured for VPRN BGP are applied to the group and are inherited by each peer, but a group parameter can be overridden on a specific basis. 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, see the 7705 SAR Routing Protocols Guide, "BGP".

7.6.5.1 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.

7.6.5.2 Configuring route reflection

Route reflection can be implemented in autonomous systems with a large internal BGP mesh to reduce the number of IBGP sessions required. One or more routers can be selected to act as focal points for internal BGP sessions. Several BGP-speaking routers can peer with a route reflector. A route reflector forms peer connections to other route reflectors. A router assumes the role as a route reflector by configuring the **cluster** *cluster-id* command. No other command is required unless disabling reflection to specific peers is needed.

If you configure the **cluster** command at the global level, all subordinate groups and neighbors are members of the cluster. The route reflector cluster ID is expressed in dotted-decimal notation. The ID should be a significant topology-specific value. No other command is required unless disabling reflection to specific peers is needed.

If a route reflector client is fully meshed, the **disable-client-reflect** command can be enabled to stop the route reflector from reflecting redundant route updates to a client.

7.6.5.3 VPRN BGP CLI syntax

The following example displays a VPRN BGP configuration. The example includes two BGP groups: one group has a static (configured) neighbor and the other group has dynamic neighbors.

```
*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-tunnel
      resolution-filter
        ldp
      exit
      resolution filter
    exit
    vrf-target target:10001:1
    interface "to-cel" create
      address 172.16.0.0/12
      sap 1/1/10:1 create
        ingress
          qos 100
          filter ip 6
        exit
        egress
          qos 1010
        exit
      exit
    exit
  static-route-entry 10.1.1.1/8
    next-hop 10.1.1.2
    no shutdown
  exit
```

```

exit
bgp
  router-id 10.0.0.1
  group "to-cel"
    export "vprnBgpExpPolCust1"
    multihop 3
    peer-as 65101
    ttl-security 10
    neighbor 172.16.0.10
  exit
  group "dynamic"
    peer-as 100
    dynamic-neighbor
      prefix 10.100.0.0/16
    dynamic-neighbor-limit 75
  exit
exit
exit
spoke-sdp 2 create
exit
no shutdown
exit
...
-----
*A:ALU-1>config>service#

```

7.6.6 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:a::123
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

```

7.6.7 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
...
```

7.6.8 Configuring OSPF or OSPFv3 for VPRN

Each VPN routing instance is isolated from any other VPN routing instance and from the routing used across the backbone. OSPF or OSPFv3 can be run with any VPRN, independently of the routing protocols used in other VPRNs, or in the backbone. For more information about the OSPF and OSPFv3 protocols and for the commands used to run OSPF or OSPFv3 over the backbone (router context), see the 7705 SAR Routing Protocols Guide.

Use the following CLI syntax to configure OSPF or OSPFv3 in the VPRN context:

CLI syntax:

```
config>service>vprn>ospf#
```

CLI syntax:

```
config>service>vprn>ospf3#
```

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#

```

7.6.9 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

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-tunnel
    resolution-filter
        ldp
    exit
    resolution filter
exit
vrf-target target:10001:1
interface "to-cel" create
    address 172.16.0.0/12
    sap 1/1/10:1 create
        ingress
            qos 100
        exit
        egress
            qos 1010
            filter ip 6
        exit
    exit
exit
rip
    export "vprnRipExpPolCust1"
    group "cel"
        neighbor "to-cel"
    exit
exit
spoke-sdp 2 create
exit
no shutdown
exit
...
-----

```

For more information about the RIP protocol, see the 7705 SAR Routing Protocols Guide.

7.6.10 Configuring IGMP for VPRN

When using the **ssm-translate** command, the group range is not created until the source is specified.

The following example displays multicast IGMP parameters under a VPRN configuration:

```

*A:Sar18 Dut-B>config>service>vprn>igmp# info detail
-----
    interface "mvpn_if"
        no import
        version 3
        subnet-check
        no max-groups
        no max-grp-sources
        no disable-router-alert-check
        ssm-translate
            grp-range 239.255.0.2 239.255.0.20
            source 192.168.0.0
        exit
    exit
    no shutdown
exit

```

```

query-interval 125
query-last-member-interval 1
query-response-interval 10
robust-count 2
no shutdown
-----
*A: Sar18 Dut-B>config>service>vprn>igmp#

```

7.6.11 Configuring PIM for VPRN

The following example displays a PIM configuration for VPRN.

```

A: ALU-1>config>service>vprn>pim# info detail
-----
no import join-policy
no import register-policy
interface "vprn_if"
  hello-interval 30
  hello-multiplier 35
  no tracking-support
  improved-assert
  no bfd-enable
  no three-way-hello
  priority 1
  multicast-senders auto
  no bsm-check-rtr-alert
  no sticky-dr
  no max-groups
  no assert-period
  no instant-prune-echo
  no shutdown
  no ipv4-multicast-disable
exit
apply-to none
rp
  no bootstrap-import
  no bootstrap-export
  static
  exit
  bsr-candidate
  shutdown
  priority 0
  hash-mask-len 30
  no address
  exit
  rp-candidate
  shutdown
  no address
  holdtime 150
  priority 192
  exit
exit
no non-dr-attract-traffic
no ssm-default-range-disable ipv4
no shutdown
no ipv4-multicast-disable
-----
A: ALU-1>config>service>vprn>pim#

```

7.6.12 Configuring MVPN for VPRN

For selective PMSI provider tunnels, mLDP must be configured before setting a **maximum-p2mp-spmsi**. Also, the **data-threshold** *c-grp-ip-addr* must be a valid multicast address.

The following example displays the MVPN parameters for VPRN configuration:

```
*A:ALU>config>service>vprn>mvpn# info detail
-----
      auto-discovery default
      c-mcast-signaling bgp
      umh-selection highest-ip
      mdt-type sender-receiver
      provider-tunnel
        inclusive
        mldp
        shutdown
      exit
    exit
  selective
    mldp
    shutdown
  exit
    maximum-p2mp-spmsi 4
    no data-delay-interval
    data-threshold 239.255.0.0/6 10
  exit
exit
vrf-target unicast
exit
-----
*A:ALU>config>service>vprn>mvpn#
```

The following example displays a VPRN service with MVPN. The MVPN in this example supports inclusive PMSI and selective PMSI. The **data-threshold** that forces a group C(S,G) to switch from I-PMSI to S-PMSI in this example is 1 kb/s.

```
vprn 1 customer 1 create
  route-distinguisher 10001:1
  auto-bind-tunnel
    resolution-filter
      ldp
      rsvp
    exit
  resolution filter
exit
vrf-target target:65000:1
interface "T0-CE-SOURCE" create
  address 172.16.0.1/12
  sap 1/1/9:100 create
  exit
exit
pim
  interface "to-ce-source"
  rp
  exit
  no shutdown
exit
mvpn
  provider-tunnel
    inclusive
```

```

        mldp
        no shutdown
    exit
    selective
        mldp
        no shutdown
    exit
    data-threshold 239.255.0.0/7 1
    exit
    vrf-target target:65000:1
    exit
    ospf
    area 0.0.0.0
    interface "T0-CE-SOURCE"
    interface-type point-to-point
    no shutdown
    exit
    exit
    exit
    no shutdown
    exit

```

7.6.13 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.

When using **mrinfo** and **mtrace** in a Layer 3 VPN context, the configuration for the VPRN should have a loopback address configured that has the same address as the core VPRN instance's system address (that is, the BGP next hop).

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-tunnel
resolution-filter
    ldp
exit
resolution filter

```



```

exit
vrf-target target:10001:1
interface "to-cel" create
    address 172.16.0.1/12
    proxy-arp policy "proxyARPolicy"
    local proxy-arp
    remote proxy-arp
exit
exit
static-route-entry 10.1.1.1/8
    next-hop 10.1.1.2
    no shutdown
exit
exit
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]
            unreachable
        if-attribute
            admin-group group-name [group-name...(up to 5 max)]
            srlg-group group-name [group-name...(up to 5 max)]
        ip-mtu octets
        ipcp
            dns ip-address [secondary ip-address]
            dns secondary ip-address
            peer-ip-address ip-address
        l4-load-balancing hashing-algorithm
        local-dhcp-service local-server-name
        local-proxy-arp

```

```

        loopback
        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
        tcp-mss mss-value
        teid-load-balancing
        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 192.168.0.0/16
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 192.168.0.0/16 broadcast host-ones
        no mac
        arp-timeout 14400
        no allow-directed-broadcasts
        icmp
            mask-reply
            unreachable 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 "proxyARPolicy"
        local proxy-arp
        remote proxy-arp
        no shutdown...

```

7.6.14 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 ipv6-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 iaaid]
          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
        unreachable 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# 2001:db8:a::123
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# unreachable 100 10
config>service>vprn>if>ipv6>icmp6# exit
config>service>vprn>if>ipv6>neighbor# 2001:db8:a::124
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
```

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
            unreachable 100 10
        exit
        address 2001:db8:a::123
        reachable-time 30
        stale-time 14400
        no dhcp6-relay
        no local-dhcp-server
        neighbor 2001:db8:a::124
        no bfd
        exit
.....
```

7.6.15 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
```

7.6.16 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 **loopback** command is enabled on the interface.

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:

- T1/E1 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 the 16-port T1/E1 ASAP Adapter card or 32-port T1/E1 ASAP Adapter card:
 - fractional T1/E1
 - clear channel T1/E1
 - any T1/E1 ASAP port or bundle on the 7705 SAR-X, 7705 SAR-M, or 7705 SAR-A:
 - fractional T1/E1
 - clear channel T1/E1
- V.35 ports in access mode with PPP encapsulation on the 12-port Serial Data Interface card, version 3, with speed set to 64 kb/s, 2048 kb/s, or any value from 128 kb/s to 1920 kb/s (every 128 kb/s)
- DS1/E1 channels on the 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card:
 - the SAP can be a PPP link over a single DS1/E1 channel
 - the SAP can be an MLPPP or MC-MLPPP bundle over multiple DS1/E1 channels
- Ethernet port in access mode:
 - any Ethernet port (null, dot1q, or qinq) on the 6-port Ethernet 10Gbps Adapter card, 8-port Gigabit Ethernet Adapter card, or 10-port 1GigE/1-port 10GigE X-Adapter card (supported on the 7705 SAR-18 only)
 - any Ethernet port (null, dot1q, or qinq) on the 7705 SAR-M, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-A, 7705 SAR-Ax, 7705 SAR-Wx, or 7705 SAR-X



Note:

- IPv6 and multicast are not supported on PPP, MLPPP, or MC-MLPPP SAPs on the 4-port OC3/STM1 / 1-port OC12/STM4 Adapter card.
- The 10-port 1GigE/1-port 10GigE X-Adapter card supports qinq only when it is in 10-port 1GigE mode.

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 10001:10
    auto-bind-tunnel
    resolution-filter
    ldp
```

```

        exit
        resolution filter
    exit
    vrf-target target:10001:1
    interface "to-cel" create
        address 172.16.0.0/12
        sap 1/1/10:1 create
            ingress
                qos 100
                filter ip 6
            exit
            egress
                qos 1010
            exit
        exit
    exit
    static-route-entry 192.168.0.0/16
        next-hop 192.168.0.1
        no shutdown
    exit
    spoke-sdp 2 create
    exit
    no shutdown
exit
...
-----
*A:ALU-1>config>service#

*A:ALU-1>config>service>vprn# info
-----
description "test VPRN for PPP SAPs"
route-distinguisher 10001:1
vrf-target target:10001:1
interface "to-cel" create
    address 172.16.0.0/12
    sap 1/1/10:1 create
    exit
exit
interface "to_ce2_ppp" create
    address 172.16.0.1/12
    bfd 100 receive 100 multiplier 3
    ipcp
        peer-ip-address 192.168.0.50
    exit
    sap 1/1/2.24 create
    exit
exit
interface "to_ce2_mlppp" create
    address 172.16.0.3/12
    bfd 100 receive 100 multiplier 3
    ipcp
        peer-ip-address 192.168.0.51
        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 172.16.0.3/12
    sap 1/2/1:25 create
    exit
exit

```

```

static-route-entry 192.168.0.0/16
  next-hop 192.168.0.5
  no shutdown
exit
exit
static-route-entry 192.168.0.1/16
  next-hop 192.168.0.6
  no shutdown
exit
exit
static-route-entry 192.168.0.2/16
  next-hop 192.168.0.7
  no shutdown
exit
exit
static-route-entry 192.168.0.3/16
  next-hop 192.168.0.8
  no shutdown
exit
exit
static-route-entry 192.168.0.3/16
  next-hop 192.168.0.9 disable
  shutdown
exit
exit
-----
*A:ALU-1>config>service>vprn#

```

7.6.17 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

```

7.6.18 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

```

7.6.19 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

```


7.6.20 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 addresses 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 overview information about VRRP and VRRP VPRN interface parameters, see the "VRRP" chapter in the 7705 SAR Router Configuration Guide.

The following displays a VPRN interface VRRP owner configuration:

```
config>service>vprn> info
#-----
...
  interface "vrrpowner"
    address 10.10.10.24
    vrrp 1 owner
      backup 10.10.10.23
      authentication-key "testabc"
    exit
  exit
...
#-----
config>service>vprn#
```

```
config>service>vprn>if># info
-----
...
  ipv6
    address 2001:db8:a::123
    vrrp 1 owner
      backup 2001:db8:a::124
    exit
  exit
  exit
...
-----
```

7.6.21 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 following CLI syntax

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-100-192.168.0.0
    exit
  nat
    pool 1 create
      description "pool 1"
      direction zone-inbound
      exit
    entry 1 create
      ip-addr interface vprn-100-203.0.113.0
      exit
    exit
  exit
  policy 1 nat pool 1
  commit
  exit
  no-shutdown
-----
A:ALU-B>config>service>ies#

```

7.6.22 Configuring serial raw socket transport within a VPRN

Configure an IP transport subservice within a VPRN service to enable the transport of serial data using raw sockets.

CLI syntax:

```

config>service

```

```

    vprn service-id [customer customer-id] [create]
    ip-transport ipt-id [create]
        description description-string
        filter-unknown-host
        local-host ip-addr ip-addr port-num port-num protocol {tcp |
udp}
        remote-host host-id [ip-addr ip-addr] [port-num port-num]
[create]
        description description-string
        name host-name
        exit
        fc fc-name profile {in | out}
        shutdown
        tcp
            inactivity-timeout seconds
            max-retries number
            retry-interval seconds
        exit
    exit
exit
exit
exit

```

The following example displays an IP transport subservice configuration output.

```

A:ALU-B>config>service>vprn# info
-----
    configure
    service vprn 100 create
        ip-transport 1/2/4.1 create
        description "ip-transport vprn"
        filter-unknown-host
        local-host ip-address 192.168.0.0 port-number 4000 protocol udp
        exit
        remote-host 1 ip-address 192.168.0.1 port-number 4001 create
        exit
    exit
    no-shutdown
-----
A:ALU-B>config>service>vprn

```

7.6.23 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 2001:db8:a::123
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 2001:db8:a::123
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#

```

7.7 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](#)

7.7.1 Modifying VPRN service parameters

Use the CLI syntax to modify VPRN parameters.

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-cel" create
      address 172.16.0.0/12
      sap 1/1/10:1 create
      exit
    exit
    static-route-entry 192.168.255.255/16
      next-hop 192.168.0.0
      no shutdown
    exit
    exit
    spoke-sdp 2 create
    exit
  exit
...
-----
*A:ALU-1>config>service>vprn#
```

7.7.2 Deleting a VPRN service

A VPRN service cannot be deleted until SAPs, interface spoke SDPs, and interfaces are shut down and deleted. If protocols 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
```

7.7.3 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-tunnel
      resolution-filter
        ldp
      exit
    resolution filter
    exit
    vrf-target target:10001:1
    interface "to-cel" create
      address 172.16.0.0/12
      sap 1/1/10:1 create
        ingress
          qos 100
          filter ip 6
        exit
        egress
          qos 1010
        exit
      exit
    exit
    static-route-entry 192.168.0.0/16
      next-hop 192.168.0.1
      no shutdown
    exit
    spoke-sdp 2 create
    exit
  exit
...
-----
*A:ALU-1>config>service#
```

7.7.4 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
```

7.8 VPRN services command reference

7.8.1 Command hierarchies

- Configuration commands
 - VPRN service configuration commands
 - BGP commands
 - OSPF commands
 - OSPFv3 commands
 - IGMP commands
 - PIM commands
 - RIP commands
 - VPRN security zone configuration commands
 - VPRN raw socket IP transport configuration commands
 - Multicast VPN commands
 - MSDP 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
 - VPRN static one-to-one NAT configuration commands
 - TWAMP Light commands
 - VPRN NTP commands
- Show commands
- Clear commands
- Debug commands

7.8.1.1 Configuration commands

7.8.1.1.1 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] [as-set] [aggregator as-
number:ip-address]
-   no aggregate ip-prefix/ip-prefix-length
-   auto-bind-tunnel
-   - ecmp max-ecmp-routes
-   - no ecmp
-   - resolution {any | filter | disabled}
-   - resolution-filter
-   - [no] gre
-   - [no] ldp
-   - [no] rsvp
-   - [no] sr-isis
-   - [no] sr-ospf
-   - [no] sr-te
-   - [no] weighted-ecmp
-   autonomous-system as-number
-   no autonomous-system
-   [no] bgp
-   description description-string
-   no description
-   dhcp (see Local DHCP and DHCPv6 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}
-   [no] entropy-label
-   grt-lookup
-   - [no] enable-grt
-   - [no] allow-local-management
-   - export-grt policy-name [policy-name...(up to 5 max)]
-   - no export-grt
-   - export-limit num-routes
-   - no export-limit
-   - export-v6-limit num-routes
-   - no export-v6-limit
-   interface ip-int-name
-   - [no] entropy-label
-   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} [fc-queue fc-name
profile {in | out}]
  - application dot1p-app-name dot1p {dot1p-priority} [fc-queue fc-name profile
{in | out}]
  - 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-entry ip-prefix/prefix-length
  - [no] black-hole
    - [no] description description-string
    - [no] metric metric
    - preference preference
    - no preference
    - [no] prefix-list prefix-list-name {all | none}
    - [no] shutdown
    - [no] tag tag
  - [no] grt
    - [no] description description-string
    - [no] metric metric
    - preference preference
    - no preference
    - [no] shutdown
  - [no] indirect ip-address
    - [no] cpe-check cpe-ip-address
      - drop-count count
      - no drop-count
      - interval seconds
      - no interval
      - [no] log
    - [no] description description-string
    - [no] metric metric
    - preference preference
    - no preference
    - [no] prefix-list prefix-list-name {all | none}
    - [no] shutdown
    - [no] tag tag
  - [no] ipsec-tunnel ipsec-tunnel-name
    - [no] description description-string
    - [no] metric metric
    - preference preference
    - no preference
    - [no] shutdown
    - [no] tag tag
  - [no] next-hop {ip-int-name | ip-address | ipv6-address}
    - [no] bfd-enable
    - [no] cpe-check cpe-ip-address
      - drop-count count
      - no drop-count
      - interval seconds
      - no interval
      - [no] log
    - [no] description description-string

```

```

- [no] metric metric
- preference preference
- no preference
- [no] prefix-list prefix-list-name {all | none}
- [no] shutdown
- [no] tag tag
- 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] weighted-ecmp
- zone zone-id [create]

```

7.8.1.1.2 BGP commands

```

config
- service
- vprn
- [no] bgp
- [no] advertise-inactive
- [no] aggregator-id-zero
- [no] as-override
- auth-keychain name
- no auth-keychain
- authentication-key {authentication-key | hash-key} [hash | hash2]
- no authentication-key
- [no] backup-path [ipv4] [ipv6]
- best-path-selection
- always-compare-med [zero | infinity]
- always-compare-med strict-as [zero | infinity]
- no always-compare-med
- as-path-ignore [ipv4] [ipv6]
- no as-path-ignore
- [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
- dynamic-neighbor-limit peers
- no dynamic-neighbor-limit
- [no] enable-bgp-vpn-backup [ipv4] [ipv6]
- [no] enable-peer-tracking
- error-handling
- [no] legacy-mode
- [no] update-fault-tolerance
- 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-route-advertisement seconds
- no min-route-advertisement
- multihop ttl-value
- no multihop
- multipath max-paths
- no multipath
- next-hop-resolution
  - policy policy-name
  - no policy
- peer-tracking-policy policy-name
- no peer-tracking-policy
- preference preference
- no preference
- [no] rapid-withdrawal
- [no] remove-private [limited]
- rib-management
  - ipv4
    - route-table-import policy-name
    - no route-table-import
  - ipv6
    - route-table-import policy-name
    - no route-table-import
- router-id ip-address
- no router-id
- [no] shutdown
- [no] split-horizon

config
- service
  - vprn
    - bgp
      - [no] group name
      - [no] advertise-inactive
      - [no] aggregator-id-zero
      - [no] as-override
      - auth-keychain name
      - no auth-keychain
      - 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
      - dynamic-neighbor
        - [no] prefix ip-prefix/ip-prefix-length
      - dynamic-neighbor-limit peers
      - no dynamic-neighbor-limit
      - [no] enable-peer-tracking
      - error-handling

```

```

- [no] update-fault-tolerance
- 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-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 family limit [threshold percentage] [idle-timeout {minutes |
forever} | log-only] [post-import]
- no prefix-limit family
- [no] remove-private [limited]
- [no] shutdown
- [no] split-horizon
- ttl-security min-ttl-value
- no ttl-security

config
- service
  - vprn
    - bgp
      - group
        - [no] neighbor ip-address
        - [no] advertise-inactive
        - [no] aggregator-id-zero
        - [no] as-override
        - auth-keychain name
        - no auth-keychain
        - 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

```

```

- error-handling
  - [no] update-fault-tolerance
- 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-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 family limit [threshold percentage] [idle-timeout
{minutes | forever} | log-only] [post-import]
- no prefix-limit family
- [no] remove-private [limited]
- [no] shutdown
- [no] split-horizon
- ttl-security min-ttl-value
- no ttl-security

```

7.8.1.1.3 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
        - interface ip-int-name [secondary]
        - no interface ip-int-name
          - [no] advertise-subnet
          - auth-keychain name
          - no auth-keychain
          - authentication-key {authentication-key | hash-key} [hash | hash2]
          - no authentication-key
          - authentication-type {password | message-digest}

```

```

- no authentication-type
- bfd-enable [remain-down-on-failure]
- no bfd-enable
- dead-interval seconds
- no dead-interval
- hello-interval seconds
- no hello-interval
- interface-type {broadcast | point-to-point}
- no interface-type
- lfa-policy-map route-nh-template template-name
- no lfa-policy-map
- load-balancing-weight weight
- no load-balancing-weight
- [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] [adjacency-check]
- no originate-default-route
- [no] redistribute-external
- [no] summaries
- [no] sham-link [ip-int-name ip-address]
- auth-keychain name
- no auth-keychain
- 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
- metric metric
- no metric
- retransmit-interval seconds
- no retransmit-interval
- [no] shutdown
- transit-delay seconds
- no transit-delay
- [no] stub
- default-metric metric
- no default-metric
- [no] summaries
- [no] virtual-link router-id transit-area area-id
- auth-keychain name

```

```

- no auth-keychain
- 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
- import policy-name [policy-name...(up to 5 max)]
- no import
- [no] loopfree-alternates
- exclude
- prefix-policy prefix-policy [prefix-policy...(up to 5 max)]
- no prefix-policy
- 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

```

7.8.1.1.4 OSPFv3 commands

```
config
```

```

- service
- vprn
- [no] ospf3
- [no] area area-id
- area-range ipv6-prefix/prefix-length [advertise | not-advertise]
- no area-range ipv6-prefix/prefix-length
- [no] blackhole-aggregate
- interface ip-int-name [secondary]
- no interface ip-int-name
- authentication bidirectional sa-name
- authentication inbound sa-name outbound sa-name
- no authentication
- bfd-enable [remain-down-on-failure]
- no bfd-enable
- dead-interval seconds
- no dead-interval
- hello-interval seconds
- no hello-interval
- interface-type {broadcast | point-to-point}
- no interface-type
- lfa-policy-map route-nh-template template-name
- no lfa-policy-map
- [no] loopfree-alternate-exclude
- 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
- key-rollover-interval key-rollover-interval
- no key-rollover-interval
- [no] loopfree-alternate-exclude
- [no] nssa
- area-range ipv6-prefix/prefix-length [advertise | not-advertise]
- no area-range ipv6-prefix/prefix-length
- originate-default-route [type-nssa] [adjacency-check]
- 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 bidirectional sa-name
- authentication inbound sa-name outbound sa-name
- no authentication
- dead-interval seconds
- no dead-interval
- hello-interval seconds
- no hello-interval
- 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-alternates
  - exclude
    - prefix-policy prefix-policy [prefix-policy...(up to 5 max)]
    - no prefix-policy
- 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] 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

```

7.8.1.1.5 IGMP commands

```

config
- service
  - vprn
    - [no] igmp
      - [no] interface ip-int-name
        - [no] disable-router-alert-check
        - import policy-name
        - no import
        - max-groups value
        - no max-groups
        - max-grp-sources max-grp-sources
        - no max-grp-sources
        - [no] shutdown
        - ssm-translate
          - [no] grp-range start end
            - [no] source ip-address
        - static
          - [no] group grp-ip-address
            - [no] source ip-address
        - [no] subnet-check
        - version version
        - no version
      - query-interval seconds
      - no query-interval
      - query-last-member-interval seconds
      - no query-last-member-interval
      - query-response-interval seconds

```

```

- no query-response-interval
- robust-count robust-count
- no robust-count
- [no] shutdown
- ssm-translate
  - [no] grp-range start end
  - [no] source ip-address

```

7.8.1.1.6 PIM commands

```

config
- service
  - vprn
    - [no] pim
      - apply-to {all | none}
      - import {join-policy | register-policy} policy-name [policy-name...(up to 5
max)]
      - no import {join-policy | register-policy}
      - [no] interface ip-int-name
        - assert-period assert-period
        - no assert-period
        - [no] bsm-check-rtr-alert
        - [no] bfd-enable [ipv4]
        - hello-interval hello-interval
        - no hello-interval
        - hello-multiplier deci-units
        - no hello-multiplier
        - [no] improved-assert
        - [no] instant-prune-echo
        - [no] ipv4-multicast-disable
        - max-groups value
        - no max-groups
        - multicast-senders {auto | always | never}
        - no multicast-senders
        - multicast-to-multicast source ip-address group-start ip-address group-
end ip-address to-multicast group-address
          - no multicast-to-multicast
          - priority dr-priority
          - no priority
          - [no] shutdown
          - sticky-dr [priority dr-priority]
          - no sticky-dr
          - [no] three-way-hello
          - [no] tracking-support
          - unicast-to-multicast unicast-start ip-address unicast-end ip-
address destination ip-address to-multicast ip-address
            - no unicast-to-multicast
            - [no] ipv4-multicast-disable
            - [no] non-dr-attract-traffic
            - rp
              - [no] anycast rp-ip-address
                - [no] rp-set-peer ip-address
              - [no] auto-rp-discovery
              - bootstrap-export policy-name [policy-name...(up to 5 max)]
              - no bootstrap-export
              - bootstrap-import policy-name [policy-name...(up to 5 max)]
              - no bootstrap-import
              - bsr-candidate
                - address ip-address
                - no address
                - hash-mask-len hash-mask-length

```

```

- no hash-mask-len
- priority bootstrap-priority
- no priority
- [no] shutdown
- rp-candidate
- address ip-address
- no address
- [no] group-range {grp-ip-address/mask | grp-ip-address [netmask]}
- holdtime holdtime
- no holdtime
- priority priority
- no priority
- [no] shutdown
- static
- [no] address ip-address
- [no] group-prefix {grp-ip-address/mask | grp-ip-address netmask}
- [no] override
- [no] shutdown
- spt-switchover-threshold {grp-ip-address/mask | grp-ip-address netmask} spt-
threshold
- no spt-switchover-threshold {grp-ip-address/mask | grp-ip-address netmask}
- [no] ssm-default-range-disable ipv4
- [no] ssm-groups
- [no] group-range {grp-ip-address/mask | grp-ip-address netmask}

```

7.8.1.1.7 RIP commands

7.8.1.1.7.1 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

```

7.8.1.1.7.2 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

```

7.8.1.1.7.3 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

```

7.8.1.1.8 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

```

7.8.1.1.9 VPRN raw socket IP transport configuration commands

```

config
- service
  - [no] vprn service-id [customer customer-id] [create]
  - ip-transport ipt-id [create]
  - no ip-transport ipt-id
    - description description-string
    - no description
    - dscp dscp-name
    - fc fc-name [profile {in | out}]
    - [no] filter-unknown-host
    - local-host ip-addr ip-addr port-num port-num protocol {tcp | udp}
    - no local-host
    - remote-host host-id [ip-addr ip-addr] [port-num port-num] [create]
    - no remote-host
      - description description-string
      - no description
      - name host-name
      - no name
    - [no] shutdown
  - tcp
    - inactivity-timeout seconds
    - max-retries number
    - retry-interval seconds

```

7.8.1.1.10 Multicast VPN commands

```

config
- service
  - vprn
    - mvpn
      - auto-discovery [default]
      - c-mcast-signaling bgp
      - mdt-type {sender-only | receiver-only | sender-receiver}
      - no mdt-type
      - provider-tunnel
        - inclusive
          - [no] mldp
          - [no] shutdown
        - selective
          - data-delay-interval value
          - no data-delay-interval
          - data-threshold {c-grp-ip-addr/mask | c-grp-ip-addr netmask} s-pmsi-
threshold
          - no data-threshold {c-grp-ip-addr/mask | c-grp-ip-addr netmask}
          - [no] maximum-p2mp-spmsi
          - [no] mldp

```

```

- [no] shutdown
- umh-selection {highest-ip | hash-based | unicast-rt-pref}
- no umh-selection
- vrf-export unicast
- vrf-export policy-name [policy-name...(up to 15 max)]
- no vrf-export
- vrf-import unicast
- vrf-import policy-name [policy-name ... (up to 15 max)]
- no vrf-import
- vrf-target {unicast | ext-community | export unicast | ext-community | import
unicast | ext-community}
- no vrf-target
- export {unicast | ext-community}
- import {unicast | ext-community}

```

7.8.1.1.11 MSDP commands

```

config
- service
- vprn service-id [customer customer-id] [create]
- no vprn service-id
- [no] msdp
- active-source-limit number
- no active-source-limit
- [no] data-encapsulation
- export policy-name [policy-name...(up to 5 max)]
- no export
- [no] group group-name
- active-source-limit number
- no active-source-limit
- export policy-name [policy-name...(up to 5 max)]
- no export
- import policy-name [policy-name...(up to 5 max)]
- no import
- local-address address
- no local-address
- mode {mesh-group | standard}
- [no] peer peer-address
- active-source-limit number
- no active-source-limit
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- [no] default-peer
- export policy-name [policy-name...(up to 5 max)]
- no export
- import policy-name [policy-name...(up to 5 max)]
- no import
- local-address address
- no local-address
- receive-msdp-msg-rate number interval seconds [threshold number]
- no receive-msdp-msg-rate
- [no] shutdown
- receive-msdp-msg-rate number interval seconds [threshold number]
- no receive-msdp-msg-rate
- [no] shutdown
- import policy-name [policy-name...(up to 5 max)]
- no import
- local-address address
- no local-address
- [no] peer peer-address
- active-source-limit number

```

```

- no active-source-limit
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- [no] default-peer
- export policy-name [policy-name...(up to 5 max)]
- no export
- import policy-name [policy-name...(up to 5 max)]
- no import
- local-address address
- no local-address
- receive-msdp-msg-rate number interval seconds [threshold number]
- no receive-msdp-msg-rate
- [no] shutdown
- receive-msdp-msg-rate number interval seconds [threshold number]
- no receive-msdp-msg-rate
- rpf-table {rtable-m | rtable-u | both}
- no rpf-table
- sa-timeout seconds
- no sa-timeout
- [no] shutdown
- [no] source ip-prefix/mask
  - active-source-limit number
  - no active-source-limit

```

7.8.1.1.12 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
        - [no] use-virtual-mac

```


7.8.1.1.13 Local DHCP and DHCPv6 server commands

For complete descriptions of all local DHCP and DHCPv6 server commands, see 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
- 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 [ip-address...(up to 4 max)]
- no dns-server
- domain-name domain-name
- no domain-name
- lease-rebind-time [days days] [hrs hours] [min minutes]
[sec seconds]
- no lease-rebind-time
- lease-renew-time [days days] [hrs hours] [min minutes]
[sec seconds]
- no lease-renew-time
- lease-time [days days] [hrs hours] [min minutes] [sec seconds]
- no lease-time
- netbios-name-server ip-address [ip-address...(up to 4 max)]
- no netbios-name-server
- netbios-node-type {B | P | M | H}
- no netbios-node-type
- subnet {ip-address/mask | ip-address netmask} [create]
- no subnet {ip-address/mask | ip-address netmask}
- [no] address-range start-ip-address end-ip-address
- [no] exclude-addresses start-ip-address [end-ip-address]
- maximum-declined maximum-declined
- no maximum-declined
- minimum-free minimum-free [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] ignore-rapid-commit
- 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
- options
- custom-option option-number address ipv6-address [ipv6-address...(up to 4 max)]
- custom-option option-number domain domain-string
- custom-option option-number hex hex-string
- custom-option option-number string ascii-string
- no custom-option option-number
- dns-server ipv6-address [ipv6-address...(up to 4 max)]
- no dns-server
- domain-name domain-name
- no domain-name
- prefix ipv6-address/prefix-length [pd] [wan-host] [create]
- no prefix ipv6-address/prefix-length
- options
- custom-option option-number address ipv6-address [ipv6-address...(up to 4 max)]
- custom-option option-number domain domain-string
- custom-option option-number hex hex-string
- custom-option option-number string ascii-string
- no custom-option option-number
- dns-server ipv6-address [ipv6-address...(up to 4 max)]
- no dns-server
- domain-name domain-name
- no domain-name
- preferred-lifetime [days days] [hrs hours] [min minutes] [sec seconds]
- no preferred-lifetime
- rebind-timer [days days] [hrs hours] [min minutes] [sec seconds]
- no rebind-timer
- renew-timer [days days] [hrs hours] [min minutes] [sec seconds]
- no renew-timer
- valid-lifetime [days days] [hrs hours] [min minutes] [sec seconds]
- no valid-lifetime
- server-id duid-en hex hex-string
- server-id duid-en string ascii-string
- server-id duid-ll
- no server-id
- [no] shutdown
- use-link-address [scope scope]
- no use-link-address
- [no] use-pool-from-client
- user-ident user-ident
- no user-ident

```

7.8.1.1.14 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
- cflowd-parameters
- sampling {unicast | multicast} type {interface} [direction {ingress-only
| egress-only | both}]
- no sampling {unicast | multicast}
- description description-string
- no description
- dhcp
- hold-time
- down ip seconds [init-only]
- down ipv6 seconds [init-only]
- no down ip
- no down ipv6
- up ip seconds
- no up ip
- up ipv6 seconds
- no up ipv6
- 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
- [no] load-balancing
- l4-load-balancing hashing-algorithm
- no l4-load-balancing
- [no] spi-load-balancing
- [no] teid-load-balancing
- [no] local-dhcp-server local-server-name
- [no] local-proxy-arp
- [no] loopback
- mac ieee-address
- no mac [ieee-address]
- [no] multicast-translation
- 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] [igmp-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
- tcp-mss value
- no tcp-mss
- unnumbered {ip-int-name | ip-address}
- no unnumbered
- vpls service-name
- no vpls
- vrrp (see Interface VRRP commands)

```

7.8.1.1.15 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 iaaid]
- 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
- unreachable [number seconds]
- no unreachable
- link-local-address ipv6-address [preferred]
- no link-local-address
- 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
- tcp-mss value
- no tcp-mss

```

```
- vrrp (see Interface VRRP commands)
```

7.8.1.1.16 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
```

7.8.1.1.17 Interface ICMP commands

```
config
- service
- vprn
- interface
- icmp
- [no] mask-reply
- ttl-expired number seconds
- no ttl-expired [number seconds]
- unreachable number seconds
- no unreachable [number seconds]
```

7.8.1.1.18 Interface SAP IPsec tunnel commands

For complete descriptions of all VPRN IPsec commands, see [IPsec command reference](#).

```
config
- service
- vprn
- interface
- sap
```

```
- ipsec-tunnel ipsec-tunnel-name [create]
```

7.8.1.1.19 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 ip-filter-id
- no v4-routed-override-filter
- v6-routed-override-filter ipv6-filter-id
- no v6-routed-override-filter
```

7.8.1.1.20 Interface VRRP commands

```
config
- service
- vprn
- interface ip-int-name
- [no] ipv6
- vrrp virtual-router-id [owner] [passive]
- no vrrp virtual-router-id
- [no] backup ipv6-address
- [no] bfd-enable service-id interface interface-name dst-ip ip-address
- [no] bfd-enable interface interface-name dst-ip 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] ntp-reply
- [no] ping-reply
- policy vrrp-policy-id
- no policy
- [no] preempt
- priority priority
- no priority
- [no] shutdown
- [no] standby-forwarding
- [no] telnet-reply
- [no] traceroute-reply
- vrrp virtual-router-id [owner] [passive]
- no vrrp virtual-router-id
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- [no] backup ip-address
- [no] bfd-enable service-id interface interface-name dst-ip ip-address
- [no] bfd-enable interface interface-name dst-ip 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] ntp-reply
- [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

```

7.8.1.1.21 VPRN static one-to-one NAT configuration commands

```

config
- service
- vprn
- [no] interface ip-int-name
- [no] static-nat-inside

```

```

config
- service
- vprn
- [no] static-nat
- [no] drop-packets-without-nat-entry
- inside
- map start ip-address end ip-address to ip-address
- no map start ip-address end ip-address
- [no] shutdown

```

7.8.1.1.22 TWAMP Light commands

```

config
- service
- vprn
- twamp-light
- reflector [udp-port udp-port-number] [create]
- no reflector
- description description-string
- [no] prefix ip-prefix/prefix-length [create]
- description description-string
- [no] shutdown

```

7.8.1.1.23 VPRN NTP commands

The **ntp-server** command is not supported in the **vprn ntp** context. When NTP is configured in a VPRN service, NTP server mode is assumed and is not optional.

```

config
- service

```



```

- vprn
  - [no] ntp
  - [no] authenticate
  - [no] authentication-check
  - authentication-key key-id key key [hash | hash2] type {des | message-digest}
  - no authentication-key key-id
  - authentication-keychain keychain-name
  - no authentication-keychain
  - broadcast {interface ip-int-name} [key-id key-id | authentication-
keychain keychain-name] [version version] [ttl ttl]
  - [no] broadcast {interface ip-int-name}
  - [no] shutdown

```

7.8.1.2 Show commands

For complete descriptions of all local DHCP and DHCPv6 server show commands, see the 7705 SAR Router Configuration Guide, "IP Router command reference, Show commands".

For complete descriptions of all OSPF and OSPFv3 show commands, see the 7705 SAR Routing Protocols Guide, "OSPF command reference, 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]
  - ip-transport ipt-id [detail | statistics]
    - remote-host host-id [detail | statistics]
  - macsec
  - 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]
  - ip-transport-using [ip-transport ipt-id]
  - 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] label-ipv4
    - damping [ip-prefix[/ip-prefix-length]] [damp-type] [detail] vpn-ipv4
    - damping [ip-prefix[/ip-prefix-length]] [damp-type] [detail] vpn-ipv6
    - damping [ip-prefix[/ip-prefix-length]] [damp-type] [detail] mvpn-ipv4
    - group [name] [detail]
    - inter-as-label (see Note)
    - neighbor [ip-address [detail]]
    - neighbor [as-number [detail]]
    - neighbor ip-address [family [type mvpn-type]] filter1 [brief]
    - neighbor ip-address [family] filter2

```

```

- neighbor as-number [family] filter2
- neighbor ip-address orf [filter3]
- neighbor ip-address graceful-restart
- neighbor [dynamic]
- next-hop [family] [ip-address] [detail]
- paths
- routes [ip-prefix/mask | ip-address]
- routes aspath-regex reg-exp {detail | longer}
- routes aspath-regex reg-exp
- routes aspath-regex reg-exp hunt
- routes brief
- routes community comm-id {detail | longer}
- routes community comm-id
- routes community comm-id hunt
- routes detail
- routes hunt [brief]
- routes ipv4 [aspath-regex reg-exp] [community comm-id] [brief] [all]
- routes ipv4 [aspath-regex reg-exp] hunt [community comm-id] [brief] [all]
- routes ipv4 [detail | longer] [aspath-regex reg-exp] [community comm-id] [all]
- routes ipv6 [aspath-regex reg-exp] [community comm-id] [brief] [all]
- routes ipv6 [aspath-regex reg-exp] hunt [community comm-id] [brief] [all]
- routes ipv6 [detail | longer] [aspath-regex reg-exp] [community comm-id] [all]
- routes label-ipv4 [aspath-regex reg-exp] [community comm-id] [brief] [all]
- routes label-ipv4 [aspath-regex reg-exp] hunt [community comm-id] [brief] [all]
- routes label-ipv4 [detail | longer] [aspath-regex reg-exp] [community comm-id]
[all]
- routes longer
- routes mvpn-ipv4 [aspath-regex reg-exp] [community comm-id] [rd rd] [brief]
[type mvpn-type] [originator-ip ip-address]
[source-ip ipv4 address | ipv6 address] [group-ip ipv4 address | ipv6 address] [source-as as-
number]
- routes mvpn-ipv4 [aspath-regex reg-exp] hunt [community comm-id] [rd rd] [brief]
[type mvpn-type] [originator-ip ip-address] [source-ip ipv4 address | ipv6 address] [group-
ip ipv4 address | ipv6 address] [source-as as-number]
- routes mvpn-ipv4 [detail | longer] [aspath-regex reg-exp] [community comm-id]
[rd rd] [type mvpn-type] [originator-ip ip-address] [source-ip ipv4 address | ipv6 address]
[group-ip ipv4 address | ipv6
address] [source-as as-number]
- routes route-target [source-as as-number] [brief] [aspath-regex reg-exp]
[community comm-id]
- routes route-target [rtc-prefix rtc-prefix] [hunt] [brief] [aspath-regex reg-exp]
[community comm-id]
- routes route-target rtc-prefix rtc-prefix [aspath-regex reg-exp] [community comm-
id]
- routes route-target [rtc-prefix rtc-prefix] [detail | longer] [aspath-regex reg-
exp] [community comm-id]
- routes vpn-ipv4 [aspath-regex reg-exp] [community comm-id] [rd rd] [brief]
- routes vpn-ipv4 [aspath-regex reg-exp] hunt [community comm-id] [rd rd] [brief]
- routes vpn-ipv4 [detail | longer] [aspath-regex reg-exp] [community comm-id]
[rd rd]
- routes vpn-ipv6 [aspath-regex reg-exp] [community comm-id] [rd rd] [brief]
- routes vpn-ipv6 [aspath-regex reg-exp] hunt [community comm-id] [rd rd] [brief]
- routes vpn-ipv6 [detail | longer] [aspath-regex reg-exp] [community comm-id]
[rd rd]
- summary [all]
- summary [family family] [neighbor ip-address]
- dhcp
- local-dhcp-server server-name
- associations
- declined-addresses ip-address[/mask] [detail]
- declined-addresses pool pool-name
- free-addresses ip-address[/mask]
- free-addresses summary [subnet ip-address[/mask]]
- free-addresses pool pool-name

```

```

- leases [detail]
- leases ip-address[/mask] address-from-user-db [detail]
- leases ip-address[/mask] [detail] [state]
- leases ip-address[/mask] dhcp-host dhcp-host-name [detail]
- pool-ext-stats [pool-name]
- server-stats
- subnet-ext-stats ip-address[/mask]
- subnet-ext-stats pool pool-name
- subnet-stats ip-address[/mask]
- subnet-stats pool pool-name
- summary
- servers [all]
- statistics [interface ip-int-name | ip-address]
- summary
- dhcp6
- local-dhcp-server server-name
- associations
- leases [ipv6-address/prefix-length] [type] [state] [detail]
- pool-ext-stats [pool-name]
- pool-stats [pool-name]
- prefix-ext-stats ipv6-address/prefix-length
- prefix-ext-stats pool pool-name
- prefix-stats ipv6-address/prefix-length
- prefix-stats pool pool-name
- server-stats
- summary
- servers [all]
- statistics
- summary
- interface {[ip-address | ip-int-name] [detail]} | summary | exclude-services]
- msdp
- group [group-name] [detail]
- peer [ip-address] [group group-name] [detail]
- source [ip-address/mask] [type {configured | dynamic | both}] [detail]
- source-active [{group ip-address | local | originator ip-address | peer ip-
address | source ip-address | group ip-address source ip-address}] [detail]
- source-active-rejected peer-group name [group ip-address] [source ip-address]
[originator ip-address] [peer ip-address]
- statistics [peer ip-address]
- status
- ospf [all]
- area [area-id] [detail] [lfa]
- capabilities [router-id]
- database [type {router | network | summary | asbr-summary | external | nssa |
all} [area area-id] [adv-router router-id] [link-state-id] [detail] [filtered]
- interface [area area-id] [detail]
- interface [ip-int-name | ip-address] [detail]
- interface [ip-int-name | ip-address] database [detail]
- lfa-coverage
- neighbor [ip-int-name | ip-address] [detail]
- neighbor overview
- neighbor [remote ip-address] [detail]
- opaque-database [area area-id | as] [adv-router router-id] [ls-id] [detail]
- prefix-sids [ip-prefix[/prefix-length]] [sid sid] [adv-router router-id]
- range [area-id]
- routes [ip-prefix[/prefix-length]] [type] [detail] [alternative] [summary]
[exclude-shortcut]
- spf [lfa]
- statistics
- status
- virtual-link database [detail]
- virtual-link [detail]
- virtual-neighbor [remote ip-address] [detail]
- ospf3 [all]

```

```

- area [area-id] [detail] [lfa]
- capabilities [router-id]
- database [type database-type] [area area-id] [adv-router router-id] [link-state-
id] [detail] [filtered]
- interface [area area-id] [detail]
- interface [ip-int-name | ip-address | ipv6-address] [detail]
- interface [ip-int-name | ip-address | ipv6-address] database [detail]
- lfa-coverage
- neighbor [ip-int-name] [router-id] [detail]
- neighbor overview
- range [area-id]
- routes [ip-prefix[/pfx-len]] [type] [detail] [alternative] [summary]
- spf [lfa]
- statistics
- status
- virtual-link [detail]
- virtual-neighbor [remote ipv6-address] [detail]
- route-table [family] [ip-prefix[/prefix-length]] [longer | exact | protocol protocol-
name] [all]] [next-hop-type type] [alternative]
- route-table [family] summary
- route-table [family] [ip-prefix[/prefix-length]] [longer | exact | protocol protocol-
name] extensive [all]
- 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 [family] [ip-prefix /mask] | preference preference | next-hop ip-address
| tag tag] [detail]
- tunnel-table summary [ipv4 | ipv6]
- tunnel-table [protocol protocol] {ipv4 | ipv6}
- tunnel-table [ip-prefix[/mask]] [alternative] [ipv4 | ipv6] detail
- tunnel-table [ip-prefix[/mask]] [alternative]
- tunnel-table [ip-prefix[/mask]] protocol protocol [detail]
- tunnel-table [ip-prefix[/mask]] sdp sdp-id

```



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.

7.8.1.3 Clear commands

```

clear
- router [service-id]
  - msdp
  - cache [peer ip-address] [group ip-address] [source ip-address] [originrp ip-
address]
  - statistics [peer ip-address]
- 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}]
  - ip-transport ipt-id
    - remote-host host-id
    - statistics
  - 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

```

7.8.1.4 Debug commands

```

debug
- router [service-id]
  - [no] msdp
    - packet [pkt-type] [peer ip-address]
    - no packet
    - pim [grp-address]
    - no pim
    - rtm [rp-address]
    - no rtm
    - sa-db [group grpAddr] [source srcAddr] [rp rpAddr]
    - no sa-db
  - 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

```

7.8.2 Command descriptions

- [VPRN service configuration commands](#)
- [Show service commands](#)
- [Show router commands](#)
- [Clear service commands](#)
- [Debug MSDP commands](#)
- [Debug service commands](#)

7.8.2.1 VPRN service configuration commands

- [Generic commands](#)
- [Global commands](#)
- [BGP commands](#)
- [OSPF and OSPFv3 commands](#)
- [IGMP commands](#)
- [PIM commands](#)
- [RIP commands](#)
- [VPRN security configuration commands](#)
- [VPRN raw socket IP transport configuration commands](#)
- [Multicast VPN commands](#)
- [MSDP 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](#)
- [VPRN static one-to-one NAT configuration commands](#)
- [TWAMP Light commands](#)
- [VPRN NTP commands](#)

7.8.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

```
config>service>vpn
config>service>vpn>bgp
config>service>vpn>bgp>group
config>service>vpn>bgp>group>neighbor
config>service>vpn>dhcp>local-dhcp-server
config>service>vpn>dhcp>local-dhcp-server>pool
config>service>vpn>interface
config>service>vpn>if>dhcp
config>service>vpn>if>ipv6>dhcp6-relay
config>service>vpn>if>sap
config>service>vpn>ip-transport
config>service>vpn>ip-transport>remote-host
config>service>vpn>rip
config>service>vpn>rip>group
config>service>vpn>rip>group>neighbor
config>service>vpn>static-route-entry>black-hole
config>service>vpn>static-route-entry>grt
config>service>vpn>static-route-entry>indirect
config>service>vpn>static-route-entry>ipsec-tunnel
config>service>vpn>static-route-entry>next-hop
config>service>vpn>twamp-light>reflector
config>service>vpn>twamp-light>reflector>prefix
config>service>vpn>zone
config>service>vpn>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 (such as #, \$, or spaces), 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>ip-transport
config>service>vprn>msdp
config>service>vprn>msdp>group
config>service>vprn>msdp>group>peer
config>service>vprn>msdp>peer
config>service>vprn>ntp
config>service>vprn>ospf
config>service>vprn>ospf>area>interface
config>service>vprn>ospf>area>sham-link
```



```
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf3
config>service>vprn>ospf3>area>interface
config>service>vprn>ospf3>area>virtual-link
config>service>vprn>pim
config>service>vprn>pim>interface
config>service>vprn>rip
config>service>vprn>rip>group
config>service>vprn>rip>group>neighbor
config>service>vprn>router-advertisement>interface
config>service>vprn>pim>rp>bsr-candidate
config>service>vprn>pim>rp>rp-candidate
config>service>vprn>spoke-sdp
config>service>vprn>static-route-entry>black-hole
config>service>vprn>static-route-entry>grt
config>service>vprn>static-route-entry>indirect
config>service>vprn>static-route-entry>ipsec-tunnel
config>service>vprn>static-route-entry>next-hop
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.

VPRN IP transport subservice

when an IP transport subservice within a VPRN service is shut down, all TCP/UDP packets received from remote hosts are dropped and any serial data received from the serial port is dropped. Any TCP connections that were up are closed and no new TCP connection requests are accepted.

It is not possible to make configuration changes to an IP transport subservice without performing a **shutdown** first.

The operational state of an IP transport subservice is relative to the operational state of the serial port for which the IP transport subservice is defined. When a serial port is shut down, the IP transport subservice associated with the serial port becomes operationally down.

When the **no shutdown** command is executed for an IP transport subservice, it becomes operationally up, serial data from the serial port is encapsulated in TCP/UDP packets destined for remote hosts, and TCP/UDP packets can be received by the local host, where raw serial data is then sent out the serial port.

7.8.2.1.2 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**] [**as-set**] [**aggregator** *as-number:ip-address*]
[**description** *description-text*]

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

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.

as-set

creates an aggregate where the path advertised for this route is an AS_SET consisting of all elements contained in all paths that are being summarized. This parameter should be used carefully as it can increase the amount of route churn due to best path changes. The parameter is only applicable to BGP.

as-number:ip-address

specifies the BGP aggregator path attribute to the aggregate route. When configuring the aggregator, a two-octet AS number used to form the aggregate route must be entered, followed by the IP address of the BGP system that created the aggregate route.

Values *as-number* 1 to 4294967295

ip-address a.b.c.d

description-text

a text description, up to 80 characters long, stored in the configuration file for a configuration context

auto-bind-tunnel

Syntax

auto-bind-tunnel

Context

config>service>vprn

Description

This command enables the context to configure automatic binding of a VPRN service using tunnels to MP-BGP peers.

The **auto-bind-tunnel** mode is simply a context to configure the binding of VPRN routes to tunnels. The user must configure the **resolution** option to enable autobinding resolution to tunnels in TTM.

When an explicit SDP to a BGP next hop is configured in a VPRN service (**config>service>vprn>spoke-sdp**), it overrides the **auto-bind-tunnel** selection for that BGP next hop only. There is no support for reverting automatically to the **auto-bind-tunnel** selection if the explicit SDP goes down. The user must delete the explicit spoke SDP in the VPRN service context to resume using the **auto-bind-tunnel** selection for the BGP next hop.

ecmp

Syntax

ecmp *max-ecmp-routes*

no ecmp

Context

config>service>vprn>auto-bind-tunnel

Description

This command configures the maximum number of routes that can be used for autobind tunnel resolution.

The **no** form of this command removes the configured value.

Parameters

max-ecmp-routes

the maximum number of routes that can be used for autobind tunnel resolution

Values 1 to 8

Default 1

resolution

Syntax

resolution {**any** | **filter** | **disabled**}

Context

config>service>vprn>auto-bind-tunnel

Description

This command configures the resolution mode in the automatic binding of a VPRN service to tunnels to MP-BGP peers.

If the **resolution** option is explicitly set to **disabled**, the autobinding to tunnels is removed.

If **resolution** is set to **any**, any supported tunnel type in the VPRN context will be selected following the TTM preference. If one or more explicit tunnel types are specified using the **resolution-filter** option, only these tunnel types will be selected again following the TTM preference.

The user must set **resolution** to **filter** to activate the list of tunnel types configured under **resolution-filter**.

Parameters

any

enables the binding to any supported tunnel type in the VPRN context following the TTM preference

filter

enables the binding to the subset of tunnel types configured under **resolution-filter**

disabled

disables the automatic binding of a VPRN service to tunnels to MP-BGP peers

resolution-filter

Syntax

resolution-filter

Context

config>service>vprn>auto-bind-tunnel

Description

This command configures the subset of tunnel types that can be used in the resolution of VPRN prefixes within the automatic binding of VPRN service to tunnels to MP-BGP peers.

The following tunnel types are supported in a VPRN context (in order of preference): RSVP (**rsvp**), segment routing TE (**sr-te**), LDP (**ldp**), segment routing OSPF (**sr-ospf**), segment routing IS-IS (**sr-isis**), and GRE (**gre**). The segment routing precedences can be configured. The selection of an SR tunnel in SR-ISIS when using multi-instance IS-IS is based on lowest instance ID.

gre

Syntax

[no] gre

Context

config>service>vprn>auto-bind-tunnel>resolution-filter

Description

This command specifies the GRE type of automatic binding for the SDP assigned to this service. When **auto-bind-tunnel** is used, a spoke SDP does not need to be configured for the service.

The **no** form of the command removes this type of automatic binding.

Default

no gre

ldp**Syntax**

[no] ldp

Context

config>service>vprn>auto-bind-tunnel>resolution-filter

Description

This command specifies the LDP tunnel type of automatic binding for the SDP assigned to this service. When **auto-bind-tunnel** is used, a spoke SDP does not need to be configured for the service.

The **ldp** value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next hop.

The **no** form of the command removes this type of automatic binding.

Default

no ldp

rsvp**Syntax**

[no] rsvp

Context

config>service>vprn>auto-bind-tunnel>resolution-filter

Description

This command specifies the RSVP tunnel type of automatic binding for the SDP assigned to this service. When **auto-bind-tunnel** is used, a spoke SDP does not need to be configured for the service.

The **rsvp** value instructs BGP to search for the best metric RSVP LSP to the address of the BGP next hop. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest **tunnel-id**.

The **no** form of the command removes this type of automatic binding.

Default

no rsvp

sr-isis

Syntax

[no] sr-isis

Context

config>service>vprn>auto-bind-tunnel>resolution-filter

Description

This command specifies the SR-ISIS tunnel type of automatic binding for the SDP assigned to this service. When **auto-bind-tunnel** is used, a spoke SDP does not need to be configured for the service.

When the **sr-isis** value is enabled, an SR tunnel to the BGP next hop is selected in the TTM from the lowest-numbered IS-IS instance.

The **no** form of the command removes this type of automatic binding.

Default

no sr-isis

sr-ospf

Syntax

[no] sr-ospf

Context

config>service>vprn>auto-bind-tunnel>resolution-filter

Description

This command specifies the SR-OSPF tunnel type of automatic binding for the SDP assigned to this service. When **auto-bind-tunnel** is used, a spoke SDP does not need to be configured for the service.

When the **sr-ospf** value is enabled, an SR tunnel to the BGP next hop is selected in the TTM from OSPF instance 0.

The **no** form of the command removes this type of automatic binding.

Default

no sr-ospf

sr-te

Syntax

[no] sr-te

Context

```
config>service>vprn>auto-bind-tunnel>resolution-filter
```

Description

This command specifies the SR-TE tunnel type of automatic binding for the SDP assigned to this service. When **auto-bind-tunnel** is used, a spoke SDP does not need to be configured for the service.

The **sr-te** value instructs the 7705 SAR to search for the best metric SR-TE LSP to the address of the BGP next hop. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple SR-TE LSPs with the same lowest metric, BGP selects the LSP with the lowest **tunnel-id**.

The **no** form of the command removes this type of automatic binding.

Default

```
no sr-te
```

weighted-ecmp

Syntax

```
[no] weighted-ecmp
```

Context

```
config>service>vprn>auto-bind-tunnel
```

Description

This command enables weighted ECMP for packets using tunnels that a VPRN automatically binds to. This command is applicable if the autobind tunnel is configured for RSVP or SR-TE using the **config>service>vprn>auto-bind-tunnel>resolution-filter>rsvp/sr-te** command. When weighted ECMP is enabled, packets are sprayed across RSVP-TE or SR-TE LSPs in the ECMP set according to the outcome of the hash algorithm and the configured **load-balancing-weight** of each LSP. See the 7705 SAR MPLS Guide, "MPLS Commands", for more information about the **load-balancing-weight** command.

The **no** form of the command disables weighted ECMP for next-hop tunnel selection.

Default

```
no weighted-ecmp
```

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 (Equal-Cost Multipath Protocol) in the VPRN service context of a VPRN service 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 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 [preference](#) command for information about 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.

Default

no ecmp

Parameters

max-ecmp-routes

specifies 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 1 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 after 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

entropy-label

Syntax

[no] entropy-label

Context

config>service>vprn

config>service>vprn>interface>spoke-sdp

Description

This command enables or disables the use of entropy labels for spoke SDPs on a VPRN.

If **entropy-label** is enabled, the entropy label and entropy label indicator (ELI) are inserted in packets for which at least one LSP in the stack for the far end of the tunnel used by the service has advertised entropy label capability.

If the tunnel type is RSVP-TE, **entropy-label** can also be controlled by disabling **entropy-label-capability** under the **config>router>rsvp** or **config>router>mpls>lsp** contexts at the far-end LER.

When the **entropy-label** and **entropy-label-capability** commands are both enabled, the entropy label value inserted at the iLER is always based on the service ID.

The entropy label and the hash label features are mutually exclusive. The entropy label cannot be configured on a spoke SDP or service where the hash label feature has already been configured.

Default

no entropy-label

grt-lookup

Syntax

grt-lookup

Context

config>service>vprn

Description

This command enters the context under which all GRT leaking commands are configured.

enable-grt

Syntax

[no] enable-grt

Context

```
config>service>vprn>grt-lookup
```

Description

This command enables the functions required for looking up routes in the GRT when the lookup in the local VRF fails. If this command is enabled without the use of the **config>service>vprn>static-route-entry>grt** command, a lookup in the local VRF is preferred over the GRT. When the local VRF returns no route table lookup matches, the result from the GRT is preferred.

The **no** form of this command disables the global routing table lookup function when the lookup in the local VRF fails.

Default

```
no enable-grt
```

allow-local-management

Syntax

```
[no] allow-local-management
```

Context

```
config>service>vprn>grt-lookup>enable-grt
```

Description

This command enables management traffic from GRT leaking-enabled VPRN instances to reach local interfaces in the base router instance. The local interfaces can be system IP interfaces or loopback interfaces. Management traffic is traffic generated by Telnet, SNMP, and SSH. For a complete list of supported management protocols, see [Table 128: IPv4 and IPv6 GRT-supported management protocols](#).

Ping and traceroute responses from the base router interfaces are supported but are not configurable. The **allow-local-management** command does not control the support for management protocols terminating on VPRN interfaces directly.

Default

```
no allow-local-management
```

export-grt

Syntax

```
export-grt policy-name [policy-name...(up to 5 max)]
```

```
no export-grt
```

Context

```
config>service>vprn>grt-lookup
```

Description

This command uses configured route policies to determine which routes are exported from the VRF to the GRT along with all the forwarding information.

On network egress, packets with a source IP address that matches the 7705 SAR system IP address and the destination IP address of the far-end node must perform a GRT lookup in order to be resolved. A route policy can be configured with the IP address prefix or loopback address of the far-end router and with the action to accept. This policy is configured under the **config>router>policy-options** context, and is installed in the GRT FIB using the **export-grt** command. The route installed in the GRT FIB will have a next hop of the IPsec tunnel.

Up to five policies can be exported to the GRT FIB.

The **no** form of the command restores the default of not exporting routes to the GRT FIB.

Default

no export-grt

Parameters

policy-name

the name of the route policy to be exported to the GRT FIB

export-limit

Syntax

export-limit *num-routes*

no export-limit

Context

config>service>vprn>grt-lookup

Description

This command limits the number of IPv4 routes that can be exported from the VRF to the GRT. Setting the limit to 0 overrides the maximum limit. Setting the value to 0 does not limit the number of routes exported from the VRF to the GRT. Configuring the **export-limit** between 1 and 256 will limit the number of routes to the specified value.

The **no** form of the command resets the limit to the default of allowing five routes per route policy to be exported from the VRF to the GRT.

Default

5

Parameters

num-routes

the number of routes per policy to be exported to the GRT

Values 0 to 256

export-v6-limit

Syntax

export-v6-limit *num-routes*

no export-v6-limit

Context

config>service>vprn>grt-lookup

Description

This command limits the number of IPv6 routes that can be exported from the VRF to the GRT. Setting the limit to 0 overrides the maximum limit. Setting the value to 0 does not limit the number of routes exported from the VRF to the GRT. Configuring the **export-limit** between 1 and 256 will limit the number of routes to the specified value.

The **no** form of the command resets the limit to the default of allowing five routes per route policy to be exported from the VRF to the GRT.

Default

5

Parameters

num-routes

the number of IPv6 routes per policy to be exported to the GRT

Values 0 to 256

maximum-ipv6-routes

Syntax

maximum-ipv6-routes *number* [**log-only**] [**threshold percent**]

no maximum-ipv6-routes

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, static, 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 ($[\text{max} + \text{mid}] / 2$). 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, static, 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 ($[\text{max} + \text{mid}] / 2$). 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

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*} [**fc-queue** *fc-name* **profile** {**in** | **out**}]

application *dot1p-app-name* **dot1p** {*dot 1p-priority*} [**fc-queue** *fc-name* **profile** {**in** | **out**}]

no application {*dscp-app-name* | *dot1p-app-name*}

Context

config>service>vprn>sgt-qos

Description

This set of commands configures DSCP marking for self-generated IP traffic or dot1p marking for self-generated non-IP traffic (specifically, IS-IS and ARP traffic). See the following table for supported DSCP marking and defaults per application.

Table 135: Applications and support for configurable DSCP or dot1p markings

Application	Supported marking	Default DSCP/dot1p
ARP	dot1p	7
IS-IS	dot1p	7
BGP	DSCP	NC1
DHCP	DSCP	NC1
DNS	DSCP	AF41
FTP	DSCP	AF41
ICMP (ping)	DSCP	BE
IGMP	DSCP	NC1
LDP (T-LDP)	DSCP	NC1
MLD	DSCP	NC1
NDIS	DSCP	NC1
NTP	DSCP	NC1
OSPF	DSCP	NC1
PIM	DSCP	NC1
1588 PTP	DSCP	NC1
RADIUS	DSCP	AF41
RIP	DSCP	NC1
RSVP	DSCP	NC1
SNMP (get, set, etc.)	DSCP	AF41
SNMP trap/log	DSCP	AF41
SSH (SCP)	DSCP	AF41
syslog	DSCP	AF41
TACACS+	DSCP	AF41
Telnet	DSCP	AF41

Application	Supported marking	Default DSCP/dot1p
TFTP	DSCP	AF41
Traceroute	DSCP	BE
VRRP	DSCP	NC1

When an IP or Layer 3 application is configured using the *dscp-app-name* parameter, the specified DSCP name or DSCP value is used for all packets generated by this application within the router instance in which it is configured. The value set in this command sets the DSCP value in the egress IP header. The egress QoS policy will not overwrite this value.

When a Layer 2 application is configured using the *dot1p-app-name* parameter, the specified dot1p priority value is used for all packets generated by this application within the router instance in which it is configured.

Only one name or value can be configured per application. If multiple entries are configured, a subsequent entry overrides the previously configured entry.

The **fc-queue** option redirects SGT applications to egress data queues rather than the default control queue by assigning them to a forwarding class. If this option is configured, the profile state must be set. All packets that are assigned to this forwarding class will be considered in-profile or out-of-profile based on the configuration. In case of congestion, the in-profile packets are preferentially queued over the out-of-profile packets.

If the **fc-queue** option is used with the *dscp-app-name* application, any configuration done using the **sgt-qos>dscp** command is ignored for packets generated by this application, as illustrated in the following examples:

```
sgt-qos>application telnet dscp cp1
```

```
sgt-qos>dscp cp1 fc af
```

```
sgt-qos>application ftp dscp cp1 fc-queue be profile out
```

```
sgt-qos>dscp cp1 fc af
```

In the first example, all packets generated by the Telnet application use DSCP CP1 and map to FC AF as configured in the **dscp** command. The dot1p bits of the outgoing packets are marked from the value that FC AF points to in the egress QoS policy.

In the second example, all packets generated by the FTP application use DSCP CP1 and map to FC BE as dictated by the **fc-queue** redirection. The dot1p bits of the outgoing packets are marked from the value that FC BE points to in the egress QoS policy. Because redirection is configured, the mapping configured with the **dscp** command is ignored.



Note: The above behavior applies to all SGT IP applications with the exception of VRRP, where the dot1p value is always set to 7, regardless of the value in the FC egress QoS policy.

If the **fc-queue** option is used with the *dot1p-app-name* application, the dot1p bits of the outgoing packets are marked with the value set with the *dot1p-priority* parameter, regardless of the value in the FC egress queue policy.

The **no** form of this command resets the DSCP or dot1p value for the application to its default value and resets the application to use the egress control queue.

Default

n/a

Parameters

dscp-app-name

the DSCP application name

Values bgp, dhcp, dns, ftp, icmp, igmp, ldp, mld, ndis, ntp, ospf, pim, ptp, radius, rip, rsvp, snmp, snmp-notification, ssh, syslog, tacplus, telnet, tftp, traceroute, vrrp



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-value

the value that maps to the DSCP name (the value **none** specifies that the default DSCP value for the application be used)

Values none | 0 to 63

dscp-name

the DSCP to be associated with the forwarding class. [Table 139: Valid DSCP names](#) lists the valid DSCP names.

dot1p-app-name

the dot1p application name

Values arp, isis

dot1p-priority

the dot1p priority (the value **none** specifies that the default dot1p value for the application be used)

Values none | 0 to 7

fc-name

the forwarding class assigned to SGT applications redirected to data queues

Values be, l2, af, l1, h2, ef, h1, nc

profile {in | out}

the profile state of packets assigned to the specified forwarding class; this parameter must be specified when the **fc-queue** parameter is configured

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. The forwarding class dot1p SAP egress QoS policy mapping is used to mark the dot1p bits of the Layer 3 or IP application. For example, configuring the *dscp-name* parameter as **be** and the *fc-name* parameter as **l1** results in marking the dot1p bits of the outgoing Ethernet frame, which is transporting self-generated IP traffic with DSCP bits set to BE, to the value that FC L1 points to in the SAP egress QoS policy (as configured in the **config>qos>sap-egress>fc** context).




Note: The dot1p class of service may not apply to all IP traffic and is dependent on the egress port encapsulation type.

Based on this configured FC, the SAP egress QoS policy for the egress forwarding complex sets the IEEE 802.1 dot1p bits.

Multiple commands can be entered to associate some or all of the 64 DSCP values with the forwarding class. For undefined code points, packets are assigned to the default forwarding class for the DSCP value. The following table lists the default forwarding class for each DSCP value.

The **no** form of the command resets the DSCP value to its default forwarding class.



Note: If the **fc-queue** option is configured in the **sgt-qos>application dscp-app-name** command, the mapping created with this command is ignored for packets generated by the applications that are configured with the option.

Table 136: DSCP-to-default forwarding class mapping

DSCP value	Default FC
be	nc
cp1	be
cp2	be
cp3	be
cp4	be
cp5	be
cp6	be

DSCP value	Default FC
cp7	be
cs1	be
cp9	be
af11	af
cp11	be
af12	af
cp13	be
af13	af
cp15	be
cs2	be
cp17	be
af21	l1
cp19	be
af22	l1
cp21	be
af23	l1
cp23	be
cs3	be
cp25	be
af31	l1
cp27	be
af32	l1
cp29	be
af33	l1
cp31	be
cs4	be
cp33	be

DSCP value	Default FC
af41	nc
cp35	be
af42	h2
cp37	be
af43	h2
cp39	be
cs5	be
cp41	be
cp42	be
cp43	be
cp44	be
cp45	be
ef	ef
cp47	be
nc1	nc
cp49	be
cp50	h2
cp51	be
cp52	be
cp53	be
cp54	be
cp55	be
nc2	nc
cp57	be
cp58	be
cp59	be
cp60	be

DSCP value	Default FC
cp61	be
cp62	be
cp63	be

Default

See the table for the default forwarding class for each DSCP value.

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

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 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 the IPv4 address

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

spoke-sdp

Syntax

[no] **spoke-sdp** *sdp-id*

Context

config>service>vprn

Description

This command binds a service to an existing service destination 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, the autobind feature can be used. With **auto-bind-tunnel**, no **vprn>spoke-sdp** configuration is required. When both **auto-bind-tunnel** and **spoke-sdp** are configured, **spoke-sdp** takes precedence. The **spoke-sdp** configuration must be deconfigured for the autobind feature 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 or 7750 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-entry

Syntax

static-route-entry {*ip-prefix*|*prefix-length*}

no static-route-entry {*ip-prefix*|*prefix-length*}

Context

config>service>vprn

Description

This command creates a static route entry within the associated router instance. A prefix and prefix length must be specified.

Once the static route context for the specified prefix and length has been created, additional parameters associated with the static routes may be specified.

When configuring a static route, multiple types of static routes (**blackhole**, **grt**, **indirect**, **ipsec-tunnel**, and **next-hop**) can be applied to the same IPv4 or IPv6 prefix. If a static route that is forwarding traffic goes down, the default route will be used instead. The **preference** parameter specifies the order in which the

routes are applied. If a blackhole static route has the same preference as another route with the same prefix, the blackhole route takes a lower precedence.

Before the static route entry can be deleted, the next hops associated with the prefix must be shut down and deleted.

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, as many parameters as are necessary to uniquely identify the static route must be entered.

Default

no static-route-entry

Parameters

ip-prefix/prefix-length

the destination address of the static route

black-hole

Syntax

[no] **black-hole**

Context

config>service>vprn>static-route-entry

Description

This command specifies that the route is a blackhole route. If the destination address on a packet matches this static route, it will be silently discarded.

If the 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 defaults for **preference** and **metric** are applied.

Before the static route entry can be deleted, the next hops associated with the prefix must be shut down and deleted.

Default

no black-hole

metric

Syntax

[no] **metric** *metric*

Context

config>service>vprn>static-route-entry>black-hole

```
config>service>vpn>static-route-entry>grt
config>service>vpn>static-route-entry>indirect
config>service>vpn>static-route-entry>ipsec-tunnel
config>service>vpn>static-route-entry>next-hop
```

Description

This command specifies the cost (metric) for the static route, expressed as a decimal integer. This value is used when importing the static route into other protocols such as OSPF. When modifying the metric of an existing static route, the preference will not change unless specified.

This value is also used to determine which static route to install in the forwarding table.

- If there are multiple static routes with the same preference but different metrics, the lower-cost (lower metric) route will be installed.
- If there are multiple static routes with equal preference and metrics, the 7705 SAR chooses the route with the lowest next-hop IP address as the best route.
- If there are multiple routes with unequal preferences, the lower preference route is installed.

The **no** form of this command returns the metric to the default value.

Default

no metric

Parameters

<i>metric</i>	
	the metric value
Values	0 to 65535
Default	1

preference

Syntax

```
preference preference
no preference
```

Context

```
config>service>vpn>static-route-entry>black-hole
config>service>vpn>static-route-entry>grt
config>service>vpn>static-route-entry>indirect
config>service>vpn>static-route-entry>ipsec-tunnel
config>service>vpn>static-route-entry>next-hop
```

Description

This command specifies the preference of this static route over routes from different sources such as BGP or OSPF. The preference is expressed as a decimal integer. A route with a lower preference value is preferred over a route with a higher preference value.

When modifying the preference value of an existing static route, the metric will not change unless specified. The **preference** command is also used to prioritize static routes applied to the same prefix. If a blackhole static route has the same preference as another route with the same prefix, the blackhole route takes a lower precedence.

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, the decision of which route to use is determined by the **ecmp** command.

The following table shows the default route preference based on the route source.

Table 137: Default route preference

Label	Preference	Configurable
Direct attached	0	No
Static route	5	Yes
OSPF internal routes	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
BGP	170	Yes

The **no** form of this command returns the static route preference to its default value.

Default

5

Parameters

preference

the route preference value

Values 1 to 255

prefix-list

Syntax

[no] **prefix-list** *prefix-list-name* {**all** | **none**}

Context

config>service>vpn>static-route-entry>black-hole

config>service>vpn>static-route-entry>indirect

config>service>vpn>static-route-entry>next-hop

Description

This command adds a constraint to the static route such that the static route is only active if **none** or **all** of the prefixes in the prefix list are present and active in the route table.

Default

no prefix-list

Parameters

prefix-list-name

the name of a currently configured prefix list

all

specifies that the static route condition is met if all prefixes in the prefix list are present in the active route table

none

specifies that the static condition is met if none of the prefixes in the prefix list are present in the active route table

tag

Syntax

[no] **tag** *tag*

Context

config>service>vpn>static-route-entry>black-hole

config>service>vpn>static-route-entry>indirect

config>service>vpn>static-route-entry>ipsec-tunnel

config>service>vpn>static-route-entry>next-hop

Description

This command 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.

Default

1

Parameters

tag

specifies an integer tag value

Values 1 to 4294967295

grt

Syntax

[no] grt

Context

config>service>vprn>static-route-entry

Description

This command creates a static route in a VPRN service context that points to the global routing context (base router). This is primarily used to allow traffic that ingresses through a VPRN service to be routed out of the global routing context.

The **grt** type of next hop cannot be used in conjunction with any other next-hop types.

Default

no grt

indirect

Syntax

[no] indirect *ip-address*

Context

config>service>vprn>static-route-entry

Description

This command specifies that the route is indirect and specifies the next-hop IP address used to reach the destination.

The configured *ip-address* is not directly connected to a network configured on this node. The destination can be reached via multiple paths. The indirect address can only be resolved via a dynamic routing protocol. Another static route cannot be used to resolve the indirect address.

The *ip-address* can be either on the network side or the access side and is typically at least one hop away from the node.

Default

no indirect

Parameters

ip-address

the IPv4 or IPv6 address of the IP interface

cpe-check

Syntax

[no] **cpe-check** *cpe-ip-address*

Context

config>service>vprn>static-route-entry>indirect

config>service>vprn>static-route-entry>next-hop

Description

This command enables CPE connectivity check and specifies the IP address of the target CPE device. ICMP pings will be sent to this target IP address. This parameter must be configured to enable the CPE connectivity feature for the static route. The *cpe-ip-address* cannot be in the same subnet as the static route subnet to avoid possible circular references. CPE check and BFD support are mutually exclusive on a static route.

If a CPE connectivity check target address is already being used as the target address in a different static route, **cpe-check** parameters must match. If they do not match, the new configuration command will be rejected.

If a **static-route-entry>indirect** command or **static-route-entry>next-hop** command is issued with no **cpe-check** target but the destination *prefix/prefix-length* and the next hop match a static route that has an associated **cpe-check**, the **cpe-check** test is removed from the static route.

The **no** form of this command disables the **cpe-check** option.

Default

no cpe-check

Parameters

cpe-ip-address

the IPv4 or IPv6 address of the CPE device

drop-count

Syntax

drop-count *count*

no drop-count

Context

config>service>vprn>static-route-entry>indirect>cpe-check

config>service>vprn>static-route-entry>next-hop>cpe-check

Description

This command specifies the number of consecutive ping replies that must be missed to declare the CPE down and to deactivate the static route.

Default

3

Parameters

count

an integer count value

Values 1 to 255

interval

Syntax

interval *seconds*

no interval

Context

config>service>vprn>static-route-entry>indirect>cpe-check

config>service>vprn>static-route-entry>next-hop>cpe-check

Description

This command specifies the interval, in seconds, between ICMP pings to the target IP address.

Default

1

Parameters

seconds

an integer interval value

Values 1 to 255

log

Syntax

[no] log

Context

config>service>vprn>static-route-entry>indirect>cpe-check

config>service>vprn>static-route-entry>next-hop>cpe-check

Description

This command enables the logging of transitions between active and inactive routes based on the CPE connectivity check. Events will be sent to the system log, syslog, and SNMP traps.

Default

no log

ipsec-tunnel

Syntax

[no] ipsec-tunnel *ipsec-tunnel-name*

Context

config>service>vprn>static-route-entry

Description

This command creates a static route in a VPRN service context that points to an IPsec tunnel.

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 default values for the **preference** and **metric** commands are applied.

Default

no ipsec-tunnel

Parameters

ipsec-tunnel-name

the IPsec tunnel name; the IPsec tunnel specifies the local and peer gateway addresses for the tunnel

next-hop

Syntax

[no] next-hop {*ip-int-name* | *ip-address* | *ipv6-address*}

Context

config>service>vprn>static-route-entry

Description

This command specifies the directly connected next-hop IP address or interface used to reach the destination. If the next hop is over an unnumbered interface, the interface name of the unnumbered interface can be used.

The configured *ip-address* can be either on the network side or the access side on the node. The address must be associated with a network that is directly connected to a network configured on the node.

Default

no next-hop

Parameters

ip-int-name, *ip-address*, *ipv6-address*

the IP interface name, IPv4 address, or IPv6 address

bfd-enable

Syntax

[no] bfd-enable

Context

config>service>vprn>static-route-entry>next-hop

Description

This command associates the static route state with a BFD session between the local system and the configured next hop. The remote end of the BFD session must also be configured to originate or accept the BFD session controlling the static route state.

The **no** form of this command removes the association of the static route state with the BFD session.

Default

no bfd-enable

type**Syntax**

type hub

no type

Context

config>service>vpn

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>vpn

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.

Aggregate routes are not advertised via MP-BGP protocols to the other MP-BGP peers.

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>vpn

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>vpn

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

weighted-ecmp

Syntax

[no] **weighted-ecmp**

Context

config>service>vprn

Description

This command enables weighted load-balancing for OSPF ECMP routes for the VPRN instance. Weighted ECMP can be performed when all next hops are configured with non-zero load-balancing weights.

The **no** form of this command restores regular ECMP spraying of packets to OSPF route destinations.

Default

no weighted-ecmp

7.8.2.1.3 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

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

auth-keychain

Syntax

auth-keychain *name*

no auth-keychain

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Description

This command associates an authentication keychain with the BGP protocol. The keychain is a collection of keys used to authenticate BGP messages from remote neighbors. The keychain allows the rollover of authentication keys during the lifetime of a session and also supports stronger authentication algorithms than clear text and MD5.

The keychain must already be defined in the **config>system>security>keychain** context.

Either the **authentication-key** command or the **auth-keychain** command can be used by BGP, but both cannot be supported at the same time. If both commands are configured, the **auth-keychain** configuration will be applied and the **authentication-key** command will be ignored.

By default, authentication is not enabled.

Default

no auth-keychain

Parameters

name

the name of an existing keychain, up to 32 characters

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.

Either the **authentication-key** command or the **auth-keychain** command can be used by BGP, but both cannot be supported at the same time. If both commands are configured, the **auth-keychain** configuration will be applied and the **authentication-key** command will be ignored.

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.

backup-path

Syntax

```
[no] backup-path [ipv4] [ipv6]
```

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 or IPv6 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.

The **no** form of the command disables BGP FRR with PIC.

Default

no backup-path

Parameters

ipv4

enables a backup path for IPv4 BGP learned prefixes

ipv6

enables a backup path for IPv6 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]

always-compare-med strict-as [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 used without the **strict-as** option, the MEDs of two paths are always compared even if the paths have a different neighbor AS.

If the **strict-as** option is used, the MEDs of two paths are compared only if they come from the same neighboring AS.

The **zero** and **infinity** options specify how to treat paths that do not have a MED attribute; for example, **always-compare-med zero** means that if one path is missing a MED attribute, it is treated as though it had a MED attribute with the value of 0. If neither option is specified, the **zero** option is implied.

The **no** form of the command means that only the MEDs of paths that have the same neighbor AS 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.

strict-as

specifies that the MEDs of two paths are compared only if they come from the same neighboring AS

as-path-ignore

Syntax

as-path-ignore [ipv4] [ipv6]

no as-path-ignore

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

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 the 7705 SAR 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

dynamic-neighbor**Syntax**

dynamic-neighbor

Context

config>service>vprn>bgp>group

Description

This command enables the context to configure dynamic BGP sessions for a peer group.

prefix**Syntax**

[no] prefix *ip-prefix/ip-prefix-length*

Context

config>service>vprn>bgp>group>dynamic-neighbor

Description

This command configures a prefix to accept dynamic BGP sessions, which are sessions from source IP addresses that do not match any configured (static) neighbor addresses. A dynamic session is associated with the group having the longest-match prefix entry for the source IP address of the peer. There is no limit on the number of prefixes that can be configured. The group association determines local parameters that apply to the session, including the local AS, local IP address, MP-BGP families, and import and export policies.

The **no** form of this command removes a prefix entry.

Default

none

Parameters

ip-prefix/ip-prefix-length

specifies a prefix from which to accept dynamic BGP sessions

dynamic-neighbor-limit

Syntax

dynamic-neighbor-limit *peers*
no dynamic-neighbor-limit

Context

config>service>vprn>bgp
config>service>vprn>bgp>group

Description

This command configures the maximum number of dynamic BGP sessions that will be accepted from remote peers associated with the global BGP instance or a specific peer group. If accepting a new dynamic session would cause either the group limit or the global limit to be exceeded, the new session attempt is rejected and a notification message is sent back to the remote peer.

The **no** form of this command removes the limit on the number of dynamic sessions.

Default

no dynamic-neighbor-limit

Parameters

peers

specifies the maximum number of dynamic BGP sessions

Values 1 to 8192

enable-bgp-vpn-backup

Syntax

[no] enable-bgp-vpn-backup [ipv4] [ipv6]

Context

config>service>vprn>bgp

Description

This command allows BGP-VPN routes imported into the VPRN to be used as backup paths for IPv4 or IPv6 BGP learned prefixes.

Parameters

ipv4

allow BGP-VPN routes to be used as backup paths for IPv4 prefixes

ipv6

allow BGP-VPN routes to be used as backup paths for IPv6 prefixes

enable-peer-tracking**Syntax**

[no] enable-peer-tracking

Context

config>service>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>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

error-handling**Syntax**

error-handling

Context

config>service>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Description

This command enables the context to configure BGP error handling.

legacy-mode**Syntax**

[no] legacy-mode

Context

```
config>service>vprn>bgp>error-handling
```

Description

This command configures the legacy fault tolerance mode for BGP error handling. When enabled, configuration for fault tolerance can be enabled or disabled at the BGP global, group, or neighbor level and applied to sessions at that level with the **update-fault-tolerance** command. When disabled, **update-fault-tolerance** configurations are ignored and updated fault protection is automatically applied to all BGP sessions.

Default

no legacy-mode

update-fault-tolerance

Syntax

```
[no] update-fault-tolerance
```

Context

```
config>service>vprn>bgp>error-handling
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
```

Description

This command enables updated fault tolerance for handling a wide range of BGP Update message errors. When enabled, the system uses the 'treat-as-withdraw' and other similarly non-disruptive error handling as described in RFC 7606 as long as there are no length errors that prevent all of the NLRI fields from being correctly identified and parsed. If the **legacy-mode** command is disabled, the **update-fault-tolerance** configuration is ignored and updated fault tolerance is automatically applied to all BGP sessions.

Default

no update-fault-tolerance

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. See the section on "Route Policy" in the 7705 SAR 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 (such as #, \$, or spaces), 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*

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 (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

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 IPv4 or IPv6 address of the BGP peer router

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 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. See the section on "Route Policy" in the 7705 SAR 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 (such as #, \$, or spaces), 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 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 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

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-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-resolution

Syntax

next-hop-resolution

Context

config>service>vprn>bgp

Description

This command enters the context to configure next-hop resolution parameters.

policy

Syntax

policy *policy-name*

no policy

Context

config>service>vprn>bgp>next-hop-res

Description

This command specifies the name of a policy statement to use with the BGP next-hop resolution process. The policy controls which IP routes in the RTM are eligible to resolve the BGP next-hop addresses of IPv4 and IPv6 routes. The policy has no effect on the resolution of BGP next hops to MPLS tunnels. If a BGP next hop of an IPv4 or IPv6 route is resolved in the RTM and the longest matching route for the next-hop address is an IP route that is rejected by the policy, the route is unresolved; if the route is accepted by the policy, it becomes the resolving route.

If the **no** form of the command is used, the default next-hop-resolution policy is to use the longest matching active route in the RTM that is not a BGP route or an aggregate route.

Default

no policy

Parameters

policy-name

specifies an existing route policy name. Route policies are configured in the **config>router>policy-options** context.

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 because 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

peer-tracking-policy

Syntax

peer-tracking-policy *policy-name*

no peer-tracking-policy

Context

config>service>vpn>bgp

Description

This command specifies the name of a policy statement to use with the BGP peer-tracking function on the BGP sessions where the **peer-tracking-policy** command is enabled. The policy controls which IP routes in the RTM are eligible to indicate reachability of IPv4 and IPv6 BGP neighbor addresses. If the longest matching route in the RTM for a BGP neighbor address is an IP route that is rejected by the policy or a BGP route accepted by the policy, or if there is no matching route, the neighbor is considered unreachable and BGP tears down the peering session and holds it in the idle state until a valid route is once again available and accepted by the policy.

The **no** form of the command defaults to using the longest matching active route in the RTM that is not an aggregate route.

Default

no peer-tracking-policy

Parameters

policy-name

specifies an existing route policy name. Route policies are configured in the **config>router>policy-options** context.

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 peers.

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 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 *family limit* [**threshold** *percentage*] [**idle-timeout** {*minutes* | **forever**} | **log-only**] [**post-import**]

no prefix-limit *family*

Context

config>service>vprn>bgp>group

```
config>service>vpn>bgp>group>neighbor
```

Description

This command configures the maximum number of BGP routes that can be received from a peer before administrative action is taken. The administrative action can be the generation of a log event or the taking down of the session. If a session is taken down, it can be brought back up automatically after an idle-timeout period or it can be configured to stay down (**forever**) until the operator performs a reset.

The **prefix-limit** command allows each address family to have its own limit; a set of address family limits can be applied to one neighbor or to all neighbors in a group.

The **no** form of the command removes the **prefix-limit**.

Default

No prefix limits for any address family

Parameters

family

specifies the address family to which the limit applies

Values ipv4, vpn-ipv4, ipv6, vpn-ipv6, mvpn-ipv4, route-target, label-ipv4

limit

specifies the number of routes that can be learned from a peer, expressed as a decimal integer

Values 1 to 4294967295

percentage

specifies the threshold value, as a percentage, that triggers a warning message to be sent

Values 1 to 100

minutes

specifies the length of time, in minutes, before automatically re-establishing a session

Values 1 to 1024

forever

specifies that the session is re-established only after the **clear router bgp** command is executed

log-only

enables a warning message to be sent at the specified threshold percentage and also when the limit is reached. However, the BGP session is not taken down.

post-import

specifies that the limit should be applied only to the number of routes that are accepted by import policies

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 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

rib-management

Syntax

rib-management

Context

config>service>vprn>bgp

Description

This command enables the context to configure RIB management parameters. Under the RIB management context are options for **ipv4** and **ipv6**.

route-table-import

Syntax

route-table-import *policy-name*

no route-table-import

Context

config>service>vprn>bgp>rib-management>ipv4

config>service>vprn>bgp>rib-management>ipv6

Description

This command specifies the name of a route policy to control the importation of active routes from the IP route table into one of the BGP RIBs.

If the **route-table-import** command is not configured, or if the command refers to an empty policy, all non-BGP routes from the IP route table are imported into the applicable RIB.

If the **route-table-import** command is configured, routes that are dropped or rejected by the configured policy are not installed in the associated RIB. Rejected routes cannot be advertised to BGP peers associated with the RIB, but they can still be used to resolve BGP next hops of routes in that RIB. If the active route for a prefix is rejected by the **route-table-import** policy, then the best BGP route for that prefix in the BGP RIB can be advertised to peers as though it is used.

Aggregate routes are always imported into the applicable RIB, independent of the **route-table-import** policy.

Route modifications specified in the actions of a **route-table-import** policy are ignored and have no effect on the imported routes.

Default

no route-table-import

Parameters

policy-name

specifies the name of a policy-statement; the policy statement must already have been created

split-horizon

Syntax

[no] split-horizon

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Description

This command enables the use of split-horizon. When applied globally, to a group, or to a specific peer, split-horizon prevents routes from being reflected back to a peer that sends the best route. It applies to routes of all address families and to any type of sending peer: confed-EBGP, EBGP, or IBGP.

The configuration default is **no split-horizon**, meaning that no effort is taken to prevent a best route from being reflected back to the sending peer.



Caution: Use of the **split-horizon** command may have a detrimental impact on peer and route scaling; therefore, operators are encouraged to use it only when absolutely needed.

The **no** form of the command disables split-horizon, which allows the lower level to inherit the setting from an upper level.

Default

no split-horizon

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

7.8.2.1.4 OSPF and OSPFv3 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

```
ospf3
```

Syntax

[no] ospf3

Context

config>service>vprn

Description

This command enables access to the context to define OSPFv3 parameters for VPRN.

When an OSPFv3 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 OSPFv3 protocol instance and removes all associated configuration parameters.

Default

no ospf3

area

Syntax

[no] **area** *area-id*

Context

config>service>vprn>ospf

config>service>vprn>ospf3

Description

This command enables the context to configure an OSPF or OSPFv3 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 or OSPFv3 configuration of all the interfaces, virtual links, sham 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 or OSPFv3 areas are defined

Parameters

area-id

the OSPF or OSPFv3 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*

area-range *ipv6-prefix/prefix-length* [**advertise** | **not-advertise**]

no area-range *ipv6-prefix/prefix-length*

Context

config>service>vprn>ospf>area

config>service>vprn>ospf3>area

config>service>vprn>ospf>area>nssa

config>service>vprn>ospf3>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 *ip-prefix/mask* parameter applies in the **ospf** context. The *ipv6-prefix/prefix-length* parameter applies in the **ospf3** context.

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/mask

the IP prefix and subnet mask length for the range

ipv6-prefix/prefix-length

the IPv6 prefix and prefix length for the range

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

config>service>vprn>ospf3>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

interface *ip-int-name* [**secondary**]

no interface *ip-int-name*

Context

config>service>vprn>ospf>area

config>service>vprn>ospf3>area

Description

This command creates a context to configure an OSPF or OSPFv3 interface.

By default, interfaces are not activated in any interior gateway protocol, such as OSPF or OSPFv3, unless explicitly configured.

The **no** form of the command deletes the OSPF or OSPFv3 interface configuration for this interface. The **shutdown** command in the **config>router>ospf>interface** context or **config>router>ospf3>interface** context 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 (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

If the IP interface exists in a different area, the configuration will be rejected with an error message unless the keyword **secondary** is specified.

secondary

enables multiple secondary adjacencies to be established over this IP interface (see the 7705 SAR Routing Protocols Guide, "Multi-area Adjacencies", for information about this feature)

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.

This command is not supported in the **ospf3** context.

The **no** form of the command disables advertising point-to-point interfaces as subnet routes, meaning they are advertised as host routes.

Default

advertise-subnet

auth-keychain

Syntax

auth-keychain *name*

no auth-keychain

Context

config>service>vprn>ospf>area>interface

config>service>vprn>ospf>area>sham-link

config>service>vprn>ospf>area>virtual-link

Description

This command associates an authentication keychain with the OSPF interface, virtual link, or sham link. The keychain is a collection of keys used to authenticate OSPF messages from remote peers. The keychain allows the rollover of authentication keys during the lifetime of a session and also supports stronger authentication algorithms than clear text and MD5.

The keychain must already be defined in the **config>system>security>keychain** context.

Either the **authentication-key** command or the **auth-keychain** command can be used by OSPF, but both cannot be supported at the same time. If both commands are configured, the **auth-keychain** configuration will be applied and the **authentication-key** command will be ignored.

This command is not supported in the **ospf3** context.

The **no** form of the command removes the authentication keychain name from the configuration.

Default

no auth-keychain

Parameters

name

the name of an existing keychain, up to 32 characters

authentication

Syntax

authentication bidirectional *sa-name*

authentication inbound *sa-name* **outbound** *sa-name*

no authentication

Context

config>service>vprn>ospf3>area>interface

config>service>vprn>ospf3>area>virtual-link

Description

This command configures an interface with a static security association (SA) used to authenticate OSPFv3 packets.

This command is not supported in the **ospf** context.

The **no** form of the command removes the SA name from the configuration.

Parameters

bidirectional *sa-name*

specifies the IPsec SA name used for transmitting and receiving OSPFv3 packets

inbound *sa-name*

specifies the IPsec SA name used for receiving OSPFv3 packets

outbound *sa-name*

specifies the IPsec SA name used for transmitting OSPFv3 packets

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>sham-link

config>service>vprn>ospf>area>virtual-link

Description

This command configures the password used by the OSPF interface, virtual link, or sham 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.

Either the **authentication-key** command or the **auth-keychain** command can be used by OSPF, but both cannot be supported at the same time. If both commands are configured, the **auth-keychain** configuration will be applied and the **authentication-key** command will be ignored.

This command is not supported in the **ospf3** context.

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>vpn>ospf>area>interface

config>service>vpn>ospf>area>sham-link

config>service>vpn>ospf>area>virtual-link

Description

This command enables authentication and specifies the type of authentication to be used on the OSPF interface, virtual link, or sham link.

Both simple **password** and **message-digest** authentication are supported.

By default, authentication is not enabled on an interface or link.

This command is not supported in the **ospf3** context.

The **no** form of the command disables authentication on the interface or link.

Default

no authentication-type

Parameters

password

enables simple password (plaintext) 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.

bfd-enable

Syntax

bfd-enable [**remain-down-on-failure**]

no bfd-enable

Context

config>service>vprn>ospf>area>interface

config>service>vprn>ospf3>area>interface

Description

This command enables the use of bidirectional forwarding detection (BFD) to control the state of the associated OSPF or OSPFv3 interface. By enabling BFD on an OSPF or OSPFv3 interface, the state of the 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 using the **bfd** command under the IP interface.

If the BFD session does not come back up within 10 s and the **remain-down-on-failure** parameter is enabled, OSPF will bring down the adjacency and wait for BFD to come up again. This behavior may cause OSPF neighbors to flap because OSPF will form the adjacency and then bring it down if the BFD session is still down. If this parameter is not configured, the OSPF adjacency will form even if the BFD session does not come back up after a failure.

The **no** form of this command removes BFD from the associated OSPF or OSPFv3 adjacency.

Default

no bfd-enable

Parameters

remain-down-on-failure

forces adjacency down on BFD failure

dead-interval

Syntax

dead-interval *seconds*

no dead-interval

Context

```
config>service>vprn>ospf>area>interface
config>service>vprn>ospf3>area>interface
config>service>vprn>ospf>area>sham-link
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf3>area>virtual-link
```

Description

This command configures the time, in seconds, that OSPF or OSPFv3 waits before declaring a neighbor router, virtual-link neighbor, or sham-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 resets the configured interval to the default value.

Default

40

Special cases

OSPF or OSPFv3 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

Sham link

if the **dead-interval** configured applies to a sham link, the interval on both endpoints of the sham 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>ospf3>area>interface
config>service>vprn>ospf>area>sham-link
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf3>area>virtual-link
```

Description

This command configures the interval between OSPF or OSPFv3 hello messages issued on the interface, virtual link, or sham link.

The hello interval, in combination with the dead interval, is used to establish and maintain the adjacency.

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 resets the configured interval to the default value.

Default

10

Special cases

OSPF or OSPFv3 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

Sham link

if the **hello-interval** configured applies to a sham link, the interval on both endpoints of the sham link must have the same hello interval

Parameters

seconds

the hello interval in seconds, expressed as a decimal integer

Values 1 to 65535

interface-type

Syntax

interface-type {**broadcast** | **point-to-point**}

no interface-type

Context

```
config>service>vprn>ospf>area>interface  
config>service>vprn>ospf3>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 or OSPFv3, 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 resets the configured interface type 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 link. 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  
config>service>vprn>ospf3>area>interface
```

Description

This command applies a route next-hop policy template to an OSPF or OSPFv3 interface.

When a route next hop policy template is applied to an interface, it is applied in all areas. However, this command can only be executed under the area in which the specified interface is primary. When the command is executed, the template is applied in that area and in all other areas where the interface is secondary. If the user attempts to execute the command under an area where the interface is secondary, the command will fail.

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 or OSPFv3 interface.

Default

no lfa-policy-map

Parameters

template-name

the name of an existing template

load-balancing-weight

Syntax

load-balancing-weight *weight*

no load-balancing-weight

Context

config>service>vprn>ospf>area>interface

Description

This command configures the load balancing weight for an OSPF PE-CE interface that is used to perform weighted ECMP for a VPRN service.

The **no** form of the command removes the configured load-balancing weight for the OSPF interface.

Default

no load-balancing-weight

Parameters

weight

specifies the load-balancing weight

Values 1 to 4294967295

loopfree-alternate-exclude

Syntax

[no] loopfree-alternate-exclude

Context

```
config>service>vprn>ospf>area
config>service>vprn>ospf3>area
config>service>vprn>ospf>area>interface
config>service>vprn>ospf3>area>interface
```

Description

This command instructs OSPF or OSPFv3 to exclude a specific interface or all interfaces participating in a specific OSPF or OSPFv3 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 calculation, it is excluded in all areas. However, this command can only be executed under the area in which the specified interface is primary. When the command is executed, the interface is excluded in that area and in all other areas where the interface is secondary. If the user attempts to execute the command under an area where the interface is secondary, the command will fail.

Default

no loopfree-alternate-exclude

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>sham-link
config>service>vprn>ospf>area>virtual-link
```

Description

This command configures a message digest key when MD5 authentication is enabled on the interface, virtual link, or sham link. Multiple message digest keys can be configured.

This command is not supported in the **ospf3** context.

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

config>service>vprn>ospf3>area>interface

config>service>vprn>ospf>area>sham-link

Description

This command configures an explicit route cost metric for the interface or sham link that overrides the metrics calculated based on the speed of the underlying link.

The **no** form of the command deletes the manually configured metric, so the interface or sham link 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 or sham link, expressed as a decimal integer

Values 1 to 65535

mtu

Syntax

mtu *bytes*

no mtu

Context

config>service>vprn>ospf>area>interface

config>service>vprn>ospf3>area>interface

Description

This command configures the OSPF or OSPFv3 interface MTU value used when negotiating an OSPF or OSPFv3 adjacency.

The operational OSPF MTU value is calculated as follows.

If this command is not configured:

- the OSPF or OSPFv3 interface operational MTU derives the MTU value from the IP interface MTU (which is derived from the port MTU); for example, port MTU minus 14 bytes for a null-encapsulated Ethernet port
 - for OSPF (not OSPFv3), if the derived MTU value is less than 576 bytes, the OSPF interface operational MTU is set to 576 bytes. If a lower interface MTU is required, you must explicitly configure it using this command.

If this command is configured:

- for OSPF (not OSPFv3):
 - if the OSPF interface MTU is less than 576 bytes, it becomes the operational OSPF MTU, regardless of the port MTU value

- if the OSPF interface MTU is equal to or greater than 576 bytes, and the derived interface MTU is less than 576 bytes, the operational OSPF MTU is set to 576 bytes
- if the OSPF interface MTU is equal to or greater than 576 bytes, and the derived interface MTU is greater than 576 bytes, the operational OSPF MTU is set to the lesser of the values configured with this command and the derived MTU

The port MTU must be set to 512 bytes or higher, since OSPF cannot support port MTU values lower than 512 bytes.

- for OSPFv3:
 - the operational OSPF MTU is set to the lesser of the values configured with this command and the derived MTU
 - this applies only when the port MTU is set to 1280 bytes or higher, since OSPFv3 cannot support port MTU values less than 1280 bytes

To determine the actual packet size, add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured with this command.

If the OSPF **mtu** command is configured to a value less than the interface or port MTU value, the OSPF MTU value will be used to transmit OSPF packets.

Use the **no** form of this command to revert to the default.

Default

no mtu – uses the value derived from the port MTU

Parameters

bytes

the MTU to be used by OSPF or OSPFv3 for this logical interface in bytes

Values OSPF: 512 to 9710 (9724 – 14) (depends on the physical media)
 OSPFv3: 1280 to 9710 (9724 – 14) (depends on the physical media)

passive

Syntax

[no] passive

Context

config>service>vprn>ospf>area>interface

config>service>vprn>ospf3>area>interface

Description

This command adds the passive property to an OSPF or OSPFv3 interface.

By default, only interface addresses that are configured for OSPF or OSPFv3 will be advertised as OSPF or OSPFv3 interfaces. The **passive** parameter allows an interface to be advertised as an OSPF or OSPFv3 interface without running the OSPF or OSPFv3 protocol.

While in passive mode, the interface will ignore ingress OSPF or OSPFv3 protocol packets and will not transmit any OSPF or OSPFv3 protocol packets.

The **no** form of the command removes the passive property from the OSPF or OSPFv3 interface.

Default

no passive

priority

Syntax

priority *number*

no priority

Context

config>service>vprn>ospf>area>interface

config>service>vprn>ospf3>area>interface

Description

This command configures the priority of the OSPF or OSPFv3 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 resets 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>ospf3>area>interface
config>service>vprn>ospf>area>sham-link
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf3>area>virtual-link
```

Description

This command specifies the length of time, in seconds, that OSPF or OSPFv3 will wait before retransmitting an unacknowledged LSA to an OSPF or OSPFv3 neighbor.

The value should be greater 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 resets the configuration 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>ospf3>area>interface
config>service>vprn>ospf>area>sham-link
config>service>vprn>ospf>area>virtual-link
config>service>vprn>ospf3>area>virtual-link
```

Description

This command configures the estimated time, in seconds, that it takes to transmit an LSA on the interface, virtual link, or sham link.

The **no** form of this command resets the configuration to the default delay time.

Default

1

Parameters*seconds*

the transit delay in seconds, expressed as a decimal integer

Values 1 to 1800**key-rollover-interval****Syntax****key-rollover-interval** *key-rollover-interval***no key-rollover-interval****Context**

config>service>vprn>ospf3>area

Description

This command configures the key rollover interval. The **no** form of the command resets the configured interval to the default setting.

Default

10

Parameters*key-rollover-interval*

specifies the time, in seconds, after which a key rollover will start

Values 10 to 300**nssa****Syntax****[no] nssa****Context**

config>service>vprn>ospf>area

config>service>vprn>ospf3>area

Description

This command enables the context to configure an OSPF or OSPFv3 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 or OSPFv3 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 or OSPFv3 domain.

Existing virtual links of a 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] [adjacency-check]

originate-default-route [type-nssa] [adjacency-check]

no originate-default-route

Context

config>service>vprn>ospf>area>nssa

config>service>vprn>ospf3>area>nssa

Description

This command enables the generation of a default route and its LSA type into an NSSA by an NSSA ABR or ASBR.

The functionality of the **type-7** parameter and the **type-nssa** parameter is the same. The **type-7** parameter is available in the **ospf** context; the **type-nssa** parameter is available in the **ospf3** context. Include the **type-7** or **type-nssa** parameter to inject a type 7 LSA default route instead of a type 3 LSA into the NSSA configured with no summaries.

To return to a type 3 LSA, enter the **originate-default-route** command without the **type-7** or **type-nssa** 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 | type-nssa

specifies that a type 7 LSA should be used for the default route

Default type 3 LSA for the default route

adjacency-check

specifies whether adjacency checks are performed before originating a default route. If this parameter is configured, an area 0 adjacency is required for the ABR to advertise the default route.

redistribute-external

Syntax

[no] redistribute-external

Context

config>service>vprn>ospf>area>nssa

config>service>vprn>ospf3>area>nssa

Description

This command enables the redistribution of external routes into the NSSA on an NSSA 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 or OSPFv3 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 or OSPFv3 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>ospf3>area>nssa

config>service>vprn>ospf>area>stub

```
config>service>vprn>ospf3>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

sham-link

Syntax

```
[no] sham-link {ip-int-name ip-address}
```

Context

```
config>service>vprn>ospf>area
```

Description

This command configures an OSPF area sham link to a far-end PE OSPF router.

The **no** form of the command removes the sham link.

Default

no sham-link

Parameters

ip-int-name

specifies the local interface name used for the sham link. This is a mandatory parameter. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within double quotes. If the IP interface name does not exist or does not have an IP address configured, an error message will be returned.

ip-address

specifies the IP address of the sham-link remote neighbor. This is a mandatory parameter. The address must be a valid IP address.

stub

Syntax

[no] **stub**

Context

config>service>vpn>ospf>area

config>service>vpn>ospf3>area

Description

This command enables access to the context to configure an OSPF or OSPFv3 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 stub area or NSSA are removed when its designation is changed to NSSA or stub.

An OSPF or OSPFv3 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>vpn>ospf>area>stub

config>service>vpn>ospf3>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 resets the configuration 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

config>service>vprn>ospf3>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

config>service>vprn>ospf3

Description

This command specifies export route policies to determine which routes are exported from the routing table manager to OSPF or OSPFv3. Export policies are only in effect if OSPF or OSPv3 is configured as an ASBR.

If no export policy is specified, routes that are not OSPF or OSPFv3 are not exported from the routing table manager to OSPF or OSPFv3.

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.

See the 7705 SAR Router Configuration Guide for information about defining route policies.

Default

no export – no export route policies specified

Parameters

policy-name

the name of an existing route policy

external-db-overflow

Syntax

external-db-overflow *limit seconds*

no external-db-overflow

Context

config>service>vprn>ospf

config>service>vprn>ospf3

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 or OSPFv3 area. OSPF or OSPFv3 stub areas and 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

config>service>vprn>ospf3

Description

This command configures the preference for OSPF or OSPFv3 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 the following table.

Table 138: Route preference defaults by route type

Route type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF/OSPFv3 internal	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
OSPF/OSPFv3 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. See the 7705 SAR Router Configuration Guide for information about ECMP.



Note: To configure a preference for static routes, use the **config>router>static-route-entry** command. See the 7705 SAR Router Configuration Guide for information.

The **no** form of the command returns the setting to the default value.

Default

external-preference 150 – OSPF or OSPFv3 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

config>service>vprn>ospf3

Description

This command specifies whether to ignore the DN (down) bit for OSPF or OSPFv3 LSA packets for this instance of OSPF or OSPFv3 on the router. When enabled, the DN bit for OSPF or OSPFv3 LSA packets will be ignored. When disabled, the DN bit will not be ignored for OSPF or OSPFv3 LSA packets.

Default

no ignore-dn-bit

import

Syntax

import *policy-name* [*policy-name...*(up to 5 max)]

no import

Context

config>service>vprn>ospf

Description

This command configures up to five import route policies that determine which routes are imported into the routing table.

When a prefix received in an OSPF LSA is accepted by an entry in an OSPF import policy, it is installed in the routing table if it is the most preferred route to the destination. When a prefix received in an OSPF LSA is rejected by an entry in an OSPF import policy, it is not installed in the routing table, even if it has the lowest preference value among all the routes to that destination.

The flooding of LSAs is not affected by OSPF import policy actions.

The **no** form of this command removes all import policies from the configuration. The default behavior then applies, that is, if an OSPF route has the lowest preference value among all routes to the destination, it is installed in the routing table.

Default

no import

Parameters

policy-name

specifies the import route policy name. The route policy names must already be defined.

loopfree-alternates

Syntax

[no] loopfree-alternates

Context

config>service>vprn>ospf

config>service>vprn>ospf3

Description

This command enables loop-free alternate (LFA) computation by SPF under the OSPFv2 or OSPFv3 routing protocol context.

When this command is enabled, the OSPF or OSPFv3 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 alternates

exclude

Syntax

exclude

Context

config>service>vprn>ospf>loopfree-alternates

config>service>vprn>ospf3>loopfree-alternates

Description

This command enables the context for identifying prefix policies to be excluded from the LFA calculation by OSPF.

prefix-policy

Syntax

prefix-policy *prefix-policy* [*prefix-policy*...(up to 5 max)]

no prefix-policy

Context

```
config>service>vprn>ospf>loopfree-alternates>exclude
```

```
config>service>vprn>ospf3>loopfree-alternates>exclude
```

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 **config>router>policy-options>prefix-list** command. For information about prefix lists, see the 7705 SAR Router Configuration Guide, "Route Policies".

The default action of the **loopfree-alternates>exclude>prefix-policy** 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 this command removes the excluded prefix policy.

Default

no prefix-policy

Parameters

prefix-policy

the name of the prefix policy to be excluded from the LFA SPF calculation for OSPF. Up to five prefixes can be specified. The specified prefix policy must already be defined.

overload

Syntax

overload [timeout seconds]

no overload

Context

```
config>service>vprn>ospf
```

```
config>service>vprn>ospf3
```

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 or OSPFv3 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

config>service>vprn>ospf3

Description

This command is used to determine if the OSPF or OSPFv3 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

```
config>service>vprn>ospf3
```

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 or OSPFv3 upon boot-up 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
```

```
config>service>vprn>ospf3
```

Description

This command configures the preference for OSPF or OSPFv3 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 138: Route preference defaults by route type](#). 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. See the 7705 SAR Router Configuration Guide for information about ECMP.

The **no** form of the command resets the preference configuration to the default value.

Default

preference 10 – OSPF or OSPFv3 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

config>service>vprn>ospf3

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** command in the **config>router>ospf>area>interface** *ip-int-name* context or **config>router>ospf3>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

config>service>vprn>ospf3

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 or OSPFv3, 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 resets the router ID 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

config>service>vprn>ospf3

Description

This command specifies whether to suppress the setting of the DN (down) bit for OSPF or OSPFv3 LSA packets generated by this instance of OSPF or OSPFv3 on the router. When enabled, the DN bit will not

be set. When disabled, this instance of the OSPF or OSPFv3 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

config>service>vprn>ospf3

Description

This command enables the context that allows for the configuration of OSPF or OSPFv3 timers. Timers control the delay between receipt of an 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

config>service>vprn>ospf3>timers

Description

This command defines the minimum delay that must pass between receipt of the same 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 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

config>service>vprn>ospf3>timers

Description

This command customizes the throttling of OSPF or OSPFv3 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

config>service>vprn>ospf3>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>vpn>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.

This command is not supported in the **ospf3** context.

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>vpn>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.

This command is not supported in the **ospf3** context.

Default

vpn-tag 0

Parameters

vpn-tag

specifies the route tag for an OSPF VPN

Values 0 to 4294967295

7.8.2.1.5 IGMP commands

igmp

Syntax

[no] igmp

Context

config>service>vpn

Description

This command enables the context to configure IGMP parameters.

The **no** form of the command disables IGMP.

Default

disabled

interface**Syntax**

[no] interface *ip-int-name*

Context

config>service>vprn>igmp

Description

This command enables the context to configure IGMP interface parameters.

Parameters

ip-int-name

specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed between double quotes.

disable-router-alert-check**Syntax**

[no] disable-router-alert-check

Context

config>service>vprn>igmp>if

Description

This command enables or disables the IGMP router alert check option.

The **no** form of the command enables the router alert check.

Default

no disable-router-alert-check

import**Syntax**

import *policy-name*

no import

Context

```
config>service>vprn>igmp>if
```

Description

This command imports a policy to filter IGMP packets on this interface.

The **no** form of the command removes the policy association from the IGMP instance.

Default

no import

Parameters

policy-name

the import route policy name. The specified names must already be defined.

max-groups

Syntax

max-groups *value*

no max-groups

Context

```
config>service>vprn>igmp>if
```

Description

This command specifies the maximum number of groups for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed.

The **no** form of the command removes the value.

Default

0 – no limit to the number of groups

Parameters

value

specifies the maximum number of groups for this interface

Values 1 to 2047

max-grp-sources

Syntax

max-grp-sources *max-grp-sources*

no max-grp-sources

Context

config>service>vprn>igmp>if

Description

This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than the currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed.

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

Default

0

Parameters

max-grp-sources

specifies the maximum number of group sources

Values 1 to 2407

ssm-translate

Syntax

ssm-translate

Context

config>service>vprn>igmp

config>service>vprn>igmp>if

Description

This command enables the context to configure group ranges that are translated to SSM (S,G) entries. If the static entry needs to be created, it must be translated from an IGMPv1 or IGMPv2 request to a Source Specific Multicast (SSM) join request. An SSM translate source can only be added when the **starg** command is not enabled. An error message is generated when trying to configure the **source** command while the **starg** command is enabled.

grp-range

Syntax

[no] **grp-range** *start end*

Context

config>service>vprn>igmp>ssm-translate

config>service>vprn>igmp>if>ssm-translate

Description

This command is used to configure group ranges that are translated to SSM (S,G) entries.

Parameters

start

specifies the start address of the multicast group range

end

specifies the end address of the multicast group range. This value should always be greater than or equal to the *start* value.

source

Syntax

[no] **source** *ip-address*

Context

config>service>vprn>igmp>ssm-translate>grp-range

config>service>vprn>igmp>if>ssm-translate>grp-range

Description

This command specifies the source IP address for the group range. Whenever a (*,G) report is received and is in the range specified by **grp-range** *start* and *end* parameters, it is translated to an (S,G) report where *ip-address* is the source address.

Parameters

ip-address

specifies the unicast source address

static

Syntax

static

Context

config>service>vprn>igmp>if

Description

This command accesses the context to test forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without having received membership reports from host members.

Default

n/a

group

Syntax

[no] **group** *grp-ip-address*

Context

config>service>vprn>igmp>if>static

Description

This command adds a static multicast group as either a (*,G) record or one or more (S,G) records. Use IGMP static group memberships to test multicast forwarding when there is no receiver host. When IGMP static groups are enabled, data is forwarded to an interface without receiving membership reports from host members.

When static IGMP group entries on point-to-point links that connect routers to a rendezvous point (RP) are configured, the static IGMP group entries do not generate join messages toward the RP.

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. The address must be in dotted-decimal notation.

source

Syntax

[no] **source** *ip-address*

Context

config>service>vprn>igmp>if>static>group

Description

This command specifies the source address of the multicast group. It is an IPv4 unicast address. By specifying the source address, a multicast receiver host signals to the router that the multicast group will only be receiving multicast traffic from this specific source.

The **source** command and the specification of individual sources for the same group are mutually exclusive.

The **source** command, in combination with the **group** command, is used to create a specific (S,G) static group entry.

The **no** form of the command removes the source from the configuration.

Default

n/a

Parameters

ip-address

specifies the IPv4 unicast address

subnet-check

Syntax

[no] **subnet-check**

Context

config>service>vprn>igmp>if

Description

This command enables or disables subnet checking for IGMP messages received on this interface. All IGMP packets with a source address that is not in the local subnet are dropped.

The **no** form of the command disables **subnet-check**.

Default

enabled

version

Syntax

version *version*

no version

Context

config>service>vprn>igmp>if

Description

This command specifies the IGMP version. If routers run different versions of IGMP, they negotiate the lowest common version of IGMP that is supported by hosts on their subnet and operate in that version. For IGMP to function correctly, all routers on a LAN should be configured to run the same version of IGMP.

For IGMPv3, a multicast router that is also a group member performs both parts of IGMPv3, receiving and responding to its own IGMP message transmissions as well as those of its neighbors.

Default

3

Parameters

version

specifies the IGMP version number

Values 1, 2, or 3

query-interval

Syntax

query-interval *seconds*

no query-interval

Context

config>service>vprn>igmp

Description

This command specifies the frequency at which the querier router transmits general host-query messages. The host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1.

Default

125

Parameters

seconds

specifies the frequency, in seconds, at which the router transmits general host-query messages

Values 2 to 1024

query-last-member-interval

Syntax

query-last-member-interval *seconds*

no query-last-member-interval

Context

config>service>vprn>igmp

Description

This command configures the frequency at which the querier sends group-specific query messages, including messages sent in response to leave-group messages; the shorter the interval, the faster the detection of the loss of the last member of a group.

Default

1

Parameters

seconds

specifies the frequency, in seconds, at which query messages are sent

Values 1 to 1023

query-response-interval

Syntax

query-response-interval *seconds*

no query-response-interval

Context

config>service>vprn>igmp

Description

This command specifies how long the querier router waits to receive a response to a host-query message from a host.

Default

10

Parameters*seconds*

specifies the length of time, in seconds, that the router waits to receive a response to the host-query message from the host

Values 1 to 1023**robust-count****Syntax****robust-count** *robust-count***no robust-count****Context**

config>service>vprn>igmp

Description

This command configures the robust count. The *robust-count* allows adjusting for the expected packet loss on a subnet. If a subnet anticipates losses, the *robust-count* can be increased.

Default

2

Parameters*robust-count*

specifies the robust count value

Values 2 to 10**7.8.2.1.6 PIM commands****pim****Syntax****[no] pim****Context**

config>service>vprn

Description

This command configures a Protocol Independent Multicast (PIM) instance in the VPRN service. When a PIM instance is created, the protocol is enabled. PIM is used for multicast routing within the network. Devices in the network can receive the requested multicast feed and non-participating routers can be pruned. The router supports PIM sparse mode (PIM-SM).

The **no** form of the command deletes the PIM protocol instance, removing all associated configuration parameters.

Default

n/a

apply-to

Syntax

apply-to {all | none}

Context

config>service>vprn>pim

Description

This command creates a PIM interface configured with default parameters.

If a manually created interface or a modified interface is deleted, the interface will be recreated when the **apply-to** command is executed. If PIM is not required on a specific interface, execute a **shutdown** command on the interface.

The **apply-to** command is saved first in the PIM configuration structure; all subsequent **apply-to** commands either create new structures or modify the defaults as created by the **apply-to** command.

Default

none (keyword)

Parameters

all

specifies that all VPRN and non-VPRN interfaces are automatically applied in PIM

none

specifies that no interfaces are automatically applied in PIM; PIM interfaces must be manually configured

import

Syntax

import {join-policy | register-policy} *policy-name* [*policy-name*...(up to 5 max)]

no import {join-policy | register-policy}

Context

config>service>vpn>pim

Description

This command specifies up to five import route policies to be used for determining which routes are accepted from peers. Route policies are configured in the **config>router>policy-options** context. When an import policy is not specified, BGP routes are accepted by default.

The **no** form of the command removes the policy associations from the PIM instance for the specified type.

Default

no import join-policy no import register-policy

Parameters

join-policy

specifies filtering PIM join messages to prevent unwanted multicast streams from traversing the network

register-policy

specifies filtering PIM messages to prevent register messages from being processed by the RP. This filter can only be defined on an RP. When a match is found, the RP immediately sends back a register-stop message.

policy-name

specifies the route policy name

interface

Syntax

[no] interface *ip-int-name*

Context

config>service>vpn>pim

Description

This command enables PIM on an interface and enables the context to configure interface-specific parameters. By default, interfaces are activated in PIM based on the [apply-to](#) command and do not have to be configured on an individual basis unless the default values must be changed.

The **no** form of the command deletes the PIM interface configuration for this interface. If the [apply-to](#) command parameter is configured, then the **no interface** form must be saved in the configuration to avoid automatic recreation of the interface after the next [apply-to](#) command is executed as part of a reboot.

The **shutdown** command can be used to disable an interface without removing the configuration for the interface.

Default

Interfaces are activated in PIM based on the **apply-to** command.

Parameters

ip-int-name

specifies the interface name up to 32 characters; if the string contains special characters (such as #, \$, or spaces), then the entire string must be enclosed between double quotes

assert-period

Syntax

assert-period *assert-period*

no assert-period

Context

config>service>vprn>pim>if

Description

This command configures the time between refreshes of PIM assert messages on an interface.

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

Default

60

Parameters

assert-period

specifies the time, in seconds, between refreshes of PIM assert messages on an interface

Values 1 to 300

bsm-check-rtr-alert

Syntax

[no] **bsm-check-rtr-alert**

Context

config>service>vprn>pim>if

Description

This command enables the checking of the router alert option in the bootstrap messages received on this interface.

Default

no bsm-check-rtr-alert

bfd-enable

Syntax

[no] **bfd-enable** [ipv4]

Context

config>service>vpn>pim>if

Description

This command enables the use of bidirectional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a 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 via the **bfd** command under the IP interface.

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

Default

no bfd-enable

hello-interval

Syntax

hello-interval *hello-interval*

no hello-interval

Context

config>service>vpn>pim>if

Description

This command configures the interval at which PIM hello messages are transmitted on this interface.

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

Default

30

Parameters

hello-interval

specifies the hello interval, in seconds; a 0 value disables the sending of hello messages

Values 0 to 255

hello-multiplier

Syntax

hello-multiplier *deci-units*
no hello-multiplier

Context

config>service>vprn>pim>if

Description

This command configures the multiplier used to determine the hold time for a PIM neighbor.

The **hello-multiplier** in conjunction with the **hello-interval** determines the hold time for a PIM neighbor.

Parameters

deci-units

specifies the value of the hello-multiplier, in multiples of 0.1, for the formula used to calculate the hello hold-time

hello hold-time = (hello-interval * hello-multiplier) / 10

This allows the PIMv2 default timeout of 3.5 s to be supported. For example, if hello-interval = 1 s, and hello-multiplier = 35 deci-units, then hold-time = (1 * 35) / 10 = 3.5 s.

Values 20 to 100

Default 35 (3.5 s)

improved-assert

Syntax

[no] improved-assert

Context

config>service>vprn>pim>if

Description

This command enables improved assert processing on this interface. The PIM assert process establishes a forwarder node for a LAN and requires interaction between the control and forwarding planes.

The assert process is started when data is received on an outgoing interface. There can be an impact on performance if data is continuously received on an outgoing interface.

When enabled, the PIM assert process occurs entirely on the control plane, with no interaction between the control and forwarding planes.

Default

enabled

instant-prune-echo**Syntax**

[no] instant-prune-echo

Context

config>service>vprn>pim>if

Description

This command enables or disables instant PruneEcho for a PIM interface.

Default

no instant-prune-echo

ipv4-multicast-disable**Syntax**

[no] ipv4-multicast-disable

Context

config>service>vprn>pim

config>service>vprn>pim>if

Description

This command administratively disables or enables PIM operation for IPv4.

Default

no ipv4-multicast-disable

max-groups**Syntax**

max-groups *value*

no max-groups

Context

config>service>vprn>pim>if

Description

This command configures the maximum number of groups for which PIM can have a downstream state based on received PIM join messages on this interface. This number does not include IGMP local receivers on the interface. When this configuration is changed dynamically to a value lower than the currently accepted number of groups, the groups that are already accepted are not deleted. Only new groups will not be allowed. When this configuration has a value of 0, there is no limit to the number of groups.

Parameters

value

specifies the maximum number of groups for this interface

Values 1 to 2047

multicast-senders

Syntax

multicast-senders {**auto** | **always** | **never**}

no multicast-senders

Context

config>service>vprn>pim>if

Description

This command configures the way subnet matching is done for incoming data packets on this interface. An IP multicast sender is a user entity to be authenticated in a receiving host.

Parameters

auto

specifies that subnet matching is automatically performed for incoming data packets on this interface

always

specifies that subnet matching is always performed for incoming data packets on this interface

never

specifies that subnet matching is never performed for incoming data packets on this interface

multicast-to-multicast

Syntax

multicast-to-multicast source *ip-address* **group-start** *ip-address* **group-end** *ip-address* **to-multicast** *group-address*

no multicast-to-multicast

Context

```
config>service>vprn>pim>interface
```

Description

This command enables multicast-to-multicast address translation by mapping a range of source IP addresses to a range of multicast group addresses. The PIM interface on the 7705 SAR translator router is the source of the multicast address. Multiple (S,G)s (for example, s1,g1 / s2,g1 / s3,g1) can be mapped to a single PIM interface, using the same source IP address for the translated (S,G) but for a range of groups.

The PIM interface on the translator router must first be enabled for multicast translation with the **config>router>interface>multicast-translation** command.

Default

no multicast-to-multicast

Parameters

source *ip-address*

the source address of the (S,G) being translated

group-start *ip-address*

the starting group address of the (S,G) being translated

group-end *ip-address*

the ending group address of the (S,G) being translated

group-address

the multicast group address used for translation

priority

Syntax

priority *dr-priority*

no priority

Context

```
config>service>vprn>pim>if
```

Description

This command sets the priority value for the interface to become the designated router (DR), so that multiple PIM routers do not exist on one Layer 2 network.

The **no** form of the command resets the priority to the default value.

Default

1 (the router is least likely to become the designated router)

Parameters

dr-priority

specifies the priority to become the designated router; greater values have higher priority

Values 1 to 4294967295

sticky-dr

Syntax

sticky-dr [**priority** *dr-priority*]

no sticky-dr

Context

config>service>vpn>pim>if

Description

This command enables **sticky-dr** operation on this interface. When enabled, the priority value used in PIM hello messages sent on this interface when elected as the designated router (DR) is changed to the value configured with this command. This is done to avoid forwarding delays caused by DR recovery, which occurs when switching back to the old DR on a LAN when the old DR comes back up.

By enabling **sticky-dr** on this interface, the interface continues to act as the DR for the LAN even after the old DR comes back up.

When **sticky-dr** is used without the **priority** keyword, the **sticky-dr priority** value is 1024 (default).

The **no** form of the command disables **sticky-dr** operation on this interface.

Default

disabled (no sticky-dr)

Parameters

dr-priority

when **sticky-dr** operation is enabled, *dr-priority* sets the DR priority sent in PIM hello messages after the election of that interface as the DR

Default 1024

Values 1 to 4294967295

three-way-hello

Syntax

[no] **three-way-hello**

Context

```
config>service>vprn>pim>if
```

Description

This command configures the compatibility mode to enable three-way hello. By default, three-way hello is disabled on all interfaces and the standard two-way hello is supported.

Default

no three-way-hello

tracking-support

Syntax

```
[no] tracking-support
```

Context

```
config>service>vprn>pim>if
```

Description

This command sets the T-bit in the LAN prune delay option of the hello message. This indicates the router's capability to disable join-message suppression.

Default

no tracking-support

unicast-to-multicast

Syntax

```
unicast-to-multicast unicast-start ip-address unicast-end ip-address destination ip-address to-multicast ip-address
```

```
no unicast-to-multicast
```

Context

```
config>service>vprn>pim>interface
```

Description

This command enables unicast-to-multicast address translation by mapping a range of unicast source addresses and a unicast destination address to a multicast group address. The unicast destination address is a loopback IP address configured on the 7705 SAR that is performing the translation. This translator router becomes the source of the multicast packets. The multicast source address is a loopback interface IP address configured on the PIM interface of the translator router. The PIM interface on the 7705 SAR translator router must first be enabled for multicast translation with the **config>service>vprn>interface>multicast-translation** command.

The unicast destination and the multicast source can be the same loopback address or different loopback addresses.

The translation can map a range of unicast source addresses to a range of multicast group addresses. For example, if the unicast source address range is 1.1.1.1 to 1.1.1.4 and the multicast group address is 230.0.0.100, the following multicast destination address range is created:

Unicast source	Multicast group
1.1.1.1	230.0.0.100
1.1.1.2	230.0.0.101
1.1.1.3	230.0.0.102
1.1.1.4	230.0.0.103

Default

no unicast-to-multicast

Parameters

- unicast-start** *ip-address*
the start of the range of unicast source addresses to be translated
- unicast-end** *ip-address*
the end of the range of unicast source addresses to be translated
- destination** *ip-address*
the destination address of the unicast stream being translated
- multicast** *ip-address*
the group and destination addresses for the multicast stream

non-dr-attract-traffic

Syntax

[no] non-dr-attract-traffic

Context

config>service>vprn>pim

Description

This command specifies whether the router should ignore the designated router state and attract traffic even when it is not the designated router.

An operator can configure an interface (router or IES or VPRN interfaces) with IGMP and PIM. The interface IGMP state will be synchronized to the backup node if it is associated with the redundant peer port. The interface can be configured to use PIM, which will cause multicast streams to be sent to the elected DR only. The DR will also be the router sending traffic to the DSLAM. Since it may be required to

attract traffic to both routers, a **non-dr-attract-traffic** flag can be used in the PIM context to have the router ignore the DR state and attract traffic when it is not the DR. While using this flag, the router may not send the stream to the DSLAM when it is not the DR.

The **no** form of the command disables the DR state. When disabled, the DR state is used.

Default

no non-dr-attract-traffic

rp

Syntax

rp

Context

config>service>vprn>pim

Description

This command enables access to the context to configure the rendezvous point (RP) of a PIM protocol instance.

A PIM router acting as an RP must respond to a PIM register message that specifies an SSM multicast group address by sending stop register messages to the first-hop router. The PIM router does not build an (S, G) shortest path tree toward the first-hop router. An SSM multicast group address can be an address either from the SSM default range of 232/8 or from a multicast group address range that was explicitly configured for SSM.

Default

rp enabled when PIM is enabled

anycast

Syntax

[no] **anycast** *rp-ip-address*

Context

config>service>vprn>pim>rp

Description

This command configures a PIM anycast protocol instance for the RP being configured. Anycast enables fast convergence when a PIM RP router fails by allowing receivers and sources to rendezvous at the closest RP.

The **no** form of the command removes the anycast instance from the configuration.

Default

n/a

Parameters

rp-ip-address

specifies the loopback IP address shared by all routes that form the RP set for this anycast instance. Only a single address can be configured. If a subsequent **anycast** command is entered with an address, the old address is replaced with the new address. If no *ip-address* is entered, the command is used to enter the anycast CLI level.

rp-set-peer

Syntax

[no] **rp-set-peer** *ip-address*

Context

config>service>vprn>pim>rp>anycast

Description

This command configures a peer in the anycast RP-set. The *ip-address* identifies the address used by the other node as the RP candidate address for the same multicast group address range as configured on this node.



Caution: This is a manual procedure. Caution should be taken to produce a consistent configuration of an RP-set for a given multicast group address range. The priority should be identical on each node and be a higher value than any other configured RP candidate that is not a member of this RP-set.

Although there is no set maximum number of addresses that can be configured in an RP-set, up to 15 multicast addresses is recommended.

The **no** form of the command removes an entry from the list.

Default

n/a

Parameters

ip-address

specifies the address used by the other node as the RP candidate address for the same multicast group address range as configured on this node

auto-rp-discovery

Syntax

[no] **auto-rp-discovery**

Context

```
config>service>vprn>pim>rp
```

Description

This command enables auto-RP protocol in discovery mode. In discovery mode, RP-mapping and RP-candidate messages are received and forwarded to downstream nodes. RP-mapping messages are received locally to learn about availability of RP nodes present in the network.

Either **bsr-candidate** for IPv4 or **auto-rp-discovery** can be configured; the two mechanisms cannot be enabled together. The **auto-rp-discovery** command cannot be enabled together with **mdt-type sender-only** or **mdt-type receiver-only**, or **wildcard-spmsi** configurations.

The **no** form of the command disables auto-RP discovery.

Default

no auto-rp-discovery

bootstrap-export

Syntax

bootstrap-export *policy-name* [*policy-name*... (up to 5 max)]

no bootstrap-export

Context

```
config>service>vprn>pim>rp
```

Description

This command exports policies to control the flow of bootstrap messages from the RP. Up to five policies can be defined.

The **no** form of this command removes the specified policy names from the configuration.

Default

n/a

Parameters

policy-name

specifies the policy name, where the policy statement must already be configured in the **config>router>policy-options** context

bootstrap-import

Syntax

bootstrap-import *policy-name* [*policy-name*... (up to 5 max)]

no bootstrap-import

Context

```
config>service>vprn>pim>rp
```

Description

This command imports policies to control the flow of bootstrap messages to the RP. Up to five policies can be defined.

The **no** form of this command removes the specified policy names from the configuration.

Default

n/a

Parameters

policy-name

specifies the policy name, where the policy statement must already be configured in the **config>router>policy-options** context

bsr-candidate

Syntax

bsr-candidate

Context

```
config>service>vprn>pim>rp
```

Description

This command enables the context to configure candidate bootstrap router (BSR) parameters.

Either **bsr-candidate** for IPv4 or **auto-rp-discovery** can be configured; the two mechanisms cannot be enabled together.

Default

bsr-candidate shutdown

address

Syntax

[no] **address** *ip-address*

Context

```
config>service>vprn>pim>rp>bsr-candidate
```

```
config>service>vprn>pim>rp>rp-candidate
```

Description

This command configures a static bootstrap or rendezvous point (RP) as long as the source is not directly attached to this router.

Use the **no** form of this command to remove the static RP from the configuration.

Default

no address

Parameters

ip-address

specifies the IP host address that will be used by the IP interface within the subnet. This address must be a unique unicast address within the subnet.

hash-mask-len

Syntax

hash-mask-len *hash-mask-length*

no hash-mask-len

Context

config>service>vprn>pim>rp>bsr-candidate

Description

This command is used to configure the length of the mask that is combined with the group address before the hash function is called. All groups with the same hash result will map to the same RP. For example, if the *hash-mask-length* value is 24, only the first 24 bits of the group addresses matter. This mechanism is used to map one group or multiple groups to an RP.

Default

30

Parameters

hash-mask-length

the hash mask length

Values 0 to 32

priority

Syntax

priority *bootstrap-priority*

Context

```
config>service>vprn>pim>rp>bsr-candidate
```

Description

This command defines the priority used when determining the rendezvous point (RP). The higher the priority value the more likely that this router becomes the RP. If there is a tie, the router with the highest IP address is elected.

Parameters

bootstrap-priority

the priority value used to become the bootstrap router

Values 0 to 255

Default 0 (the router is not eligible to be the bootstrap router)

rp-candidate

Syntax

rp-candidate

Context

```
config>service>vprn>pim>rp
```

Description

This command enables the context to configure the candidate rendezvous point (RP) parameters.

Default

enabled when PIM is enabled

group-range

Syntax

[no] group-range {*grp-ip-address/mask* | *grp-ip-address netmask*}

Context

```
config>service>vprn>pim>rp>rp-candidate
```

Description

This command configures the group address or range of group addresses for which this router can be the rendezvous point (RP).

Use the **no** form of this command to remove the group address or range of group addresses for which this router can be the RP from the configuration.

Default

n/a

Parameters

grp-ip-address/mask | *grp-ip-address*

specifies the multicast group IP address or the IP address and mask length

netmask

specifies the subnet mask in dotted-decimal notation

holdtime**Syntax**

holdtime *holdtime*

no holdtime *holdtime*

Context

config>service>vprn>pim>rp>rp-candidate

Description

This command defines the length of time a neighboring router considers this router to be up.

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

Default

150

Parameters

holdtime

specifies the length of time, in seconds, that neighbor should consider the sending router to be operational

Values 5 to 255

priority**Syntax**

priority *priority*

no priority *priority*

Context

config>service>vprn>pim>rp>rp-candidate

Description

This command defines the priority used to determine the rendezvous point (RP). The higher the priority value, the more likely that this router will become the RP.

Use the **no** form of this command to revert to the default value.

Default

192

Parameters

priority

specifies the priority to become the designated router

Values 0 to 255

static

Syntax

static

Context

config>service>vpn>pim>rp

Description

This command enables access to the context to configure a static rendezvous point (RP) for a PIM-SM protocol instance.

Default

n/a

address

Syntax

[no] address *ip-address*

Context

config>service>vpn>pim>rp>static

Description

This command configures the static rendezvous point (RP) address.

The **no** form of this command removes the static RP entry from the configuration.

Default

n/a

Parameters

ip-address

specifies the IP host address

group-prefix**Syntax**

[no] **group-prefix** {*grp-ip-address/mask* | *grp-ip-address netmask*}

Context

config>service>vprn>pim>rp>static

Description

The command defines a range of multicast IP addresses for which a certain RP is applicable.

The **no** form of the command removes the criterion.

Default

n/a

Parameters

grp-ip-address/mask | *grp-ip-address*

specifies the multicast group IPv4 address or the IPv4 address and mask length

netmask

the subnet mask in dotted-decimal notation

override**Syntax**

[no] **override**

Context

config>service>vprn>pim>rp>static

Description

This command changes the precedence of static RP over dynamically-learned RP.

When enabled, the static group-to-RP mappings take precedence over the dynamically-learned mappings.

Default

no override

spt-switchover-threshold**Syntax****spt-switchover-threshold** {*grp-ip-address/mask* | *grp-ip-address netmask*} *spt-threshold***no spt-switchover-threshold** {*grp-ip-address/mask* | *grp-ip-address netmask*}**Context**

config>service>vprn>pim

Description

This command configures a shortest path tree (SPT) switchover threshold for a group prefix.

PIM-SM routers with directly connected routers receive multicast traffic initially on a shared tree rooted at the rendezvous point (RP). Once the traffic arrives on the shared tree and the source of the traffic is known, a switchover to the SPT tree rooted at the source is attempted.

For a group that falls in the range of a prefix configured in the table, the corresponding threshold value determines when the router should switch over from the shared tree to the source specific tree. The switchover is attempted only if the traffic rate on the shared tree for the group exceeds the configured threshold.

In the absence of any matching prefix in the table, the default behavior is to switchover when the first packet is seen. In the presence of multiple prefixes matching a given group, the most specific entry is used.

Parameters*grp-ip-address/mask* | *grp-ip-address*

specifies the multicast group IP address or the IP address and mask length

netmask

specifies the subnet mask in dotted-decimal notation

spt-threshold

specifies the configured threshold, in kilobits per second (kb/s), for the group to which this (S,G) belongs. For a group (G) configured with a threshold, switchover to SPT for an (S,G) is attempted only if the (S,G) rate exceeds this configured threshold. When the **infinity** keyword is specified, no switchover will occur at any time, regardless of the traffic level detected.

Values 1 to 4294967294 | infinity (threshold in kbps)**ssm-default-range-disable****Syntax****[no] ssm-default-range-disable ipv4**

Context

```
config>service>vprn>pim
```

Description

This command specifies whether to disable the use of default range (232/8) for SSM so that it can be used by ASM to process (*,G). When enabled, the use of the default range is disabled for SSM and it (the default range) can be used by ASM. When disabled, the SSM default range is enabled.

The **no** form of the command enables the use of the default range.

Default

no ssm-default-range-disable (enabled)

ssm-groups

Syntax

```
[no] ssm-groups
```

Context

```
config>service>vprn>pim
```

Description

This command enables access to the context to enable a source-specific multicast (SSM) configuration instance.

Default

n/a

group-range

Syntax

```
[no] group-range {ip-prefix/mask | ip-prefix netmask}
```

Context

```
config>service>vprn>pim>ssm-groups
```

Description

This command configures the group address or range of group addresses for which this router can be the rendezvous point (RP).

Use the **no** form of this command to remove the group address or range of group addresses for which this router can be the RP from the configuration.

Default

n/a

Parameters

ip-prefix/mask | ip-prefix

specifies the IP address or address ranges or the IP address or address ranges and mask length for which this router can be an RP

netmask

specifies the subnet mask in dotted-decimal notation

7.8.2.1.7 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 (plaintext) 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 (such as #, \$, or spaces), 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 (such as #, \$, or spaces), 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 (such as #, \$, or spaces), 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 resets 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 (such as #, \$, or spaces), the entire string must be enclosed within double quotes.

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>vpn>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>vpn>rip

config>service>vpn>rip>group

config>service>vpn>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 are 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 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

7.8.2.1.8 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 autobind 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>vpn>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>vpn>zone

Description

This command applies a security log to the specified zone. The security log must already be configured in the **config>security>policy** context.

The **no** form of this command removes logging for the zone.

Parameters

- log-id

the identifier for the log

Values1 to 32 characters
- name

the name of the log

Values1 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

Values1 to 32 characters (must start with a letter). If the string contains special characters (such as #, \$, or spaces), 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

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 (such as #, \$, or spaces), 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>vpn>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

7.8.2.1.9 VPRN raw socket IP transport configuration commands

ip-transport

Syntax

[no] **ip-transport** *ipt-id* [create]

no ip-transport *ipt-id*

Context

config>service>vpn

Description

This command creates an IP transport subservice within a VPRN service. An IP transport subservice is used to transmit serial raw socket data to and from a local host and remote host.

All IP transport subservices must be explicitly created using the **create** keyword. An IP transport subservice is owned by the service within which it is created. An IP transport subservice can only be associated with a single service. The **create** keyword is not needed when editing parameters for an existing IP transport subservice. An IP transport subservice must be first shut down before changes can be made to the configured parameters.

The **no** form of this command deletes the IP transport subservice with the specified *ipt-id*. When an IP transport subservice is deleted, all configured parameters for the IP transport subservice are also deleted.

Default

no ip-transport

Parameters

ipt-id

the IP transport subservice physical port identifier. The *ipt-id* must reference an RS-232 serial port that has been configured as a **socket** and has its encapsulation type set to **raw**. See the 7705 SAR Interface Configuration Guide, "Serial commands", for more information.

Values value in the format *slot/mda/port.channel*

create

creates this IP transport subservice

dscp

Syntax

dscp *dscp-name*

Context

config>service>vprn>ip-transport

Description

This command configures the DSCP name used to mark the DSCP field in IP transport packets originating from this node.

Raw socket traffic redirection to a specific queue is enabled by the **fc** command.

Default

ef

Parameters

dscp-name

the DSCP name used to mark the DSCP field in IP transport packets

Table 139: Valid DSCP names

dscp-name
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

Syntax

fc fc-name profile {in | out}

Context

config>service>vprn>ip-transport

Description

This command configures the forwarding class and profile marking for IP transport packets originating from this node.

Default

ef for fc, in for profile

Parameters

fc-name
the forwarding class name to use for the IP transport packets

Values be, l2, af, l1, h2, ef, h1, nc

profile {in| out}
specifies the profile marking for the IP transport packets, either in or out

filter-unknown-host

Syntax

[no] filter-unknown-host

Context

config>service>vprn>ip-transport

Description

This command filters connections from unknown hosts. An unknown host is any host that is not configured as a remote host.

The **no** form of this command disables the filter.

Default

no filter-unknown-host

local-host

Syntax

local-host *ip-addr ip-addr port-num port-num protocol {tcp | udp}*

no local-host

Context

config>service>vprn>ip-transport

Description

This command creates the local host within the IP transport subservice.

The local host is required to accept TCP/UDP sessions initiated from far-end remote hosts, and for the node to initiate sessions toward the far-end remote hosts.

The **no** form of this command deletes the local host.

Default

no local-host

Parameters

ip-addr

the IP address that is used for this local host. The IP address must be the same as a loopback or local interface IP address that is already configured within this service.

port-num

the port number that is used by remote hosts to establish TCP/UDP sessions to this local host

Values 1026 to 49150

protocol {tcp | udp}

the protocol type that is used for all sessions to and from this local host, either tcp or udp

remote-host

Syntax

remote-host *host-id* **ip-addr** *ip-addr* **port-num** *port-num* [**create**]

no remote-host *host-id*

Context

config>service>vprn>ip-transport

Description

This command creates a remote host within the IP transport subservice. Multiple remote hosts may be created in order to send serial raw socket IP transport data to multiple destinations. The **create** keyword must be used for each remote host that is created.

The **no** form of this command deletes the remote host.

Default

no remote-host

Parameters

host-id

the remote host identifier

Values 1 to 2147483647 or a name string up to 64 characters long

ip-addr

the IP address that is used to reach the remote host in order to route IP transport packets to that remote host

port-num

the destination port number that is used to reach the serial port socket on the remote host

Values 1 to 65535

create

creates this remote host

name

Syntax

name *host-name*

no name

Context

config>service>vprn>ip-transport>remote-host

Description

This command configures a unique name for this remote host.

The **no** form of this command deletes the remote host name.

Default

n/a

Parameters

host-name

a unique name for this remote host, up to 64 characters long

```
tcp
```

Syntax

tcp

Context

config>service>vprn>ip-transport

Description

This command enables the context to configure TCP parameters within this IP transport subservice.

Default

n/a

```
inactivity-timeout
```

Syntax

inactivity-timeout *seconds*

Context

config>service>vprn>ip-transport>tcp

Description

This command specifies how long to wait before disconnecting a TCP connection due to traffic inactivity over the connection.

Default

30 s

Parameters

seconds

how long to wait, in seconds, before disconnecting a TCP connection

Values 1 to 65535

max-retries

Syntax

max-retries *number*

Context

config>service>vprn>ip-transport>tcp

Description

This command specifies the number of times that a remote host, acting as a client, tries to establish a TCP connection after the initial attempt fails.

Default

5

Parameters

number

the number of attempts to establish a TCP connection after the initial attempt fails

Values 1 to 10

retry-interval

Syntax

retry-interval *seconds*

Context

config>service>vprn>ip-transport>tcp

Description

This command specifies how long to wait before each TCP **max-retries** attempt.

Default

5 s

Parameters

seconds

how long to wait, in seconds, before each TCP **max-retries** attempt

Values 1 to 300

7.8.2.1.10 Multicast VPN commands

mvpn

Syntax

mvpn

Context

config>service>vpn

Description

This command enables the context to configure MVPN-related parameters for the IP VPN.

auto-discovery

Syntax

auto-discovery [default]

Context

config>service>vpn>mvpn

Description

This command enables MVPN membership auto-discovery through BGP. When auto-discovery is enabled, PIM peering on the inclusive provider tunnel is disabled. Changing the auto-discovery configuration requires a shutdown of this VPRN instance.

Default

default

Parameters

default

enables auto-discovery route exchange based on the format defined in NG-MVPN (RFC 6514)

c-mcast-signaling

Syntax

c-mcast-signaling **bgp**

Context

config>service>vprn>mvpn

Description

This command specifies BGP for PE-to-PE signaling of CE multicast states.

Default

bgp

Parameters

bgp

specifies to use BGP for PE-to-PE signaling of CE multicast states. Auto-discovery must be enabled.

mdt-type

Syntax

mdt-type {**sender-only** | **receiver-only** | **sender-receiver**}

no mdt-type

Context

config>service>vprn>mvpn

Description

This command allows the restriction of an MVPN instance per PE node to a specific role. By default, an MVPN instance on a PE node assumes the role of a sender as well as a receiver. This creates a mesh of MDT/PMSI across all PE nodes from this PE.

This command provides an option to configure either a sender-only or receiver-only mode per PE node. Restricting the role of a PE node avoids creating a full mesh of MDT/PMSI across all PE nodes that are participating in the MVPN instance.

The **no** version of this command restores the default (sender-receiver).

Default

sender-receiver

Parameters**sender-only**

MVPN has only senders connected to the PE node

receiver-only

MVPN has only receivers connected to the PE node

sender-receiver

MVPN has both senders and receivers connected to the PE node

provider-tunnel**Syntax**

provider-tunnel

Context

config>service>vprn>mvpn

Description

This command enables the context to configure tunnel parameters for the MVPN.

inclusive**Syntax**

inclusive

Context

config>service>vprn>mvpn>pt

Description

This command enables the context for specifying inclusive provider tunnels.

mldp**Syntax**

[no] mldp

Context

config>service>vprn>mvpn>pt>inclusive

config>service>vprn>mvpn>provider-tunnel>selective

Description

This command enables the use of an mLDP LSP for the provider tunnel.

Default

no mldp

shutdown

Syntax

[no] shutdown

Context

config>service>vprn>mvpn>ptl>inclusive>mldp

config>service>vprn>mvpn>provider-tunnel>selective>mldp

Description

This command administratively disables or enables the use of an mLDP LSP for the provider tunnel.

Default

no shutdown

selective

Syntax

selective

Context

config>service>vprn>mvpn>provider-tunnel

Description

This command enables the context to specify selective provider tunnel parameters.

Default

n/a

data-delay-interval

Syntax

data-delay-interval *value*

no data-delay-interval

Context

```
config>service>vprn>mvpn>provider-tunnel>selective
```

Description

This command specifies the interval, in seconds, before a PE router connected to the source switches traffic from the inclusive provider tunnel to the selective provider tunnel.

The **no** form of the command resets the value to the default.

Default

3 s

Parameters

value

specifies the data delay interval, in seconds

Values 3 to 180

data-threshold**Syntax**

data-threshold {*c-grp-ip-addr/mask* | *c-grp-ip-addr netmask*} *s-pmsi-threshold*

no data-threshold {*c-grp-ip-addr/mask* | *c-grp-ip-addr netmask*}

Context

```
config>service>vprn>mvpn>provider-tunnel>selective
```

Description

This command specifies the data rate threshold that triggers the switch from the inclusive provider tunnel to the selective provider tunnel for (C-S, C-G) within the group range. Optionally, PE thresholds for creating or deleting NG-MVPN S-PMSI may also be specified. Omitting the PE thresholds preserves the currently set value (or defaults, if never set). Multiple statements (one per unique group) are allowed in the configuration.

The **no** form of the command removes the values from the configuration.

Default

no data-threshold

Parameters

c-grp-ip-addr/mask | *c-grp-ip-addr*

specifies the IPv4 multicast group address or the IPv4 address and mask length

netmask

specifies the subnet mask in dotted-decimal notation

s-pmsi-threshold

specifies the rate, in kb/s. If the rate for a (C-S, C-G) within the specified group range exceeds the threshold, traffic for the (C-S, C-G) will be switched to the selective provider tunnel.

Values 1 to 4294967294

maximum-p2mp-spmsi

Syntax

maximum-p2mp-spmsi *range*

no maximum-p2mp-spmsi

Context

config>service>vprn>mvpn>provider-tunnel>selective

Description

This command specifies the maximum number of LDP point-to-multipoint S-PMSI tunnels for the MVPN. When the limit is reached, no more LDP point-to-multipoint S-PMSI tunnels are created and traffic over the data threshold will stay on I-PMSI.

Default

10

Parameters

number

specifies the maximum number of LDP point-to-multipoint S-PMSI tunnels for the MVPN

Values 1 to 4000

Default 10

umh-selection

Syntax

umh-selection {highest-ip | hash-based | unicast-rt-pref}

no umh-selection

Context

config>service>vprn>mvpn

Description

This command specifies which upstream multicast hop (UMH) selection mechanism to use, highest IP address, hash-based, or preferred unicast route.

The **no** form of the command resets it back to the default.

Default

umh-selection highest-ip

Parameters

highest-ip

specifies that the highest IP address is selected as the UMH

hash-based

specifies that the UMH selection is based on the hash based procedures

unicast-rt-pref

when selected, preferred unicast route will decide which UMH is chosen. All PE routers must prefer the same route to the UMH for the UMH selection criterion (for example, BGP path selection criteria must not influence one PE to choose a different UMH from another PE).

vrf-export

Syntax

vrf-export unicast

vrf-export *policy-name* [*policy-name...* (up to 15 max)]

no vrf-export

Context

config>service>vprn>mvpn

Description

This command specifies the export policy (up to 15) to control MVPN routes exported from the local VRF to other VRFs on the same or remote PE routers.

Default

vrf-export unicast

Parameters

unicast

specifies to use the unicast VRF export policy for the MVPN

policy-name

the route policy name

vrf-import

Syntax

vrf-import unicast

vrf-import *policy-name* [*policy-name...* (up to 15 max)]

no vrf-import

Context

config>service>vprn>mvpn

Description

This command specifies the import policy (up to 15) to control MVPN routes imported to the local VRF from other VRFs on the same or remote PE routers.

Default

vrf-import unicast

Parameters

unicast

specifies to use a unicast VRF import policy for the MVPN

policy-name

the route policy name

vrf-target

Syntax

vrf-target {**unicast** | *ext-community* | **export unicast** | *ext-community* | **import unicast** | *ext-community*}

no vrf-target

Context

config>service>vprn>mvpn

Description

This command specifies the route target to be added to the advertised routes or compared against the received routes from other VRFs on the same or remote PE routers. The VRF import or VRF export policies override the VRF target policy.

The **no** form of the command removes the VRF target.

Default

no vrf-target

Parameters

unicast

specifies to use the unicast **vrf-target** *ext-community* for the multicast VPN

ext-community

an extended BGP community in the **type:x:y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. The **x** and **y** values are 16-bit integers.

Values

target:{*ip-address:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*}

<i>ip-address:</i>	a.b.c.d
<i>comm-val:</i>	0 to 65535
<i>2byte-asnumber:</i>	0 to 65535
<i>ext-comm-val:</i>	0 to 4294967295
<i>4byte-asnumber</i>	0 to 4294967295

export

Syntax

export {**unicast** | *ext-community*}

Context

config>service>vpn>mvpn>vrf-target

Description

This command specifies communities to be sent to peers.

Parameters

unicast

specifies to use the unicast **vrf-target** *ext-community* for the multicast VPN

ext-community

an extended BGP community in the **type:x:y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. The **x** and **y** values are 16-bit integers.

Values

target:{*ip-address:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*}

<i>ip-address:</i>	a.b.c.d
<i>comm-val:</i>	0 to 65535
<i>2byte-asnumber:</i>	0 to 65535

<i>ext-comm-val:</i>	0 to 4294967295
<i>4byte-asnumber</i>	0 to 4294967295

import

Syntax

import {unicast | *ext-community*}

Context

config>service>vpn>mvpn>vrf-target

Description

This command specifies communities to be accepted from peers.

Parameters

unicast

specifies to use the unicast **vrf-target** *ext-community* for the multicast VPN

ext-community

an extended BGP community in the **type:x:y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. The **x** and **y** values are 16-bit integers.

Values

target:{*ip-address:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*}

<i>ip-address:</i>	a.b.c.d
<i>comm-val:</i>	0 to 65535
<i>2byte-asnumber:</i>	0 to 65535
<i>ext-comm-val:</i>	0 to 4294967295
<i>4byte-asnumber</i>	0 to 4294967295

7.8.2.1.11 MSDP commands

msdp

Syntax

[no] msdp

Context

```
config>service>vprn
```

Description

This command enables a Multicast Source Discovery Protocol (MSDP) instance. When an MSDP instance is created, the protocol is enabled. To start or suspend execution of the MSDP protocol without affecting the configuration, use the **[no] shutdown** command.

For MSDP to function, at least one peer must be configured.

When MSDP is configured and started, an event message is generated.

Before the **no** form of the command is executed, all sessions are terminated and an event message is generated.

When all peering sessions are terminated, event messages are not generated for each peer.

The **no** form of the command deletes the MSDP instance, removing all associated configuration parameters.

Default

```
no msdp
```

active-source-limit

Syntax

```
active-source-limit number
```

```
no active-source-limit
```

Context

```
config>service>vprn>msdp
```

```
config>service>vprn>msdp>group
```

```
config>service>vprn>msdp>group>peer
```

```
config>service>vprn>msdp>peer
```

```
config>service>vprn>msdp>source
```

Description

This command controls the maximum number of source-active (SA) messages that will be accepted by MSDP, which controls the number of active sources that can be stored on the system.

The **no** form of this command resets the SA message limit to its default operation.

Default

```
no active-source-limit
```

Parameters

number

defines how many active sources can be maintained by MSDP

Values 0 to 1000000

data-encapsulation

Syntax

[no] data-encapsulation

Context

config>service>vprn>msdp

Description

This command configures a rendezvous point (RP) that uses MSDP to encapsulate multicast data received in MSDP register messages inside forwarded MSDP SA messages.

Default

data-encapsulation

export

Syntax

export *policy-name* [*policy-name...*(up to 5 max)]

no export

Context

config>service>vprn>msdp

config>service>vprn>msdp>peer

config>service>vprn>msdp>group

config>service>vprn>msdp>group>peer

Description

This command specifies the policies to export the SA state from the SA list into MSDP.

If multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five policy names can be specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command.

If you configure an export policy at the global level, each individual peer inherits the global policy. If you configure an export policy at the group level, each individual peer in a group inherits the group's policy. If you configure an export policy at the peer level, the policy only applies to the peer where it is configured.

The **no** form of the command removes all policies from the configuration and all SA entries are allowed.

Default

no export

Parameters

policy-name

specifies the export policy name. Up to five policy names can be specified.

group

Syntax

[no] group *group-name*

Context

config>service>vpn>msdp

Description

This command enables access to the context to create or modify an MSDP group. To configure multiple MSDP groups, multiple group statements must be included in the configuration.

By default, the group's parameter settings are inherited from the global MSDP parameter settings. To override the global settings, group-specific settings within the group can be configured.

If the specified group name is already configured, this command enables the context to configure or modify group-specific parameters.

If the specified group name is not already configured, this command creates the group and enables the context to configure the group-specific parameters.

For a group to be functional, at least one peer must be configured.

Default

no group

Parameters

group-name

specifies a unique name for the MSDP group

import

Syntax

import *policy-name* [*policy-name...*(up to 5 max)]

no import

Context

```
config>service>vprn>msdp
config>service>vprn>msdp>peer
config>service>vprn>msdp>group
config>service>vprn>msdp>group>peer
```

Description

This command specifies the policies to import the SA state from MSDP into the SA list.

If multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five policy names can be specified. The first policy that matches is applied. If multiple import commands are issued, the last command entered will override the previous command.

If you configure an import policy at the global level, each individual peer inherits the global policy.

If you configure an import policy at the group level, each individual peer in a group inherits the group's policy.

If you configure an import policy at the peer level, the policy only applies to the peer where it is configured.

The **no** form of the command removes all policies from the configuration and all SA messages are allowed.

Default

no import

Parameters

policy-name

specifies the import policy name. Up to five policy names can be specified.

local-address

Syntax

local-address *address*

no local-address

Context

```
config>service>vprn>msdp
config>service>vprn>msdp>peer
config>service>vprn>msdp>group
config>service>vprn>msdp>group>peer
```

Description

This command configures the local end of an MSDP session. For MSDP to function, at least one peer must be configured. When configuring a peer, you must include this **local-address** command. This address must be present on the node and is used to validate incoming connections to the peer and to establish connections to the remote peer.

When the address is configured, it is validated and will be used as the local address for MSDP peers from that point. If a subsequent **local-address** command is entered, it will replace the existing configuration and existing sessions will be terminated.

Similarly, when the **no** form of this command is entered, the existing **local-address** will be removed from the configuration and the existing sessions will be terminated.

Whenever a session is terminated, all information pertaining to and learned from that peer will be removed.

Whenever a new peering session is created or a peering session is lost, an event message is generated.

The **no** form of this command removes the **local-address** from the configuration.

Default

no local-address

Parameters

address

specifies an existing address on the node

mode

Syntax

mode {**mesh-group** | **standard**}

Context

config>service>vprn>msdp>group

Description

This command configures groups of peers either in non-meshed mode or in a full mesh topology to limit excessive flooding of SA messages to neighboring peers. When the mode is specified as **mesh-group**, SA messages received from a mesh group member are always accepted but are not flooded to other members of the same mesh group. These SA messages are only flooded to non-mesh-group peers or members of other mesh groups.

In a meshed configuration, all members of the group must have a peer connection with every other mesh group member. If this rule is not adhered to, unpredictable results may occur.

Default

standard

Parameters

mesh-group

specifies that all members of the group have full mesh MSDP connectivity with each other

standard

specifies a non-meshed mode

peer

Syntax

[no] **peer** *peer-address*

Context

config>service>vprn>msdp

config>service>vprn>msdp>group

Description

This command configures an MSDP peer or MDSP group peer. MSDP must have at least one peer configured. A peer is defined by configuring a **local-address** that is used by the local node to set up a peering session and by configuring the address of a remote MSDP router. It is the address of this remote peer that is configured with this command.

After peer relationships are established, the MSDP peers exchange messages to advertise active multicast sources. If multiple peering sessions are required, multiple peer statements should be included in the configuration.

By default, the parameters applied to a peer are inherited from the global or group level. To override these inherited settings, the parameters must be configured at the peer level.

If the specified peer address is already a configured peer, this command enables the context to configure or modify the peer-specific parameters.

If the specified peer address is not already a configured peer, this command creates the peer instance and enables the context to configure the peer-specific parameters.

The peer address is validated and, if valid, will be used as the remote address for an MSDP peering session.

When the **no** form of this command is entered, the existing peering address is removed from the configuration and the existing session is terminated. Whenever a session is terminated, all SA information pertaining to and learned from that peer is removed. Whenever a new peering session is created or a peering session is lost, an event message is generated.

Default

n/a

Parameters

peer-address

specifies the peer address that identifies the remote MSDP router with which the peering session will be established

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

```
config>service>vprn>msdp>group>peer
```

```
config>service>vprn>msdp>peer
```

Description

This command configures a Message Digest 5 (MD5) authentication key to be used with a specific MSDP peering session. The authentication key must be configured per peer; therefore, no global or group configuration is possible.

Using the **no** form of the command accepts all MSDP messages and disables the MD5 signature option authentication key.

Default

no authentication-key

Parameters

authentication-key

specifies the authentication key. Allowed values are any string up to 256 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed in quotation marks (" ").

hash-key

specifies the hash key. The key can be any combination of ASCII characters up to 451 characters in length (encrypted). If spaces are used in the string, the entire string must be enclosed in quotation marks (" ").

This parameter 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** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

specifies that the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

default-peer

Syntax

[no] default-peer

Context

```
config>service>vprn>msdp>peer  
config>service>vprn>msdp>group>peer
```

Description

This command enables the default peer mechanism, where a peer can be selected as the default MSDP peer. As a result, all SA messages from the peer will be accepted without the usual peer reverse path forwarding (RPF) check.

The MSDP peer-RPF check is different from the normal multicast RPF checks. The peer-RPF check is used to stop SA messages from looping. A router validates SA messages originated from other routers in a deterministic fashion.

A set of rules is applied to validate received SA messages, and the first rule that applies determines the peer-RPF neighbor. All SA messages from other routers are rejected. The following rules are applied to SA messages originating at router_S and received at router_R from router_N.

- If router_N and router_S are the same, the message is originated by a direct peer-RPF neighbor and is accepted.
- If router_N is a configured peer or a member of the router_R mesh group, its SA messages are accepted.
- If router_N is the BGP next hop of the active multicast RPF route toward router_S, then router_N is the peer-RPF neighbor and its SA messages are accepted.
- If router_N is an external BGP peer of router_R and the last autonomous system (AS) number in the BGP AS-path to router_S is the same as router_N's AS number, then router_N is the peer-RPF neighbor and its SA messages are accepted.
- If router_N uses the same next hop as the next hop to router_S, then router_N is the peer-RPF neighbor and its SA messages are accepted.
- If router_N fits none of the above rules, then router_N is not a peer-RPF neighbor and its SA messages are rejected.

When the **no** form the command is issued, no default peer is established and all SA messages are RPF checked.

Default

```
no default-peer
```

receive-msdp-msg-rate

Syntax

```
receive-msdp-msg-rate number interval seconds [threshold number]  
no receive-msdp-msg-rate
```

Context

```
config>service>vprn>msdp  
config>service>vprn>msdp>group
```

```
config>service>vprn>msdp>group>peer
config>service>vprn>msdp>peer
```

Description

This command limits the number of MSDP messages that are read from the TCP session to prevent an MSDP RP router from receiving a large number of MSDP message packets in an SA message.

After the number of MSDP packets (including SA messages) defined by the **threshold** *number* have been processed, all other MSDP packets are rate-limited. Messages from the TCP session are no longer accepted until the configured **interval** *seconds* has elapsed. Setting the threshold is useful during at system startup and initialization. No limit is placed on the number of MSDP and SA messages that will be accepted.

The **no** form of this command resets the message limit to its default operation.

Default

n/a

Parameters

receive-msdp-msg-rate *number*

specifies the number of MSDP messages (including SA messages) that are read from the TCP session per **interval** *seconds*

Values 10 to 10000

Default 0

seconds

specifies the interval of time in which the number of MSDP messages set by the **receive-msdp-msg-rate** *number* parameter are read from the TCP session

Values 1 to 600

Default 0

threshold *number*

specifies the number of MSDP messages that can be processed before the MSDP message rate-limiting function is activated

Values 1 to 1000000

Default 0

rpf-table

Syntax

rpf-table {*rtable-m* | *rtable-u* | **both**}

no rpf-table

Context

```
config>service>vprn>msdp
```

Description

This command configures the sequence of route tables used to find an RPF interface for a particular multicast route.

By default, only the unicast route table is looked up to calculate an RPF interface toward the source/rendezvous point. However, the operator can specify one of the following options:

- use the unicast route table only
- use the multicast route table only
- use both route tables

Default

rtable-u

Parameters

rtable-m

specifies that only the multicast route table is used by the multicast protocol (PIM) for IPv4 RPF checks. This route table contains routes submitted by static routes, ISIS, and OSPF.

rtable-u

specifies that only the unicast route table is used by the multicast protocol (PIM) for IPv4 RPF checks. This route table contains routes submitted by all unicast routing protocols.

both

specifies that the first lookup is always in the multicast route table, and if there is a route, it will use it. If PIM does not find a route in the first lookup, it tries to find it in the unicast route table.

sa-timeout

Syntax

sa-timeout *seconds*

no sa-timeout

Context

```
config>service>vprn>msdp
```

Description

This command configures the timeout value for the SA entries in the cache. If these entries are not refreshed within the timeout value, they are removed from the cache. Normally, the entries are refreshed at least once a minute. However, under high load with many MSDP peers, the refresh cycle could be incomplete. A higher timeout value (more than 90 seconds) could be useful to prevent instabilities in the MSDP cache.

Default

90

Parameters

seconds

specifies the time, in seconds, to wait for a response from the peer before declaring the peer unavailable

Values 90 to 600

source

Syntax

[no] source *ip-prefix/mask*

Context

config>service>vpn>msdp

Description

This command configures an MSDP source.

If the specified prefix and mask is already configured, this command enables the context to configure or modify the source-specific parameters.

If the specified prefix and mask is not already configured, this command creates the source node instance and enables the context to configure the source-specific parameters.

The SA messages are not rate-limited based on the source address range.

The **no** form of this command removes the sources in the address range.

Default

n/a

Parameters

ip-prefix/mask

specifies the IP prefix and mask length for the MSDP source

7.8.2.1.12 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 (such as #, \$, or spaces), 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 about 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 and prefix length

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. The address generated from a prefix that is no longer preferred 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)

use-virtual-mac

Syntax

[no] use-virtual-mac

Context

config>service>vprn>router-advertisement>interface

Description

This command enables the sending of router advertisement messages using the VRRP virtual MAC address, provided that the virtual router is currently the master.

If the virtual router is not the master, no router advertisement messages are sent.

The **no** form of the command disables the sending of router advertisement messages.

Default

no use-virtual-mac

7.8.2.1.13 Local DHCP and DHCPv6 server commands

For complete descriptions of all local DHCP and DHCPv6 server commands, see 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

7.8.2.1.14 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). When 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 the following table.

Table 140: 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/mask | ip-address
    the IP address or the IP address and subnet mask length of the IP interface
```


netmask

the subnet mask in dotted-decimal notation

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 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>if>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.

cflowd-parameters

Syntax

cflowd-parameters

Context

config>service>vprn>interface

Description

This command enables the context to configure cflowd parameters for the specified IP interface.

Cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement.

Default

n/a

sampling

Syntax

sampling {unicast | multicast} type {interface} [direction {ingress-only | egress-only | both}]

no sampling {unicast | multicast}

Context

config>service>vprn>if>cflowd-parameters

Description

This command configures the cflowd sampling behavior to collect traffic flow samples through a router for analysis.

This command can be used to configure the sampling parameters for unicast and multicast traffic separately.

If cflowd sampling is enabled with no **direction** parameter specified, **ingress-only** sampling is enabled by default.

The **no** form of the command disables the specified type of traffic sampling on the interface.

Default

no sampling unicast

no sampling multicast

Parameters

unicast

cflowd will sample unicast traffic on the interface

multicast

cflowd will sample multicast traffic on the interface

interface

specifies that all traffic entering or exiting the interface is subject to sampling. Interface is the only sampling type supported on the 7705 SAR and must be specified with this command.

direction

specifies the direction in which to collect traffic flow samples: **ingress-only**, **egress-only**, or **both**

hold-time

Syntax

hold-time

Context

config>service>vprn>interface

Description

This command enables the CLI context to configure interface hold-up or hold-down timers.

Default

n/a

down

Syntax

down ip *seconds* [**init-only**]

no down ip

down ipv6 *seconds* [**init-only**]

no down ipv6

Context

config>service>vprn>if>hold-time

Description

This command enables a delay in the activation of the IPv4 or IPv6 interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface up,

unless the **init-only** option is configured. If the **init-only** option is configured, the delay is only applied when the IP interface is first configured or after a system reboot.

The **no** form of this command disables the delay in the activation of the IPv4 or IPv6 interface. Removing the configuration during an active delay period stops the delay period immediately.

Default

n/a

Parameters

ip

specifies that the configured **down** delay is applied to an IPv4 interface

ipv6

specifies that the configured **down** delay is applied to an IPv6 interface

seconds

specifies the time delay, in seconds, before the interface is activated

Values 1 to 1200

init-only

specifies that the configured **down** delay is applied only when the interface is first configured or after a reboot

up

Syntax

up ip seconds

no up ip

up ipv6 seconds

no up ipv6

Context

config>service>vprn>if>hold-time

Description

This command enables a delay in the deactivation of the IPv4 or IPv6 interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface down.

The **no** form of this command disables the delay in the deactivation of the IPv4 or IPv6 interface. Removing the configuration during an active delay period stops the delay period immediately.

Default

n/a

Parameters

- ip**
specifies that the configured **up** delay applies to an IPv4 interface
 - ipv6**
specifies that the configured **up** delay applies to an IPv6 interface
 - seconds**
specifies the time delay, in seconds, before the interface is deactivated
- Values** 1 to 1200

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 – uses the value derived from the port MTU

Parameters

- octets**
specifies the MTU for this interface
- Values** 128 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 addresses 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

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

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

load-balancing

Syntax

load-balancing

Context

config>service>vprn>interface

Description

This command enables the context to configure load balancing hashing options on the interface. The options enabled at the interface level overwrite parallel system-level configurations.

Default

n/a

I4-load-balancing

Syntax

I4-load-balancing *hashing-algorithm*

no I4-load-balancing

Context

config>service>vprn>interface>load-balancing

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

hashing-algorithm

specifies that Layer 4 source and destination port fields are included in or excluded from the hashing calculation

Values includeL4: include Layer 4 source and destination port fields in the hashing calculation for TCP/UDP packets
 excludeL4: exclude Layer 4 source and destination port fields in the hashing calculation for TCP/UDP packets

Default the system configuration setting (under the **config>system** context)

spi-load-balancing

Syntax

[no] spi-load-balancing

Context

config>service>vprn>interface>load-balancing

Description

This command enables use of the SPI in hashing for ESP/AH encrypted IPv4 or IPv6 traffic at the interface level.

The **no** form of this command disables SPI hashing.

Default

no spi-load-balancing

teid-load-balancing

Syntax

[no] teid-load-balancing

Context

config>service>vprn>interface>load-balancing

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

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 interface is a loopback interface that has no associated physical interface. If this command is enabled, a SAP cannot be defined on the interface.

Default

no loopback

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

a 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 and cannot be all zeros. Allowed values are any non-broadcast, non-multicast MAC, and non-IEEE reserved MAC addresses.

multicast-translation

Syntax

[**no**] **multicast-translation**

Context

config>service>vprn>interface

Description

This command enables multicast address translation on the 7705 SAR that is the translator router for unicast-to-multicast or multicast-to-multicast translation.

When enabled for unicast-to-multicast translation, the 7705 SAR will try to find the source and destination address of the packet in the unicast-to-multicast translation table. If the source and destination address is not found, the packet is processed as a regular IP packet. To disable unicast-to-multicast translation, all entries must be removed from the translation table and then the command must be set to **no multicast-translation**.

When enabled for multicast-to-multicast translation, the static group configuration is used for multicast PDUs that arrive on the node and are to be translated via the translation table. If the command is enabled and an arriving PDU does not match an entry in the translation table, the multicast PDU is dropped. If the (S,G) arrives from another interface via a dynamic protocol while this command is enabled, the interface that the dynamic (S,G) arrived from will be added as an outgoing interface but it will not forward traffic. Only the outgoing loopback interface on the translation router will forward the translated PDU.

For multicast-to-multicast translation, if this command is not enabled, the node will function as a leaf for the static group configuration. To disable multicast-to-multicast translation, the interface must be shut down before the **no** version of this command is issued.

Default

no multicast-translation

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 (see the 7705 SAR 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 (such as #, \$, or spaces), 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/mask | *ip-address*

the IP address or the IP address and mask length of the IP interface

netmask

the subnet mask in dotted-decimal notation

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 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 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 IPv4 address for the static ARP

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 and cannot be all zeros. 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.

tcp-mss

Syntax

tcp-mss *value*

no tcp-mss

Context

config>service>vprn>interface

config>service>vprn>if>ipv6

Description

This command configures the maximum segment size (MSS) in a TCP SYN or SYN-ACK packet during the establishment of a TCP connection. A **tcp-mss** value can be specified on an ingress interface, egress interface, or both. When configured on two interfaces, the smaller of the two values is used. If the TCP SYN packet has no TCP MSS field, the 7705 SAR assigns it the MSS value configured on the interface and recalculates the IP checksum. If the TCP SYN or SYN-ACK packet has an MSS field and the value is greater than the value configured on the interface, the 7705 SAR overwrites the packet MSS value with the lower value. If the MSS value is less than the value configured on the interface, the packet MSS value does not change. See the 7705 Router Configuration Guide, "TCP MSS Configuration and Adjustment", for more information.

This command is supported on interfaces with IPv4 and IPv6 traffic, and a different MSS value can be configured for the IPv4 and IPv6 interfaces. This command is supported on IPsec private interfaces in a VPRN.

Default

no tcp-mss

Parameters

value

the MSS, in bytes, to be used in a TCP SYN or SYN-ACK packet

Values 384 to 9732

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

7.8.2.1.15 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 IPv6 interface address and prefix length

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 address may be 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 (such as #, \$, or spaces), 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)

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

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

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

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**]

no link-local-address

Context

config>service>vprn>if>ipv6

Description

This command configures the IPv6 link-local address.

The **no** form of the command removes the configured link-local address, and the router automatically generates a default link-local address.

Removing a manually configured link-local address may impact routing protocols that have a dependency on that address.

Default

n/a

Parameters

ipv6-address

the IPv6 link-local address

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 address may be 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

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

7.8.2.1.16 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

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 (toward 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 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 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) or 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 the *vlan-id* and dot1p bits, in addition to the **ascii-tuple**. The format is supported on dot1q and qinq ports only. 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 Nokia 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 Nokia vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the MAC address in the Nokia 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 Nokia vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the SAP ID in the Nokia 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 Nokia vendor-specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the Nokia 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 Nokia 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>vpn>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 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

7.8.2.1.17 Interface ICMP commands

icmp

Syntax

icmp

Context

config>service>vpn>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>vpn>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

7.8.2.1.18 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](#).

If the VPRN interface has been configured as a loopback interface with the [loopback](#) command, a SAP cannot be defined on the interface.

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 44: SAP ID configurations](#) 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>vpn>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>vpn>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) Ethernet adapter card, the scheduler mode must be set to 16-priority. On a third-generation (Gen-3) Ethernet adapter card, the scheduler mode is always 4-priority. For information about adapter card generations, see the "Evolution of Ethernet Adapter Cards, Modules, and Platforms" section in the 7705 SAR 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**.

If the specified SAP is a LAG SAP, *agg-rate* and *cir-rate* can be configured regardless of the scheduler mode setting on Gen-2 or Gen-3 hardware. If the active port is on a Gen-3 card or platform, *agg-rate* and *cir-rate* are applicable. If the active port is on a Gen-2 card or platform, *agg-rate* and *cir-rate* apply when the scheduler mode is set to 16-priority. For details on the behavior of a mix-and-match LAG SAP, see the "LAG Support on Third-Generation Ethernet Adapter Cards, Ports, and Platforms" and "Network LAG Traffic Management" sections in the 7705 SAR Interface Configuration Guide.

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.

The following table 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-dot1p 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.

Table 141: 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

Port/SAP type	Match-qinq-dot1p setting ¹	Existing packet tags	P-bits used for match
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.

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.

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).

The following table 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 142: 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
Dot1q/ 0 ²	n/a	Outer tag
QinQ/ X.Y ¹	False	Two outer tags ³
	True	Outer tag ³
QinQ/ X.* ¹	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.

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, which only support 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.

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

7.8.2.1.19 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 destination 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 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

7.8.2.1.20 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 *ip-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

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)

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)

7.8.2.1.21 Interface VRRP commands

vrrp

Syntax

vrrp *virtual-router-id* [**owner**] [**passive**]

no vrrp *virtual-router-id*

Context

config>service>vprn>interface

config>service>vprn>if>ipv6

Description

This command creates or edits a virtual router ID (VRID) 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

owner

keyword used to identify this virtual router instance as owning the virtual router IP addresses. If the **owner** keyword is not specified at the time of VRID creation, the **vrrp backup** command must be used to define the virtual router IP addresses. The **owner** keyword is not required when entering the VRID for editing purposes. When created as **owner**, a VRID on an IP interface cannot have the **owner** parameter removed. The VRID must be deleted, and then recreated without the **owner** keyword, to remove ownership.

passive

keyword used to identify this virtual router instance as **passive**, owning the virtual router IP addresses. A **passive** VRID does not send or receive VRRP advertisement messages and is always in either the master state (if the interface is operationally up), or the initialize state (if the interface is operationally down). The **passive** keyword is not required when entering the VRID for editing purposes. When a VRID on an IP interface is created as

passive, the parameter cannot be removed from the VRID. The VRID must be deleted, and then recreated without the **passive** keyword, to remove the parameter.

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 alphanumeric 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) for a hash key or up to 121 characters for a hash2 key. If spaces are used in the string, the entire string must be enclosed 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 that 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 that 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*

[no] **backup** *ipv6-address*

Context

config>service>vprn>if>vrrp

config>service>vprn>if>ipv6>vrrp

Description

This command configures virtual router IP addresses for backup.

Default

n/a

Parameters

ip-address

specifies the destination IPv4 address for the backup interface

ipv6-address

specifies the destination IPv6 address for the backup interface

bfd-enable

Syntax

[no] **bfd-enable interface** *interface-name* **dst-ip** *ip-address*

[no] **bfd-enable service-id interface** *interface-name* **dst-ip** *ip-address*

Context

config>service>vprn>if>vrrp

config>service>vprn>if>ipv6>vrrp

Description

This command assigns a BFD session that provides a heartbeat mechanism for a VRRP instance. Only one BFD session can be assigned to a VRRP instance, but multiple VRRP instances can use the same BFD session.

BFD controls the state of the associated interface. By enabling BFD on a 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 specified in this command.

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 214748690 or *service-name*

interface-name

specifies the name of the interface running BFD

ip-address

specifies the destination IPv4 or IPv6 address to be used for the BFD session

init-delay

Syntax

init-delay *seconds*

no init-delay

Context

```
config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp
```

Description

This command configures a VRRP initialization delay timer.

Default

no init-delay

Parameters

seconds

specifies the length of time in seconds for the initialization delay timer for VRRP

Values 1 to 65535

mac

Syntax

mac *mac-address*

no mac

Context

```
config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>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

mac-address

specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff, where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

master-int-inherit

Syntax

[no] master-int-inherit

Context

config>service>vpn>if>vrrp

config>service>vpn>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 the backup router 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. See "Master Down Interval" in the "VRRP" chapter of the 7705 SAR Router Configuration Guide for details.

Default

no master-int-inherit

message-interval

Syntax

message-interval {[seconds] [milliseconds milliseconds]}

no message-interval

Context

config>service>vpn>if>vrrp

config>service>vpn>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.

Configuring the message interval value can be done in three ways: using only the milliseconds value, using only the seconds value, or using a combination of the two values. The following table shows the ranges for each way of configuring the message interval.

Table 143: Message interval configuration ranges

Configuration	IPv4	IPv6
Using milliseconds value only	100 to 900 ms	10 to 990 ms
Using seconds value only	1 to 255 s	1 to 40 s
Using combination milliseconds and seconds values	1 s 100 ms to 255 s 900 ms (1.1 s to 255.9 s)	1 s 10 ms to 40s 990 ms (1.01 s to 40.99 s)
Default setting	1 s	1 s

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.

Default

1 s

Parameters

seconds

the time interval, in seconds, between sending advertisement messages

Values IPv4: 1 to 255
IPv6: 1 to 40

milliseconds

specifies the time interval, in milliseconds, between sending advertisement messages. This parameter is not supported on non-redundant chassis.

Values IPv4: 100 to 900
IPv6: 10 to 990

ntp-reply

Syntax

[no] ntp-reply

Context

config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp

Description

This command enables the reception of and response to Network Time Protocol (NTP) requests directed at the VRRP virtual IP address. This behavior only applies to the router currently acting as the master VRRP.

The **no** form of this command disables NTP requests from being processed.

Default

no ntp-reply

ping-reply

Syntax

[no] ping-reply

Context

config>service>vprn>if>vrrp

config>service>vprn>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 parent 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.

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

config>service>vprn>if>ipv6>vrrp

Description

This command associates a VRRP priority control policy with the virtual router instance (non-owner context only). VRRP policies are defined under the **config>vrrp>policy** context. For details, see the "VRRP" chapter in the 7705 SAR Router Configuration Guide.

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

config>service>vprn>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 effect 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
config>service>vprn>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 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
Default	100

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 parent 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

config>service>vprn>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>vprn>if>vrrp

config>service>vprn>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 parent 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

config>service>vprn>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

7.8.2.1.22 VPRN static one-to-one NAT configuration commands

static-nat-inside

Syntax

[no] static-nat-inside

Context

```
config>service>vprn>interface
```

Description

This command configures an interface as an inside (private) interface.

By default, all interfaces are outside (public) interfaces. The **no** form of this command returns the interface to the default setting.

Default

```
no static-nat-inside
```

static-nat

Syntax

```
[no] static-nat
```

Context

```
config>service>vprn
```

Description

This command enables the context to configure static one-to-one NAT.

The **no** form of this command disables static one-to-one NAT.

Default

```
no static-nat
```

drop-packets-without-nat-entry

Syntax

```
[no] drop-packets-without-nat-entry
```

Context

```
config>service>vprn>static-nat
```

Description

This command configures the router to drop packets that are traveling from either an inside network to an outside network or an outside network to an inside network that do not have a NAT mapping entry.

By default, packets traveling from either an inside network to an outside network or an outside network to an inside network are forwarded whether or not there is a NAT mapping entry.

The **no** form of this command returns the router to the default behavior.

Default

no drop-packets-without-nat-entry

inside

Syntax

inside

Context

config>service>vprn>static-nat

Description

This command creates a static one-to-one NAT mapping from an inside network to an outside network. When configured, a packet traveling from an inside network to an outside network that matches a NAT mapping entry will have NAT applied to its source IP address. Similarly, a packet traveling from an outside network to an inside network that matches a NAT mapping entry will have NAT applied to its destination IP address.

Default

n/a

map

Syntax

map start *ip-address end ip-address to ip-address*

no map start *ip-address end ip-address*

Context

config>service>vprn>static-nat>inside

Description

This command maps a range of inside source IP addresses that will undergo NAT to a specified outside IP address range.

For example, to map the entire range of inside addresses within 192.168.0.0/16 to the outside address 10.10.0.0/16, the configuration would be:

map start 192.168.0.0 **end** 192.168.255.255 **to** 10.10.0.0

The 7705 SAR will then map each inside source IP address to its corresponding outside IP address sequentially; for example:

- inside address 192.168.0.1 maps to 10.10.0.1
- inside address 192.168.10.10 maps to 10.10.10.10
- inside address 192.168.254.100 maps to 10.10.254.100

The **no** form of this command removes the NAT mapping.

Default

no map start *ip-address* end *ip-address*

Parameters

start *ip-address*

identifies the start of the range of inside IPv4 addresses that will undergo NAT to an outside address

end *ip-address*

identifies the end of the range of inside IPv4 addresses that will undergo NAT to an outside address

to *ip-address*

identifies the outside IPv4 address that the range of inside addresses maps to

shutdown

Syntax

[no] shutdown

Context

config>service>vprn>static-nat>inside>map

Description

This command administratively disables the static NAT map entry.

The **no** form of this command administratively enables the static NAT map entry.

Default

no shutdown

7.8.2.1.23 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 [**udp-port** *udp-port-number*] [**create**]

no reflector

Context

config>service>vprn>twamp-light

Description

This command configures the TWAMP Light reflector function. The UDP port number is mandatory when creating a TWAMP Light reflector. The reflector functionality is enabled using the **no shutdown** command.

Default

disabled

Parameters

udp-port-number

the 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.

Values 862, 64364 to 64373

create

mandatory keyword when creating a TWAMP Light reflector

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 and prefix length

7.8.2.1.24 VPRN NTP commands

```
ntp
```

Syntax

[no] ntp

Context

config>service>vprn

Description

This command enables the context to configure Network Time Protocol (NTP) and its operation. It also enables NTP server mode within the VPRN routing instance so that the router will respond to NTP requests received from external clients in the VPRN.

The **no** form of this command stops the execution of NTP and removes its configuration.

Default

n/a

```
authenticate
```

Syntax

[no] authenticate

Context

config>service>vprn>ntp

Description

This command enables authentication for the NTP server.

Default

n/a

authentication-check

Syntax

[no] authentication-check

Context

config>service>vprn>ntp

Description

This command provides the option to skip the rejection of NTP PDUs that do not match the authentication key ID, type, or key values.

When authentication is configured, NTP PDUs received on an interface or the management port are authenticated on receipt and rejected if there is a mismatch in the authentication key ID, type, or key value.

When **authentication-check** is enabled, NTP PDUs are authenticated on receipt and rejected if there is a mismatch in the authentication key ID, type, or key value. Any mismatches cause a counter to be incremented: one counter for type, one for key ID, and one for key value mismatches. These counters are visible in the **show>system>ntp** command output.

The **no** form of this command allows mismatched packets to be accepted (overriding authentication); however, the counters are maintained.

Default

authentication-check

authentication-key

Syntax

authentication-key *key-id* **key** *key* [**hash** | **hash2**] **type** {**des** | **message-digest**}

no authentication-key *key-id*

Context

config>service>vprn>ntp

Description

This command sets the authentication key ID, type, and key value used to authenticate NTP PDUs that are either sent by the broadcast server function toward external clients or received from external unicast clients

within the VPRN routing instance. For authentication to work, the configured authentication key ID, type, and key values must match those of the NTP PDUs.

Configuring the **authentication-key** with a *key-id* value that matches an existing key will override the existing entry.

Recipients of the NTP packets must have the same authentication key ID, type, and key values in order to use the data transmitted by this node.

The **no** form of this command removes the authentication key.

Default

n/a

Parameters

key-id

the authentication key identifier used by the node when transmitting or receiving NTP packets

Values 1 to 255

key

the authentication key associated with the configured key ID. The configured value is the actual value used by other network elements to authenticate the NTP packet.

Values 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

specifies that the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

specifies that the key is entered in a more complex encrypted form that involves more variables than the key value alone. This means that the **hash2** encrypted key cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

des

specifies that DES authentication is used for this key. The **des** value is not supported in FIPS-140-2 mode.

message-digest

specifies that MD5 authentication in accordance with RFC 2104 is used for this key.

authentication-keychain

Syntax

authentication-keychain *keychain-name*
no authentication-keychain

Context

config>service>vprn>ntp

Description

This command configures the authentication keychain used to handle unsolicited NTP requests. If the system receives a request with a key ID that matches both the configured key and the keychain, the system first checks the MAC ID using the key information. If the key authentication fails, the system then checks the MAC ID using the keychain information.

The **no** form of the command removes the authentication keychain.

Default

no authentication-keychain

Parameters

keychain-name
the name of the keychain, up to 32 characters

broadcast

Syntax

broadcast {**interface** *ip-int-name*} [**key-id** *key-id* | **authentication-keychain** *keychain-name*] [**version** *version*] [**ttl** *ttl*]
no broadcast {**interface** *ip-int-name*}

Context

config>service>vprn>ntp

Description

This command configures the node to transmit NTP broadcast packets on the specified interface. Because broadcast messages can easily be spoofed, authentication is strongly recommended.

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

Default

n/a

Parameters

ip-int-name

the local interface on which to transmit NTP broadcast packets. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

Values 32 character maximum

key-id

identifies the configured authentication key and authentication type used by this node to receive and transmit NTP packets from and to an NTP server and peers. If an NTP packet is received by this node, both the authentication key and authentication type must be valid; otherwise, the packet will be rejected and an event or trap will be generated.

Values 1 to 255

keychain-name

specifies the name of the configured authentication keychain

version

the NTP version number that is generated by this node. This parameter does not need to be configured when the node is in NTP client mode because all versions will be accepted.

Values 2 to 4

Default 4

tll tll

the IP Time To Live (TTL) value

Values 1 to 255

7.8.2.2 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 144: Service egress field descriptions](#) describes the fields.

Output example

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 144: Service egress field descriptions

Label	Description
Svc Id	The ID that identifies a service
Sdp Binding	The ID that identifies an SDP

Label	Description
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. When 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
- macsec**
MACsec related information about this service
- sap**
SAPs associated with 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 145: Service ID all field descriptions](#) describes the fields.

Output example

```
*A:7705:Dut-A# show service id 1 all
=====
Service Detailed Information
=====
Service Id       : 1
Service Type     : VPRN
Name            : XYZ Vprn 1
Description      : Default Description For VPRN ID 1
Customer Id      : 1
Creation Origin  : manual
Last Status Change: 07/23/2018 15:31:42
Last Mgmt Change : 07/23/2018 15:31:42
Admin State      : Up
Oper State       : Up
Route Dist.      : 10.20.1.1:1
VPRN Type        : regular
AS Number        : None
Router Id        : 10.20.1.1
ECMP              : Enabled
ECMP Max Routes  : 1
```

```

Max IPv4 Routes      : No Limit

Auto Bind Tunnel
Resolution           : filter
Filter Protocol      : sr-te

Max IPv6 Routes      : No Limit
Ignore NH Metric     : Disabled
Hash Label           : Disabled
Entropy Label        : Disabled
Vrf Target           : target:1:1
Vrf Import           : None
Vrf Export           : None
MVPN Vrf Target      : None
MVPN Vrf Import      : None
MVPN Vrf Export      : None
Label mode           : vrf
BGP VPN Backup       : Disabled

SAP Count            : 1                SDP Bind Count      : 0
IPT Count            : 0

-----
Service Destination Points(SDPs)
-----
No Matching Entries
-----
Service Access Points
-----

SAP 1/1/2:1
-----
Service Id           : 1
SAP                  : 1/1/2:1           Encap                : q-tag
Description          : sap-1-10.10.100.1
Admin State          : Up                Oper State           : Up
Flags                : None
Multi Svc Site       : None
Last Status Change   : 07/23/2018 15:31:47
Last Mgmt Change     : 07/23/2018 15:31:42
Sub Type             : regular
Dot1Q Ethertype      : 0x8100           QinQ Ethertype       : 0x8100
Split Horizon Group   : (Not Specified)
Admin MTU            : 1518             Oper MTU             : 1518
Ingr IP Fltr-Id      : n/a             Egr IP Fltr-Id       : n/a
Ingr Mac Fltr-Id     : n/a             Egr Mac Fltr-Id      : n/a
Ingr IPv6 Fltr-Id    : n/a             Egr IPv6 Fltr-Id     : n/a
qinq-pbit-marking    : both
Ing Scheduler Mode    : 4-priority       Egr Scheduler Mode   : 4-priority
Ing Agg Rate Limit    : max             Egr Agg Rate Limit   : max
Ing Agg cir           : 0               Egr Agg cir          : 0
Ing Shaper Group      : default         Egr Shaper Group     : default
Q Frame-Based Acct    : Disabled

Acct. Pol            : None             Collect Stats         : Disabled

-----
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 (Ingress)		
Dropped	: 0	0
Off. HiPrio	: 0	0
Off. LowPrio	: 938	380828
Off. Managed	: 0	0

Queueing Stats(Ingress QoS Policy 1)		
Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 938	380828

Queueing Stats(Egress QoS Policy 1)		
Dro. InProf	: 0	0
Dro. OutProf	: 0	0
For. InProf	: 0	0
For. OutProf	: 958	373620

Sap per Queue stats

	Packets	Octets
Ingress Queue 1 (Priority)		
Off. Combined	: 0	0
Dro. HiPrio	: 0	0
Dro. LowPrio	: 0	0
For. InProf	: 0	0
For. OutProf	: 938	380828

Ingress Queue 3 (Profile)		
Off. ColorIn	: 0	0
Off. ColorOut	: 0	0
Off. Uncolor	: 0	0
Dro. ColorOut	: 0	0
Dro. ColorIn/Uncolor	: 0	0
For. InProf	: 0	0
For. OutProf	: 0	0

Egress Queue 1		
For. InProf	: 0	0
For. OutProf	: 958	373620
Dro. InProf	: 0	0
Dro. OutProf	: 0	0

Service Interfaces

Interface

If Name	: ies-1-10.10.100.1		
Admin State	: Up	Oper (v4/v6)	: Up/Up
Protocols	: None		
IP Addr/mask	: 10.10.100.1/24	Address Type	: Primary
IGP Inhibit	: Disabled	Broadcast Address	: Host-ones
HoldUp-Time	: 0	Track Srrp Inst	: 0
IPv6 Address	: 2001:10:10:100::1/64		

```

IPv6 Addr State : PREFERRED
HoldUp-Time : 0
Link Lcl Address : fe80::be6b:4dff:fe52:b374/64
Link Lcl State : PREFERRED
Description : N/A
Ignore Port State : None
-----
Details
-----
Description : (Not Specified)
If Index : 2
Last Oper Chg : 07/23/2018 15:31:47
Srrp En Rtng : Disabled
SAP Id : 1/1/2:1
TOS Marking : Trusted
SNTP B.Cast : False
MAC Address : bc:6b:4d:52:b3:74
TCP MSS V4 : 0
ARP Timeout : 14400s
ARP Retry Timer : 5000ms
ARP Limit : Disabled
ARP Threshold : Disabled
ARP Limit Log Only : Disabled
IP Oper MTU : 1500
ARP Populate : Disabled
LdpSyncTimer : None
LSR Load Balance : system
Hashing Treatment : system
Use Ingress Port : system
TEID Load Balance : Disabled
SPI Load Balance : Disabled
L4 Load Balance : system
Reassem. Profile : none
uRPF Chk : disabled
Virt. If Index : 2
Global If Index : 257
Hold time : N/A
If Type : VPRN
TCP MSS V6 : 0
IPv6 Nbr ReachTime : 30s
IPv6 stale time : 14400s
IPv6 Nbr Limit : Disabled
IPv6 Nbr Threshold : Disabled
IPv6 Nbr Log Only : Disabled
Host Conn Verify : Disabled
-----
Statistics
-----
Rx Pkts : 958
Rx V4 Pkts : 469
Rx V4 Discard Pkts : 0
Inv Hdr CRC Pkts : 0
Inv Length Pkts : 0
Inv GRE Protocol* : 0
Dest Unreach Pkts : 0
Inv Mcast Addr P* : 0
Directed Bcast P* : 0
Src Martian Addr* : 0
Dest Martian Addr* : 0
Black Hole Pkts : 0
FltrActionDrop P* : 0
FltrNHUnreach Pk* : 0
FltrNHNotDirect * : 0
TTL Expired Pkts : 0
Slowpath Pkts : 0
MTU Exceeded Pkts : 0
Queue Pkts : 0
EncryptionDrop P* : 0
Last Tunnel : (Not Specified)
Other Discards P* : 0
Rx V6 Pkts : 489
Rx V6 Discard Pkts : 8
Inv Length Pkts : 0
Dest Unreach Pkts : 0
Inv Mcast Addr P* : 8
Rx Bytes : 388400
Rx V4 Bytes : 193228
Rx V4 Discard Byt* : 0
Inv Hdr CRC Bytes : 0
Inv Length Bytes : 0
Inv GRE Protocol* : 0
Dest Unreach Byt* : 0
Inv Mcast Addr B* : 0
Directed Bcast B* : 0
Src Martian Addr* : 0
Dest Martian Addr* : 0
Black Hole Bytes : 0
FltrActionDrop B* : 0
FltrNHUnreach Byt* : 0
FltrNHNotDirect * : 0
TTL Expired Bytes : 0
Slowpath Bytes : 0
MTU Exceeded Byt* : 0
Queue Bytes : 0
EncryptionDrop B* : 0
Other Discards B* : 0
Rx V6 Bytes : 195172
Rx V6 Discard Byt* : 688
Inv Length Bytes : 0
Dest Unreach Byt* : 0
Inv Mcast Addr B* : 688

```

Src Martian Addr*: 0	Src Martian Addr*: 0
Dest Martian Addr*: 0	Dest Martian Addr*: 0
Black Hole Pkts : 0	Black Hole Bytes : 0
FltrActionDrop P*: 0	FltrActionDrop B*: 0
TTL Expired Pkts : 0	TTL Expired Bytes: 0
Slowpath Pkts : 0	Slowpath Bytes : 0
MTU Exceeded Pkts: 0	MTU Exceeded Byt*: 0
Queue Pkts : 0	Queue Bytes : 0
EncryptionDrop P*: 0	EncryptionDrop B*: 0
Last Tunnel : (Not Specified)	
Other Discards P*: 0	Other Discards B*: 0
Tx Pkts : 964	Tx Bytes : 395212
Tx V4 Pkts : 479	Tx V4 Bytes : 197348
Tx V4 Discard Pkts: 0	Tx V4 Discard Byt*: 0
FltrActionDrop P*: 0	FltrActionDrop B*: 0
MTU Exceeded Pkts: 0	MTU Exceeded Byt*: 0
Queue Pkts : 0	Queue Bytes : 0
EncryptionDrop P*: 0	EncryptionDrop B*: 0
Last Tunnel : (Not Specified)	
Other Discards P*: 0	Other Discards B*: 0
Tx V6 Pkts : 485	Tx V6 Bytes : 197864
Tx V6 Discard Pkts: 0	Tx V6 Discard Byt*: 0
FltrActionDrop P*: 0	FltrActionDrop B*: 0
MTU Exceeded Pkts: 0	MTU Exceeded Byt*: 0
Queue Pkts : 0	Queue Bytes : 0
EncryptionDrop P*: 0	EncryptionDrop B*: 0
Last Tunnel : (Not Specified)	
Other Discards P*: 0	Other Discards B*: 0
 Security Details	
Admin Zone : None	Oper Zone : None
Bypass : No	
Rx V4 Discard Pkts: 0	Rx V4 Discard Byt*: 0
Unsup Proto Pkts : 0	Unsup Proto Bytes: 0
Unsup Svc Pkts : 0	Unsup Svc Bytes : 0
Unsup ICMP Type *: 0	Unsup ICMP Type *: 0
Fragment Pkts : 0	Fragment Bytes : 0
No Session Pkts : 0	No Session Bytes : 0
NAT Rte Loop Pkts: 0	NAT Rte Loop Byt*: 0
Other Discards P*: 0	Other Discards B*: 0
 Proxy ARP Details	
Rem Proxy ARP : Disabled	Local Proxy ARP : Disabled
Policies : none	
 Proxy Neighbor Discovery Details	
Local Pxy ND : Disabled	
Policies : none	
 DHCP no local server	
 DHCP Details	
Description : (Not Specified)	
Admin State : Down	Lease Populate : 0
Action : Keep	Trusted : Disabled
 DHCP6 Relay Details	
Description : (Not Specified)	
Admin State : Down	Lease Populate : 0
Oper State : Down	Nbr Resolution : Disabled
If-Id Option : None	Remote Id : Disabled
Src Addr : Not configured	
 DHCP6 Server Details	

Admin State : Down Max. Lease States : 8000

ICMP Details

Unreachables	: Number - 100	Time (seconds)	- 10
TTL Expired	: Number - 100	Time (seconds)	- 10
Parameter Problem	: Number - 100	Time (seconds)	- 10
ICMP Mask Reply	: True		

ICMPv6 Details

Packet Too Big	: Number - 100	Time (seconds)	- 10
Parameter Problem	: Number - 100	Time (seconds)	- 10
Redirects	: Disabled		
Time Exceeded	: Number - 100	Time (seconds)	- 10
Unreachables	: 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

Admin Groups

No Matching Entries

Srlg Groups

No Matching Entries

Group Encryption

Inbound Keygroup *: N/A
 Outbound Keygroup*: N/A

IP Transports

No Matching Entries

=====

* indicates that the corresponding row element may have been truncated.

*A:7705:Dut-A# show service id 1000 all

Service Detailed Information

=====

Service Id	: 1000		
Service Type	: Epipe		
Name	: XYZ Epipe 1000		
Description	: Default epipe description for service id 1000		
Customer Id	: 1	Creation Origin	: manual
Last Status Change	: 07/23/2018 18:46:02		
Last Mgmt Change	: 07/23/2018 18:44:07		
Admin State	: Up	Oper State	: Up
MTU	: 1514		
Vc Switching	: False		
SAP Count	: 1	SDP Bind Count	: 1
Per Svc Hashing	: Disabled		
TEID Hashing	: Disabled	L4 Hashing	: Disabled
Force QTag Fwd	: Disabled		

```

-----
BGP Information
-----

-----
Service Destination Points(SDPs)
-----

Sdp Id 1000:1000  -(10.20.1.6)
-----
Description      : Default sdp description
SDP Id           : 1000:1000                Type           : Spoke
VC Type          : Ether                    VC Tag           : n/a
Admin Path MTU   : 0                       Oper Path MTU    : 1546
Delivery         : MPLS
Far End          : 10.20.1.6
Tunnel Far End   : n/a                     LSP Types        : SR-ISIS
Entropy Label    : Disabled

Admin State      : Up                      Oper State       : Up
MinReqd SdpOperMTU : 1514
Acct. Pol       : None                    Collect Stats    : Disabled
Ingress Label   : 131052                  Egress Label     : 131062
Ingr Mac Fltr-Id : n/a                    Egr Mac Fltr-Id  : n/a
Ingr IP Fltr-Id  : n/a                    Egr IP Fltr-Id   : n/a
Ingr IPv6 Fltr-Id : n/a                   Egr IPv6 Fltr-Id : n/a
Admin ControlWord : Not Preferred          Oper ControlWord  : False
Admin BW(Kbps)   : 0                      Oper BW(Kbps)     : 0
Last Status Change : 07/23/2018 18:46:02   Signaling        : TLDP
Last Mgmt Change  : 07/23/2018 18:44:07
Endpoint         : N/A                    Precedence       : 4
PW Status Sig    : Enabled
Force Vlan-Vc    : Disabled
Class Fwding State : Down
Flags            : None
Local Pw Bits    : None
Peer Pw Bits     : None
Peer Fault Ip    : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : mplsRouterAlertLabel

Application Profile: None
Standby Sig Slave  : False

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

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

Statistics          :
I. Fwd. Pkts.       : 938                  I. Dro. Pkts.      : 0
I. Fwd. Octs.       : 399588                I. Dro. Octs.      : n/a
E. Fwd. Pkts.       : 938                  E. Fwd. Octets     : 410844
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

```

```

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

-----
Segment Routing
-----
ISIS                : enabled                LSP Id                : 524327
Oper Instance Id    : 0
OSPF                : disabled
TE-LSP              : disabled

-----
Number of SDPs : 1
-----
* indicates that the corresponding row element may have been truncated.
-----
Service Access Points
-----

-----
SAP 1/1/2:1000
-----
Service Id          : 1000
SAP                 : 1/1/2:1000              Encap                  : q-tag
Description         : Default sap description for service id 1000
Admin State         : Up                      Oper State             : Up
Flags               : None
Multi Svc Site      : None
Last Status Change  : 07/23/2018 18:44:13
Last Mgmt Change    : 07/23/2018 18:44:07
Sub Type            : regular
Dot1Q Ethertype     : 0x8100                 QinQ Ethertype         : 0x8100
Split Horizon Group: (Not Specified)

Hold Meps Up        : Disabled
Admin MTU           : 1518
Ingr IP Fltr-Id     : n/a                    Oper MTU               : 1518
Ingr Mac Fltr-Id    : n/a                    Egr IP Fltr-Id         : n/a
Ingr IPv6 Fltr-Id   : n/a                    Egr Mac Fltr-Id        : n/a
Ingr IPv6 Fltr-Id   : n/a                    Egr IPv6 Fltr-Id       : n/a
qinq-pbit-marking   : both
Ing Scheduler Mode   : 4-priority              Egr Scheduler Mode     : 4-priority
Ing Agg Rate Limit   : max                    Egr Agg Rate Limit     : max
Ing Agg cir          : 0                      Egr Agg cir            : 0
Ing Shaper Group     : default                 Egr Shaper Group       : default
Endpoint            : N/A
Q Frame-Based Acct   : Disabled
Vlan-translation    : None

Acct. Pol           : None                    Collect Stats           : Disabled

Ignore Oper Down     : Disabled

Loopback            : None
Swap Mac Addr       : Disabled                Loopback Time Left: unspecified

-----
ETH-CFM SAP specifics
-----
Hold Meps Up        : Disabled
-----
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 pncy      : n/a                     Multipoint shared  : Disabled
-----

Sap Statistics
-----
Last Cleared Time   : N/A

                Packets                      Octets
Forwarding Engine Stats (Ingress)
Dropped            : 0                      0
Off. HiPrio        : 0                      0
Off. LowPrio       : 938                    386456
Off. Managed       : 0                      0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio        : 0                      0
Dro. LowPrio       : 0                      0
For. InProf        : 0                      0
For. OutProf       : 938                    386456

Forwarding Engine Stats (Egress)
Dropped            : 0                      n/a

Queueing Stats(Egress QoS Policy 1)
Dro. InProf        : 0                      0
Dro. OutProf       : 0                      0
For. InProf        : 0                      0
For. OutProf       : 938                    386456
-----

Sap per Queue stats
-----
                Packets                      Octets
Ingress Queue 1 (Priority)
Off. Combined      : 0                      0
Dro. HiPrio        : 0                      0
Dro. LowPrio       : 0                      0
For. InProf        : 0                      0
For. OutProf       : 938                    386456

Egress Queue 1
For. InProf        : 0                      0
For. OutProf       : 938                    386456
Dro. InProf        : 0                      0
Dro. OutProf       : 0                      0
-----

Service Endpoints
-----
No Endpoints found.
-----

=====
VLL Sites
=====
Site           Site-Id  Dest           Admin          Oper   Fwdr
-----
No Matching Entries
=====
*A:7705:Dut-A#

```

Table 145: Service ID all field descriptions

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
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

Label	Description
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
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
Tunnel Far End	n/a
LSP Types	Indicates the supported LSP types: R = RSVP, L = LDP, B = BGP, I = SR-ISIS, O = SR-OSPF, T = SR-TE, n/a = not applicable
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

Label	Description
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
Class Fwding State	Not applicable
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdmin Down, SapAdminDown, InterfaceAdminDown, Port OperDown, PortMTUTooSmall, L2OperDown, Sap IngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdmin Down, NoSapIpPipeCelpAddr, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngressNamedPool Mismatch, SapEgressNamedPoolMismatch, NoSap EpipeRingNode
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

Label	Description
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
Associated LSP LIST	<p>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.</p> <p>If the SDP type is GRE, the following message displays: SDP delivery mechanism is not MPLS</p>
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: ServiceAdmin Down, SapAdminDown, InterfaceAdminDown, Port OperDown, PortMTUTooSmall, L2OperDown, Sap IngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdmin Down, NoSapIpipeCelpAddr, TodResourceUnavail,

Label	Description
	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
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
Ing Scheduler Mode	Indicates the ingress scheduler mode for the SAP
Egr Scheduler Mode	Indicates the egress scheduler mode for the SAP
Ing Agg Rate Limit	Indicates the PIR rate limit in the access ingress direction for the aggregate of the SAP queues

Label	Description
Egr Agg Rate Limit	Indicates the PIR rate limit in the access egress direction for the aggregate of the SAP queues
Ing Agg cir	Indicates the CIR rate limit in the access ingress direction for the aggregate of the SAP queues
Egr Agg cir	Indicates the CIR rate limit in the access egress direction for the aggregate of the SAP queues
Ing Shaper Group	Indicates the ingress shaper group for the SAP
Egr Shaper Group	Indicates the egress shaper group for the SAP
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
Segment Routing	
ISIS	Indicates the state of segment routing for IS-IS: enabled or disabled
LSP Id	The LSP identifier
Oper Instance Id	The IS-IS instance identifier for the SR IS-IS instance
OSPF	Indicates the state of segment routing for OSPF: enabled or disabled
TE-LSP	Indicates the state of segment routing for TE LSP: enabled or disabled
TWAMP-Light Reflector	
Admin State	Displays one of the following:

Label	Description
	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 (Ingress)	
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

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 (Priority)	The index of the ingress QoS queue of this SAP, where n is the index number
Off. Combined	The combined total number of high-priority and low-priority packets or octets offered to the forwarding engine
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
Ingress Queue n (Profile)	The index of the ingress QoS queue of this SAP, where n is the index number
Off. ColorIn	The number of packets or octets colored as in-profile for the SAP (offered)
Off. ColorOut	The number of packets or octets colored as out-of-profile for the SAP (offered)

Label	Description
Off. Uncolor	The number of packets or octets that are unprofiled for the SAP (offered)
Dro. ColorOut	The number of packets or octets colored as out-of-profile that were dropped for the SAP
Dro. ColorIn/Uncolor	The number of packets or octets that were colored as in-profile or were unprofiled that were dropped for the SAP
For. InProf	The number of forwarded packets or octets colored as in-profile (FC profile set to "in" or "no profile" and rate less than or equal to CIR)
For. OutProf	The number of forwarded packets or octets that were colored as out-of-profile (FC profile set to "out" or "no profile" and rate above CIR)
Egress Queue <i>n</i>	The index of the egress QoS queue of the SAP, where <i>n</i> 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
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	

Label	Description
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 statements 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):

Label	Description
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

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 44: SAP ID configurations](#) 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 146: Service ID ARP field descriptions](#) describes the fields.

Output example

```
*A:ALU-12>show>service>id# arp
=====
ARP Table
=====
IP Address      MAC Address      Type    Expiry    Interface      SAP
-----
192.168.86.45   98-90-96-B9-9D-61 Other    00:00:00   ies-100-190.11.1 1/1/11:0
=====
*A:ALU-12>show>service>id#
```

Table 146: Service ID ARP field descriptions

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
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 147: Service ID base field descriptions](#) describes the fields.

Output example

```
*A:7705:Dut-C>config>router>mpls# show service id 3 base
=====
Service Basic Information
=====
Service Id :          3
Service Type :        VPRN
Name :              XYZ Vprn 3
Description :        Default Description For VPRN ID 3
Customer Id :        1          Creation Origin : manual
Last Status Change: 05/17/2021 18:48:14
Last Mgmt Change : 05/17/2021 18:48:14
Admin State :
Router Oper State : Up          Up Oper State : Up
Route Dist. : 10.10.10.2:3      VPRN Type : regular
AS Number : 65000              Router Id : 10.10.10.2
ECMP : Enabled                 ECMP Max Routes : 8
Max IPv4 Routes : No Limit
Auto Bind Tunnel
Resolution : filter
Filter Protocol : rsvp
Weighted ECMP : Enabled        ECMP Max Routes : 8
Max IPv6 Routes : No Limit
Ignore NH Metric : Disabled
Hash Label : Disabled
Entropy Label : Disabled
Vrf Target : target:65000:1
Vrf Import : None
Vrf Export : None
MVPN Vrf Target : None
MVPN Vrf Import : None
MVPN Vrf Export : None
Label mode : vrf
BGP VPN Backup : Disabled

SAP Count : 1          SDP Bind Count : 0
IPT Count : 0

-----
Service Access & Destination Points
-----
Identifier                Type    AdmMTU    OprMTU    Adm    Opr
-----
sap:1/1/2                 null    1572     1572     Up     Up
-----
IP Transports
-----
IptId    Adm    Opr
```

 No Matching Entries
 =====

Table 147: Service ID base field descriptions

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
Creation Origin	The method used to create this service, either manual or automatic
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
Up Oper State	The operating state of the service
Router Oper State	The operating state of the router
Route Dist.	The largest frame size (in octets) that this 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	Specifies whether equal cost multipath is enabled or disabled for this service
ECMP Max Routes	The maximum number of ECMP routes that can be received from the neighbors in the group or for the specific neighbor
Max IPv4 Routes	The maximum number of IPv4 routes that can be used for path sharing
Auto Bind Tunnel Resolution	The autobind resolution mode: any, filter, or disabled

Label	Description
Filter Protocol	The autobind filter protocol
Weighted ECMP	Specifies whether autobind weighted ECMP option is enabled or disabled
ECMP Max Routes	The maximum number of weighted ECMP routes that can be received from the neighbors in the group or for the specific neighbor
Max IPv6 Routes	The maximum number of IPv6 routes that can be used for path sharing
Ignore NH Metric	Specifies whether the ignore next hop metric option is enabled or disabled
Hash Label	Specifies whether the hash label option is enabled or disabled
Entropy Label	Specifies whether the entropy label option is enabled or disabled
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
MVPN Vrf Target	The MVPN route target in the VRF applied to this service
MVPN Vrf Import	The MVPN VRF import policy applied to this service
MVPN Vrf Export	The MVPN VRF export policy applied to this service
Label mode	The label mode for this service
BGP VPN Backup	Specifies whether the BGP VPN backup option is enabled or disabled
SAP Count	The number of SAPs defined on this service
SDP Bind Count	The number of SDPs bound to this service
IPT Count	The number of IP transport subservices associated with this service
Service Access & Destination Points	
Identifier	The service access (SAP) and destination (SDP) points
Type	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP

Label	Description
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
IP Transports	
IptId	The IP transport subservice ID
Adm	The administrative state of the IP transport subservice
Opr	The operating state of the IP transport subservice

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

Output

The following output is an example of server ID DHCP statistics information, and [Table 148: Service ID DHCP statistics field descriptions](#) describes the fields.

Output example

```
*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 148: Service ID DHCP statistics field descriptions

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

Label	Description
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 149: Service ID DHCP summary field descriptions](#) describes the fields.

Output example

```
*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 149: Service ID DHCP summary field descriptions

Label	Description
DHCP Summary, service x	

Label	Description
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 with the service are displayed.

Parameters

ip-address

the IPv4 or IPv6 address of the interface for which to display information

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 150: Service ID interface detailed field descriptions](#) describes the fields.

Output example

```
*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
Protocols     : None
Oper (v4/v6)  : Down/Down

IP Addr/mask  : Not Assigned
-----
Details
-----
If Index       : 2
Last Oper Chg  : 02/03/2010 21:59:02
SAP Id         : 1/5/2
Virt. If Index : 2
Global If Index : 125
```

```

TOS Marking : Trusted
SNTP B.Cast : False
MAC Address :
IP MTU : 1500
Arp Populate : Disabled
LdpSyncTimer : None

If Type : VPRN
Arp Timeout : 14400
ICMP Mask Reply : True
Host Conn Verify : Disabled

Proxy ARP Details
Rem Proxy ARP: Disabled
Policies : none
Local Proxy ARP : Disabled

Proxy Neighbor Discovery Details
Local Pxy ND : Disabled
Policies : none

DHCP Details
Admin State : Down
Action : Keep
Lease Populate : 0
Trusted : Disabled

DHCP6 Relay Details
Admin State : Down
Oper State : Down
If-Id Option : None
Src Addr : Not configured
Lease Populate : 0
Nbr Resolution : Disabled
Remote Id : Disabled

DHCP6 Server Details
Admin State : Down
Max. Lease States: 8000

ICMP Details
Redirects : Number - 100
Unreachables : Number - 100
TTL Expired : Number - 100
Time (seconds) - 10
Time (seconds) - 10
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 150: Service ID interface detailed field descriptions

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	

Label	Description
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 statements applied to proxy ARP
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

Label	Description
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

ip-transport

Syntax

ip-transport *ipt-id* [**detail** | **statistics**]

Context

show>service>id

Description

This command displays information for a specified IP transport subservice within this VPRN service. If no IP transport subservice is specified, summary information is displayed for all IP transport subservices associated with the VPRN service.

Parameters

ipt-id

the physical port associated with the IP transport subservice, in the format *slot/mda/port.channel*

detail

displays detailed information for the specified IP transport subservice

statistics

displays statistical information for the specified IP transport subservice

Output

The following output is an example of IP transport subservice summary information for a specified service, and [Table 151: Service IP transport subservice summary field descriptions](#) describes the fields.

Output example

```
*A:ALU-12# show service id 100 ip-transport
=====
IP Transport (Summary), Service 100
=====
IptId      LocalIP      LocalPort Proto RemHost DSCP FC FltrUnkn Adm  Opr
-----
1/2/4.1    192.168.1.1    3000    tcp    2      ef    ef disabled Up   Up
-----
Entries found: 1
=====
*A:ALU-12#
```

Table 151: Service IP transport subservice summary field descriptions

Label	Description
IP Transport (Summary), Service x	
IptId	The IP transport subservice physical port identifier
LocalIP	The IP address (IPv4) that is used for the local host
LocalPort	The port number that is used by remote hosts to establish TCP/UDP sessions to the local host
Proto	The protocol type that is used for all sessions to and from the local host (either TCP or UDP)
RemHost	The number of remote hosts associated with the IP transport subservice
DSCP	The DSCP name used to mark the DSCP field in IP transport packets
FC	The FC name used for IP transport packets
FltrUnkn	Indicates whether the filter-unknown-host command is enabled or disabled on the IP transport subservice
Adm	The administrative state of the IP transport subservice
Opr	The operational state of the IP transport subservice
Entries found:	The number of IP transport subservices associated with this service

The following output is an example of detailed information for a specified IP transport subservice within a specified service, and [Table 152: Service IP transport subservice detailed field descriptions](#) describes the fields.

Output example

```

*A:7705:Dut-C# show service id 100 ip-transport 1/2/4.1 detail
=====
IP Transport
=====
Service Id       : 100 (VPRN)
IP Transport Id  : 1/2/4.1
Description      : (Not Specified)
Admin State     : Up
Oper Flags       : (Not Specified)
Local IP Address : 192.168.1.1
Local IP Protocol : tcp
DSCP             : ef
FC              : ef
TCP Inact Timeout : 30
TCP Max Retries  : 5
TCP Retry Interval : 5
Num Remote Hosts : 0
Last Mgmt Change : 12/07/2016 16:48:22
Last Oper Change : 12/07/2016 16:48:22
=====
Oper State       : Up
Local Port Number : 3000
Filter Unknown Host : enabled
Profile          : in
=====
IP Transport Accumulated Statistics
-----
Known Remote Hosts
Packets sent           : 44
Characters sent        : 66000
Packets received       : 67
Characters received    : 51114
Connections           : 2
  To                   : 2
  From                 : 0
Connection retries     : 20
Connection failures    : 2
Currently connected    : 0
Unknown Remote Hosts
Packets sent           : 119
Characters sent        : 178500
Packets received       : 153
Characters received    : 116039
Successful connections from : 2
Rejected due to unknown host filter : 37
Rejected due to out of resources : 0
Inactivity timeouts    : 0
Last RemIp:RemPort    : 192.168.1.7:4001
Currently connected    : 0
Dropped packets due to no remote hosts : 27
=====
*A:7705:Dut-C#

```

Table 152: Service IP transport subservice detailed field descriptions

Label	Description
IP Transport	
Service Id	The ID that identifies the service (the service type is shown in brackets)
IP Transport Id	The physical port identifier for this IP transport subservice

Label	Description
Description	The description associated with this IP transport subservice
Admin State	The administrative state of this IP transport subservice
Oper State	The operational state of this IP transport subservice
Oper Flags	The operational flags associated with this IP transport subservice
Local IP Address	The IP address (IPv4) that is used for the local host
Local Port Number	The port number that is used by remote hosts to establish TCP/UDP sessions to the local host
Local IP Protocol	The protocol type that is used for all sessions to and from the local host (either TCP or UDP)
DSCP	The DSCP name used to mark the DSCP field in IP transport packets
Filter Unknown Host	Indicates whether the filter-unknown-host command is enabled or disabled for this IP transport subservice
FC	The FC name used for IP transport packets
Profile	The profile marking for the IP transport packets (in or out)
TCP Inact Timeout	The configured inactivity timeout value for TCP connections
TCP Max Retries	The configured maximum retry value for TCP connections
TCP Retry Interval	The configured retry interval value for TCP connections
Num Remote Hosts	The number of remote hosts associated with this IP transport subservice
Last Mgmt Change	The date and time of the most recent management-initiated change to this IP transport subservice
Last Oper Change	The date and time of the most recent operational status change for this IP transport subservice
IP Transport Accumulated Statistics	
Known Remote Hosts	
Packets sent	The number of packets sent to the host
Characters sent	The number of data characters sent to the host
Packets received	The number of packets received from the host
Characters received	The number of data characters received from the host

Label	Description
Connections To From	The number of connections to and from the host
Connection retries	The number of connection retries to the host
Connection failures	The number of connection failures to the host
Currently connected	The number of hosts currently connected
Unknown Remote Hosts	
Packets sent	The number of packets sent to the host
Characters sent	The number of data characters sent to the host
Packets received	The number of packets received from the host
Characters received	The number of data characters received from the host
Successful connections from	The number of successful connections from the host
Rejected due to unknown host filter	The number of rejected connection attempts from the host due to the filter-unknown-host command being enabled
Rejected due to out of resource	The number of connection attempts from the host that were rejected due to the unavailability of resources
Inactivity timeouts	The number of connections from the host that timed out due to inactivity
Last RemIp:RemPort	The IP address (IPv4) and port number used by the host for the last connection
Currently connected	The number of hosts that are currently connected
Dropped packets due to no remote hosts	The number of packets dropped due to no hosts being connected

remote-host

Syntax

remote-host *host-id* [**detail** | **statistics**]

Context

show>service>id>ip-transport

Description

This command displays information for a specified remote host within this IP transport subservice within this service. If no remote host is specified, summary information is displayed for all remote hosts within this IP transport subservice.

Parameters

- host-id*

the remote host identifier

Values1 to 2147483647or a name string up to 64 characters long
- detail**

displays detailed information for a specified remote host
- statistics**

displays summary information for a specified remote host

Output

The following output is an example of IP transport subservice remote host summary information when no remote host is specified, and [Table 153: IP transport subservice remote host summary field descriptions](#) describes the fields.

Output example

```
*A:7705:Dut-C# show service id 100 ip-transport 1/6/4.1 remote-host
=====
IPT Remote Host (Summary), Service 100 IPT 1/6/4.1
=====
RemId      RemIp:RemPort      Rcvd Chars  Sent Chars  Drop Chars  State
      Rcvd Pkts  Sent Pkts  Drop Pkts  Up Time
-----
1          192.168.1.1:3000      2555        2044        0          connected
                        5            4            0          00h01m21s
(unknown)  192.168.1.7:4000      0           2044       5110        connected
                        0            4            10         00h00m42s
-----
Number of known remote hosts: 1
Number of unknown remote hosts: 1
Total entries found: 2
=====
*A:7705:Dut-C#
```

Table 153: IP transport subservice remote host summary field descriptions

Label	Description
IP Remote Host (Summary), Service x IPT x/x/x.x	
RemId	The remote host identifier
RemIp:RemPort	The IP address (IPv4) and port number used by the remote host
Rcvd Chars	The number of data characters received from the remote host

Label	Description
Sent Chars	The number of data characters sent to the remote host
Drop Chars	The number of data characters destined for the remote host that were dropped
State	The operational state of the packet transport session connection to the remote host
Rcvd Pkts	The number of packets received from the remote host
Sent Pkts	The number of packets sent to the remote host
Drop Pkts	The number of packets destined for the remote host that were dropped
Up Time	The amount of time that the remote host has been connected
Number of known remote hosts	The number of known remote hosts associated with the IP transport subservice
Number of unknown remote hosts	The number of unknown remote hosts associated with the IP transport subservice
Total entries found	The total number of hosts associated with the IP transport subservice

The following output is an example of IP transport subservice detailed information for a specified remote host, and [Table 154: IP transport subservice remote host detailed field descriptions](#) describes the fields.

Output example

```
*A:7705:Dut-C# show service id 100 ip-transport 1/2/4.1 remote-host 1 detail
=====
IPT Remote Host
=====
Service Id       : 100 (VPRN)
IP Transport Id  : 1/2/4.1
Remote Host Id   : 1
Name             : (Not Specified)
Description      : (Not Specified)
IP Address       : 192.168.1.6           Port Number      : 4000
Last Mgmt Change : 12/07/2016 16:48:44    Up Time          : 00h01m44s
Session State    : connected
Last Connect     : successful
-----
IPT Remote Host Statistics
-----
Sent Pkts       : 134           Sent Chars       : 201000
Dropped Pkts    : 0            Dropped Chars    : 0
Rcvd Pkts       : 267          Rcvd Chars       : 201000
Session information
  Connections    : 2
  To             : 1
  From          : 1
  Connection retries : 0
  Connection failures : 0
  Closed by far end : 1
```

```

Inactivity timeouts          : 0
=====
*A:7705:Dut-C#

```

Table 154: IP transport subservice remote host detailed field descriptions

Label	Description
IP Remote Host	
Service Id	The ID that identifies the service (the service type is shown in brackets)
IP Transport Id	The physical port identifier for the IP transport subservice
Remote host Id	The host identifier associated with this remote host
Name	The name associated with this remote host
Description	The description associated with this remote host
IP Address	The IP address associated with this remote host
Port Number	The port number associated with this remote host
Last Mgmt Change	The date and time of the most recent management-initiated change to this remote host
Session State	The operational state of the packet transport session to this host
Up Time	The amount of time that this remote host has been connected
Last Connect	Indicates whether the last connection attempt to this remote host was successful or unsuccessful
IP Remote Host Statistics	
Sent Pkts	The number of packets sent to this remote host
Sent Chars	The number of data characters sent to this remote host
Dropped Pkts	The number of packets destined for this remote host that were dropped
Dropped Chars	The number of data characters destined for this remote host that were dropped
Rcvd Pkts	The number of packets received from this remote host
Rcvd Chars	The number of data characters received from this remote host
Session information	
Connections	

Label	Description
To	The number of connections made to this host
From	The number of connections made from the host
Connection retries	The number of connection retries to this host
Connection failures	The number of connection failures to this host
Closed by far end	The number of connections closed by the far end
Inactivity timeouts	The number of connection that were timed out due to inactivity

macsec

Syntax

macsec

Context

show>service>id

Description

This command displays MACsec security information for the specified service.

Output

The following output is an example of MACsec information, and [Table 155: Service-ID MACsec field descriptions](#) describes the fields.

Output example

```
*A:ALU-12# show service id 1 macsec
=====
MACsec (Summary), Service 1
=====
SAP          MACsec    MACsec    Encap    CA tags   CA-name
              port      sub-port  match    in-clear
-----
1/1/3        1/1/3      1         all      0         cal
=====
*A:ALU-12#
```

Table 155: Service-ID MACsec field descriptions

Label	Description
SAP	The service SAP
MACsec port	The port enabled for MACsec

Label	Description
MACsec sub-port	The subport enabled for MACsec
Encap match	The traffic encapsulation type to match: all traffic, untagged-only traffic, single-tag or dot1q traffic, double-tag or QinQ traffic
CA tags in-clear	The number of tags in clear text for this CA
CA-name	The name of the MACsec connectivity association for this SAP

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 44: SAP ID configurations](#) 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 156: Service ID SAP detailed field descriptions](#) describes the fields.

Output example

```
*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
Ing Agg Rate Limit : max          Egr Agg Rate Limit: 7000
Ing Agg cir      : 0              Egr Agg cir     : 700
Ing Shaper Group  : default       Egr Shaper Group : default
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 (Ingress)
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
Admin State     : Up
Encap           : null
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 (Ingress)		
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
Admin State     : Up
Encap           : null
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 156: Service ID SAP detailed field descriptions

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, SapAdmin Down, InterfaceAdminDown, PortOperDown, PortMTUToo Small, L2OperDown, SapIngressQoSMismatch, SapEgress QoSMismatch, RelearnLimitExceeded, RxProtSrcMac, Parent IfAdminDown, NoSapIpAddr, TodResourceUnavail, Tod MssResourceUnavail, SapParamMismatch, CemSapNoEcid OrMacAddr, StandByForMcRing, ServiceMTUTooSmall, Sap IngressNamedPoolMismatch, 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

Label	Description
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
Ingr 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
Ingr 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
Ingr 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
Ingr Shaper Group	The ingress shaper group for the SAP
Egr Shaper Group	The egress shaper group for the SAP
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 (Ingress)	

Label	Description
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 (Priority)	The index of the ingress QoS queue of this SAP, where n is the index number
Off. Combined	The combined total number of high-priority and low-priority packets or octets offered to the forwarding engine
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

Label	Description
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
Ingress Queue <i>n</i> (Profile)	The index of the ingress QoS queue of this SAP, where <i>n</i> is the index number
Off. ColorIn	The number of packets or octets colored as in-profile for the SAP (offered)
Off. ColorOut	The number of packets or octets colored as out-of-profile for the SAP (offered)
Off. Uncolor	The number of packets or octets that are unprofiled for the SAP (offered)
Dro. ColorOut	The number of packets or octets colored as out-of-profile that were dropped for the SAP
Dro. ColorIn/Uncolor	The number of packets or octets that were colored as in-profile or unprofiled that were dropped for the SAP
For. InProf	The number of forwarded packets or octets colored as in-profile (FC profile set to "in" or "no profile" and rate less than or equal to CIR)
For. OutProf	The number of forwarded packets or octets that were colored as out-of-profile (FC profile set to "out" or "no profile" and rate above CIR)
Egress Queue <i>n</i>	The index of the egress QoS queue of the SAP, where <i>n</i> 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

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 157: Service ingress label field descriptions](#) describes the fields.

Output example

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 157: Service ingress label field descriptions

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

ip-transport-using

Syntax

ip-transport-using [**ip-transport** *ipt-id*]

Context

show>service

Description

This command displays IP transport subservice information for a specified port. If no port is specified, the command displays a summary of all IP transport subservices defined for the VPRN service.

Parameters

ipt-id

the physical port associated with the IP transport subservice, in the format *slot/mda/port.channel*

Output

The following output is an example of **ip-transport-using** information, and [Table 158: IP transport-using field descriptions](#) describes the fields.

Output example

```
*A:ALU-48# show service ip-transport-using
=====
IP Transports
=====
IptId          SvcId      Type  Adm  Opr
-----
1/2/4.1        100        VPRN  Up   Up
-----
Entries found: 1
-----
*A:ALU-48#
```

Table 158: IP transport-using field descriptions

Label	Description
IP Transports	
IptId	The IP transport subservice physical port identifier
SvcId	The service identifier
Type	The type of service
Adm	The administrative state of the IP transport subservice
Opr	The operational state of the IP transport subservice
Entries found	The number of IP transport subservices using this service

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 159: Service ID SDP detailed field descriptions](#) describes the fields.

Output example

```
*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 159: Service ID SDP detailed field descriptions

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)

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	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 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: ServiceAdmin

Label	Description
	Down, SapAdminDown, InterfaceAdminDown, PortOper Down, PortMTUTooSmall, L2OperDown, SapIngress QoSMismatch, SapEgressQoSMismatch, RelearnLimit Exceeded, RxProtSrcMac, ParentIfAdminDown, NoSap IpipeCelpAddr, TodResourceUnavail, TodMssResource Unavail, SapParamMismatch, CemSapNoEcidOrMacAddr, StandByForMcRing, ServiceMTUTooSmall, SapIngress NamedPoolMismatch, SapEgressNamedPoolMismatch, No SapEpipeRingNode
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
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 160: TWAMP Light field descriptions](#) describes the fields.

Output example

```
*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 160: TWAMP Light field descriptions

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

Label	Description
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 161: Service service-using field descriptions](#) describes the fields.

Output example

```
*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 161: Service service-using field descriptions

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
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

7.8.2.3 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 162: Aggregate route field descriptions](#) describes the fields.

Output example

```
*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 162: Aggregate route field descriptions

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 163: ARP table field descriptions](#) describes the fields.

Output example

```
*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   0th      system
10.10.4.3       00:00:00:00:00:00 00h00m00s   0th      ALU-1-2
-----
No. of ARP Entries: 2
=====
*A:ALU-12#
```

Table 163: ARP table field descriptions

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

Label	Description
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**

damping [*ip-prefix*[/*ip-prefix-length*]] [*damp-type*] [**detail**] **mvpn-ipv4**

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 IPv4 or IPv6 address
- 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

mvpn-ipv4
displays dampened routes for the MVPN-IPv4 address family

Output

The following output is an example of BGP damping information, and [Table 164: BGP damping field descriptions](#) describes the fields.

Output example

```
*A:ALU-12# show router 6 bgp neighbor damping
=====
BGP Router ID:10.0.0.14      AS:65206      Local AS:65206
=====
Legend -
Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP Damped Routes
=====
Flag  Network          From          Reuse          AS-Path
-----
ud*i  10.149.7.0/24        10.0.28.1     00h00m00s      60203 65001 19855 3356
                        1239 22406
si    10.155.6.0/23      10.0.28.1     00h43m41s      60203 65001 19855 3356
                        2914 7459
si    10.155.8.0/22      10.0.28.1     00h38m31s      60203 65001 19855 3356
                        2914 7459
si    10.155.12.0/22     10.0.28.1     00h35m41s      60203 65001 19855 3356
                        2914 7459
si    10.155.22.0/23     10.0.28.1     00h35m41s      60203 65001 19855 3356
                        2914 7459
si    10.155.24.0/22     10.0.28.1     00h35m41s      60203 65001 19855 3356
                        2914 7459
si    10.155.28.0/22     10.0.28.1     00h34m31s      60203 65001 19855 3356
```

```

si    10.155.40.0/21    10.0.28.1    00h28m24s    2914 7459
60203 65001 19855 3356
7911 7459
si    10.155.48.0/20    10.0.28.1    00h28m24s    60203 65001 19855 3356
7911 7459
ud*i  10.8.140.0/24     10.0.28.1    00h00m00s    60203 65001 19855 3356
4637 17447
ud*i  10.8.141.0/24     10.0.28.1    00h00m00s    60203 65001 19855 3356
4637 17447
ud*i  10.9.0.0/18       10.0.28.1    00h00m00s    60203 65001 19855 3356
3561 9658 6163
. . .
ud*i  10.213.184.0/23   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 : 10.0.0.0      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 : 10.149.7.0/24
-----
Network      : 10.149.7.0/24      Peer      : 10.0.28.1
NextHop      : 10.0.28.1         Reuse time : 00h00m00s
Peer AS      : 60203             Peer Router-Id : 10.10.10.203
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 : 10.142.48.0/20
-----
Network      : 10.142.48.0/20    Peer      : 10.0.28.1
NextHop      : 10.0.28.1         Reuse time : 00h00m00s
Peer AS      : 60203             Peer Router-Id : 10.10.10.203
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 : 10.200.128.0/19
-----
Network      : 10.200.128.0/19   Peer      : 10.0.28.1
NextHop      : 10.0.28.1         Reuse time : 00h00m00s
Peer AS      : 60203             Peer Router-Id : 10.10.10.203
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 : 10.203.192.0/18
-----
```

```
Network      : 10.203.192.0/18      Peer      : 10.0.28.1
NextHop      : 10.0.28.1            Reuse time : 00h00m00s
Peer AS      : 60203                Peer Router-Id : 10.10.10.203
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 neighbor 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 : 10.142.48.0/20
-----
```

```
Network      : 10.142.48.0/20      Peer      : 10.0.28.1
NextHop      : 10.0.28.1            Reuse time : 00h29m22s
Peer AS      : 60203                Peer Router-Id : 10.10.10.203
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 : 10.200.128.0/19
-----
```

```
Network      : 10.200.128.0/19      Peer      : 10.0.28.1
NextHop      : 10.0.28.1            Reuse time : 00h29m22s
Peer AS      : 60203                Peer Router-Id : 10.10.10.203
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 : 10.203.240.0/20
-----
```

```
Network      : 10.203.240.0/20      Peer      : 10.0.28.1
NextHop      : 10.0.28.1            Reuse time : 00h29m22s
Peer AS      : 60203                Peer Router-Id : 10.10.10.203
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 : 10.206.0.0/17
-----
```

```

-----
Network      : 10.206.0.0/17      Peer      : 10.0.28.1
NextHop      : 10.0.28.1         Reuse time : 00h29m22s
Peer AS      : 60203             Peer Router-Id : 10.10.10.203
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 164: BGP damping field descriptions

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/Flags	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
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

Label	Description
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 165: BGP group field descriptions](#) describes the fields.

Output example

```
*A:ALU-12# show router 6 bgp neighbor 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
Min Hold Time        : 10
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
Creation Origin       : manual
Split Horizon        : Disabled

List of Peers
- 10.44.44.44 :

Total Peers          : 1                  Established       : 0
-----
Peer Groups : 1

```

*A:ALU-12# show router 6 bgp neighbor 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
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
Min Hold Time        : 10
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
Creation Origin       : manual
Split Horizon        : Disabled

List of Peers
- 10.44.44.44 :

Total Peers          : 1                  Established       : 0
-----
Peer Groups : 1
=====

```

*A: ALU-12

Table 165: BGP group field descriptions

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
Prefix Limit	No Limit: no route limit assigned to the BGP peer group

Label	Description
	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
Min Hold Time	The configured minimum hold-time 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

Label	Description
NLRI	The type of network layer reachability information that the specified peer group can accept Unicast: IPv4 unicast routing information can be carried
Preference	The configured route preference value for the peer group
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
Creation Origin	The creation method of the peer group
Split Horizon	The configured split-horizon setting
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 [detail]]
neighbor [as-number [detail]]
neighbor ip-address [family [type mvpn-type]] filter1 [brief]
neighbor ip-address [family] filter2
neighbor as-number [family] filter2
neighbor ip-address orf [filter3]
neighbor ip-address graceful-restart
neighbor [dynamic]

```

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 about 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 IPv4 or IPv6 address for which to display information

detail

displays detailed information

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 VPN-IPv4 routes
 label-ipv4 – displays only those BGP peers that have the label-IPv4 family enabled
 vpn-ipv4 – displays only those BGP peers that have the VPN-IPv4 family enabled
 mvpn-ipv4 – displays only those BGP peers that have the MVPN-IPv4 family enabled

type

displays information for the MVPN type

mvpn-type

the specified MVPN type for which to display information

Values intra-ad | inter-ad | spmsi-ad | leaf-ad | source-ad | shared-join | source-join

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

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



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.

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

dynamic

displays information for dynamic BGP neighbors

Output

The following outputs are examples of BGP neighbor information:

- BGP neighbor (standard, detailed, and dynamic) ([Output example - BGP neighbor \(standard, detailed, and dynamic\)](#), [Table 166: BGP neighbor \(standard, detailed, and dynamic\) field descriptions](#))
- BGP neighbor (advertised and received) ([Output example - BGP neighbor \(advertised-routes and received-routes\)](#), [Table 167: BGP neighbor \(advertised-routes and received-routes\) field descriptions](#))
- BGP neighbor (graceful restart) ([Output example - BGP neighbor \(graceful restart\)](#), [Table 168: BGP neighbor \(graceful restart\) field descriptions](#))

Output example - BGP neighbor (standard, detailed, and dynamic)

```
*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 Port       : 179
Local AS          : 65000           Local Port       : 179
Local Address     : 10.10.10.1      Dynamic Peer     : No
Peer Type         : Internal        Last State       : Established
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 Port       : 179
Local AS          : 65000           Local Port       : 179
Local Address     : 10.10.10.1      Dynamic Peer     : No
Peer Type         : Internal        Last State       : Established
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:7705_ALU-2>show>router>bgp# neighbor dynamic
```

```
=====
BGP Neighbor
=====

```

```
-----
Peer           : 10.100.1.3
Description    : (Not Specified)
Group          : dynamic

```

```
-----
Peer AS        : 65000      Peer Port      : 51374
Peer Address   : 10.100.1.3
Local AS       : 65000      Local Port     : 179
Local Address  : 10.100.1.2

```

```

Peer Type           : Internal      Dynamic Peer       : Yes
State               : Established   Last State          : Established
Last Event          : recvKeepAlive
Last Error           : Cease (Connection Collision Resolution)
Local Family         : IPv4 VPN-IPv4 IPv6 VPN-IPv6
Remote Family        : IPv4 VPN-IPv4 IPv6 VPN-IPv6
Hold Time            : 90           Keep Alive           : 30
Min Hold Time        : 0
Active Hold Time     : 90           Active Keep Alive    : 30
Cluster Id           : None
Preference           : 170          Num of Update Flaps  : 0
Recd. Paths          : 6
IPv4 Recd. Prefixes : 5            IPv4 Active Prefixes : 0
IPv4 Suppressed Pfxs : 0            VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs  : 0            VPN-IPv4 Active Pfxs : 0
IPv6 Suppressed Pfxs : 0
IPv6 Recd. Prefixes  : 2            IPv6 Active Prefixes : 0
VPN-IPv6 Recd. Pfxs  : 0            VPN-IPv6 Active Pfxs : 0
VPN-IPv6 Suppr. Pfxs : 0
MVPN-IPv4 Suppr. Pfxs : 0          MVPN-IPv4 Recd. Pfxs : 0
MVPN-IPv4 Active Pfxs : 0
Flow-IPv4 Suppr. Pfxs : 0          Flow-IPv4 Recd. Pfxs : 0
Flow-IPv4 Active Pfxs : 0          Rte-Tgt Suppr. Pfxs  : 0
Rte-Tgt Recd. Pfxs   : 0            Rte-Tgt Active Pfxs  : 0
Backup IPv4 Pfxs      : 0            Backup IPv6 Pfxs      : 0
Mc Vpn Ipv4 Suppr. P*: 0
Backup Vpn IPv4 Pfxs : 0            Backup Vpn IPv6 Pfxs : 0
Input Queue           : 0            Output Queue          : 0
i/p Messages          : 11           o/p Messages          : 9
i/p Octets             : 951          o/p Octets             : 445
i/p Updates            : 6            o/p Updates            : 4
Evpn Suppr. Pfxs      : 0            Evpn Recd. Pfxs       : 0
Evpn Active Pfxs      : 0
MS-PW Suppr. Pfxs     : 0            MS-PW Recd. Pfxs      : 0
MS-PW Active Pfxs     : 0
TTL Security          : Disabled      Min TTL Value          : n/a
Graceful Restart       : Disabled      Stale Routes Time      : n/a
Advertise Inactive     : Disabled      Peer Tracking          : Disabled
Auth key chain         : n/a
Disable Cap Nego       : Disabled      Bfd Enabled            : Disabled
Flowspec Validate      : Disabled      Default Route Tgt      : Disabled
Aigp Metric            : Disabled      Split Horizon          : Enabled
Local Capability        : RtRefresh MPBGP 4byte ASN
Remote Capability       : RtRefresh MPBGP 4byte ASN
Local AddPath Capabi*: Disabled
Remote AddPath Capab*: Send - None
                      : Receive - None
Import Policy           : None Specified / Inherited
Export Policy           : expPol
Origin Validation       : N/A
EBGP Link Bandwidth    : n/a
IPv4 Rej. Pfxs         : 0            IPv6 Rej. Pfxs         : 0
VPN-IPv4 Rej. Pfxs     : 0            VPN-IPv6 Rej. Pfxs     : 0
Mc IPv4 Rej. Pfxs      : 0            Mc IPv6 Rej. Pfxs      : 0
MVPN-IPv4 Rej. Pfxs    : 0            MVPN-IPv6 Rej. Pfxs    : 0
Flow-IPv4 Rej. Pfxs    : 0            Flow-IPv6 Rej. Pfxs    : 0
L2-VPN Rej. Pfxs       : 0            MDT-SAFI Rej. Pfxs     : 0
Rte-Tgt Rej. Pfxs      : 0            MS-PW Rej. Pfxs        : 0
Mc Vpn Ipv4 Rej. Pfxs : 0            Evpn Rej. Pfxs         : 0
Label-v4 Suppr. Pfxs   : 0            Label-v4 Recd. Pfxs     : 0
Label-v4 Active Pfxs   : 0            Label-v4 Rej. Pfxs      : 0
Label-v6 Suppr. Pfxs   : 0            Label-v6 Recd. Pfxs     : 0
Label-v6 Active Pfxs   : 0            Label-v6 Rej. Pfxs      : 0
Bgp-Ls Suppr. Pfxs     : 0            Bgp-Ls Recd. Pfxs      : 0

```

Bgp-Ls Active Pfxs : 0 Bgp-Ls Rej. Pfxs : 0

Table 166: BGP neighbor (standard, detailed, and dynamic) field descriptions

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
Dynamic Peer	Yes: the session is dynamic (that is, unconfigured)
	No: the session is statically configured
State	Idle: The BGP peer is not accepting connections. (Shutdown) is also displayed if the peer is administratively disabled.
	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

Label	Description
	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
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

Label	Description
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
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

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
Pref Limit Idle Time*	The length of time that the session is held in the idle state after it is taken down as a result of reaching the prefix limit
Hold Time	The configured hold-time setting
Keep Alive	The configured keepalive setting
Min Hold Time	The configured minimum hold-time setting
Active Hold Time	The negotiated hold time, if the BGP neighbor is in an established state
Active Keep Alive	The negotiated keepalive time, if the BGP neighbor is in an established state
Cluster Id	The configured route reflector cluster ID
	None: no cluster ID has been configured
Client Reflect	Disabled: The BGP route reflector is configured not to reflect routes to this neighbor
	Enabled: The BGP route reflector is configured to reflect routes to this neighbor
Preference	The configured route preference value for the peer group
Num of Flaps	The number of route flaps in the neighbor connection
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

Label	Description
VPN-IPv4 Active Pfxs	The number of VPN-IPv4 routes received from the BGP neighbor and active in the forwarding table
IPv6 Suppressed. Pfxs	The number of unique sets of IPv6 path attributes received from the BGP neighbor and suppressed due to route damping
IPv6 Recd. Prefixes	The number of unique sets of IPv6 path attributes received from the BGP neighbor
IPv6 Active Prefixes	The number of IPv6 routes received from the BGP neighbor and active in the forwarding table
VPN-IPv6 Recd. Pfxs	The number of unique sets of VPN-IPv6 path attributes received from the BGP neighbor
VPN-IPv6 Active Pfxs	The number of VPN-IPv6 routes received from the BGP neighbor and active in the forwarding table
VPN-IPv6 Suppr. Pfxs	The number of unique sets of VPN-IPv6 path attributes received from the BGP neighbor and suppressed due to route damping
MVPN-IPv4 Suppr. Pfxs	The number of unique sets of MVPN-IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
MVPN-IPv4 Recd. Pfxs	The number of unique sets of MVPN-IPv4 path attributes received from the BGP neighbor
MVPN-IPv4 Active Pfxs	The number of MVPN-IPv4 routes received from the BGP neighbor and active in the forwarding table
Flow-IPv4 Suppr. Pfxs	N/A
Flow-IPv4 Recd. Pfxs	N/A
Flow-IPv4 Active Pfxs	N/A
Rte-Tgt Suppr. Pfxs	The number of unique sets of route target path attributes received from the BGP neighbor and suppressed due to route damping
Rte-Tgt Recd. Pfxs	The number of unique sets of route target path attributes received from the BGP neighbor
Rte-Tgt Active. Pfxs	The number of route target routes received from the BGP neighbor and active in the forwarding table
Backup IPv4 Pfxs	The number of BGP fast reroute backup path IPv4 prefixes
Backup IPv6 Pfxs	The number of BGP fast reroute backup path IPv6 prefixes

Label	Description
Mc Vpn Ipv4 Suppr. P*	N/A
Backup Vpn IPv4 Pfxs	The number of BGP fast reroute backup path VPN IPv4 prefixes
Backup Vpn IPv6 Pfxs	The number of BGP fast reroute backup path VPN IPv6 prefixes
Input Queue	The number of BGP messages to be processed
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
Evpn Suppr. Pfxs	The number of unique sets of EVPN-IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
Evpn Recd. Pfxs	The number of unique sets of EVPN-IPv4 path attributes received from the BGP neighbor
Evpn Active Pfxs	The number of EVPN-IPv4 routes received from the BGP neighbor and active in the forwarding table
MS-PW Suppr. Pfxs	N/A
MS-PW Recd. Pfxs	N/A
MS-PW Active Pfxs	N/A
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
Advertise Inactive	The state of advertising inactive BGP routes to other BGP peers (enabled or disabled)

Label	Description
Peer Tracking	The state of tracking a neighbor IP address in the routing table for a BGP session
Advertise Label	Indicates the enabled address family for supporting RFC 3107 BGP label capability
Auth key chain	The value for the authentication key chain
Disable Cap Nego	N/A
Bfd Enabled	Enabled: BFD is enabled Disabled: BFD is disabled
Flowspec Validate	N/A
Default Route Tgt	Indicates that the default RTC route (zero prefix length) is originated toward the selected peers
Aigp Metric	Indicates whether accumulated IGP (AIGP) path attribute support with one or more BGP peers is enabled or disabled
Split Horizon	Indicates whether split horizon is enabled or disabled, When enabled, split horizon prevents routes from being reflected back to a peer that sends the best route.
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
Local AddPath Capabi*	The state of the local BGP add-paths capabilities. The add-paths capability allows the router to send and receive multiple paths per prefix to or from a peer.
Remote AddPath Capab*	The state of the remote BGP add-paths capabilities
Import Policy	The configured import policies for the peer group
Export Policy	The configured export policies for the peer group
Origin Validation	N/A
EBGP Link Bandwidth	N/A
IPv4 Rej. Pfxs	The number of unique sets of IPv4 path attributes received from the BGP neighbor and rejected by the router
IPv6 Rej. Pfxs	The number of unique sets of IPv6 path attributes received from the BGP neighbor and rejected by the router

Label	Description
VPN-IPv4 Rej. Pfxs	The number of unique sets of VPN-IPv4 path attributes received from the BGP neighbor and rejected by the router
VPN-IPv6 Rej. Pfxs	The number of unique sets of VPN-IPv6 path attributes received from the BGP neighbor and rejected by the router
Mc IPv4 Rej. Pfxs	The number of unique sets of MC IPv4 path attributes received from the BGP neighbor and rejected by the router
Mc IPv6 Rej. Pfxs	The number of unique sets of MC IPv6 path attributes received from the BGP neighbor and rejected by the router
MVPN-IPv4 Rej. Pfxs	The number of unique sets of MVPN-IPv4 path attributes received from the BGP neighbor and rejected by the router
MVPN-IPv6 Rej. Pfxs	The number of unique sets of MVPN-IPv6 path attributes received from the BGP neighbor and rejected by the router
Flow-IPv4 Rej. Pfxs	The number of unique sets of path attributes received from the BGP neighbor and rejected by the router
Flow-IPv6 Rej. Pfxs	The number of unique sets of Flow-IPv6 path attributes received from the BGP neighbor and rejected by the router
L2-VPN Rej. Pfxs	The number of unique sets of L2-VPN path attributes received from the BGP neighbor and rejected by the router
MDT-SAFI Rej. Pfxs	The number of unique sets of MDT-SAFI path attributes received from the BGP neighbor and rejected by the router
Rte-Tgt Rej. Pfxs	The number of unique sets of route target path attributes received from the BGP neighbor and rejected by the router
MS-PW Rej. Pfxs	The number of unique sets of MS-PW path attributes received from the BGP neighbor and rejected by the router
Mc Vpn Ipv4 Rej. Pfxs	The number of unique sets of MC VPN IPv4 path attributes received from the BGP neighbor and rejected by the router
Evpn Rej. Pfxs	The number of unique sets of EVPN path attributes received from the BGP neighbor and rejected by the router
Label-v4 Suppr. Pfxs	The number of unique sets of label-IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
Label-v4 Recd. Pfxs	The number of unique sets of label-IPv4 path attributes received from the BGP neighbor
Label-v4 Active Pfxs	The number of label-IPv4 routes received from the BGP neighbor and active in the forwarding table

Label	Description
Label-v4 Rej. Pfxs	The number of unique sets of label-IPv4 path attributes received from the BGP neighbor and rejected by the router
Label-v6 Suppr. Pfxs	The number of unique sets of label-IPv6 path attributes received from the BGP neighbor and suppressed due to route damping
Label-v6 Recd. Pfxs	The number of unique sets of label-IPv6 path attributes received from the BGP neighbor
Label-v6 Active Pfxs	The number of label-IPv6 routes received from the BGP neighbor and active in the forwarding table
Label-v6 Rej. Pfxs	The number of unique sets of label-IPv6 path attributes received from the BGP neighbor and rejected by the router
Bgp-Ls Suppr. Pfxs	The number of unique sets of BGP LS path attributes received from the BGP neighbor and suppressed due to route damping
Bgp-Ls Recd. Pfxs	The number of unique sets of BGP LS path attributes received from the BGP neighbor
Bgp-Ls Active Pfxs	The number of BGP LS routes received from the BGP neighbor and active in the forwarding table
Bgp-Ls Rej. Pfxs	The number of unique sets of BGP LS path attributes received from the BGP neighbor and rejected by the router

Output example - BGP neighbor (advertised-routes and received-routes)

```

*A:ALU-12# show router 6 bgp neighbor 10.44.44.44 advertised-routes
=====
BGP Router ID : 10.55.55.55      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              Path-Id    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
?    10.10.10.0/24
?    10.10.10.1/24
?    10.10.10.2/24
?    10.10.10.3/24
-----
Routes : 5
=====
*A:ALU-12#

*A:ALU-12# show router 6 bgp neighbor 10.44.44.44 received-routes
=====
BGP Router ID : 10.55.55.55        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                Path-Id    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 167: BGP neighbor (advertised-routes and received-routes) field descriptions

Label	Description
BGP Router ID	The local BGP router ID

Label	Description
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/Flags	<p>Legend:</p> <p>Status codes:</p> <p>u - used</p> <p>s - suppressed</p> <p>h - history</p> <p>d - decayed</p> <p>* - valid</p> <p>If an * is not present, then the status is invalid</p> <p>Origin codes:</p> <p>i - IGP</p> <p>e - EGP</p> <p>? - incomplete</p> <p>> - best</p>
Network	The route IP prefix and mask length for the route
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

Output example - BGP neighbor (graceful restart)

```
*A:ALU-12# show router 6 bgp neighbor 10.10.120.44 graceful-restart
=====
BGP Neighbor 10.10.120.44 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 168: BGP neighbor (graceful restart) field descriptions

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
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 VPN-IPv4 routes

ipv6 – displays only those BGP peers that have the IPv6 family enabled and not those capable of exchanging VPN-IPv6 routes

label-ipv4 – displays only those BGP peers that support the IPv4 unicast (labeled) address family

vpn-ipv4 – displays only those BGP peers that have the VPN-IPv4 family enabled

vpn-ipv6 – displays only those BGP peers that have the VPN-IPv4 family enabled

mvpn-ipv4 – displays only those BGP peers that have the MVPN-IPv4 family enabled

ip-address

displays the next hop information for the specified IPv4 or IPv6 address

detail

displays the more detailed version of the output

Output

The following output is an example of BGP next-hop information, and [Table 169: BGP next-hop field descriptions](#) describes the fields.

Output example

```
*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                            7    ISIS
  10.20.1.1/32                          10
  10.10.2.1                             592
10.10.10.12                            7    ISIS
  10.20.1.2/32                          10
  10.10.3.2                             592
```

```
10.20.1.4                                     7    ISIS
  10.20.0.0/32                               20
  10.10.11.4                                 8
-----
Next Hops : 3
=====
A:ALU-12#

*A:ALU-12# show router 6 bgp next-hop 10.0.0.1
=====
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
-----
10.0.0.1                                     15    ISIS
  10.0.0.0/24                               20
  10.88.1.2                                 8
10.0.0.1                                     15    ISIS
  10.0.0.0/24                               20
  10.88.2.2                                 8
-----
Next Hops : 1
=====
A:ALU-12#

*A:ALU-12# show router 6 bgp next-hop 10.0.0.1 detail
=====
BGP Router ID : 10.10.10.1      AS : 65000   Local AS : 65000
=====
BGP Next Hop
=====
Next Hop: 10.0.0.1
-----
Resolving Prefix : 10.0.0.0/24
Preference       : 15                Metric       : 20
Reference Count  : 8                Owner        : ISIS
Resolved Next Hop: 10.88.1.2
Egress Label    : N/A
Resolved Next Hop: 10.88.2.2
Egress Label    : N/A
Resolved Next Hop: 10.88.3.2
Egress Label    : N/A
-----
Next Hops : 1
=====
A:ALU-12#
```

Table 169: BGP next-hop field descriptions

Label	Description
BGP Router ID	The local BGP router ID
AS	The configured autonomous system number

Label	Description
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 170: BGP path field descriptions](#) describes the fields.

Output example

```
*A:ALU-12# show router 6 bgp paths
=====
BGP Router ID : 10.55.55.55 AS : 65000   Local AS : 65000
=====
BGP Paths
=====
Path: No As-Path
-----
Next Hop      : 10.44.10.12
Origin        : Incomplete                Segments      : 0
MED           : None                      Local Preference : 4294967295
```

```

Refs      : 1080          ASes      : 0
Flags     : IBGP-learned
-----
Path: No As-Path
-----
Next Hop   : 10.88.1.2
MED        : 10          Local Preference : None
Refs      : 4          ASes      : 0
Flags     : Imported
-----
Path: No As-Path
-----
Next Hop   : 10.44.10.21
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 170: BGP path field descriptions

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

Label	Description
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 [*ip-prefix/mask* | *ip-address*]

routes aspath-regex *reg-exp* {**detail** | **longer**}

routes aspath-regex *reg-exp*

routes aspath-regex *reg-exp* **hunt**

routes brief

routes community *comm-id* {**detail** | **longer**}

routes community *comm-id* **hunt**

routes detail

routes hunt [**brief**]

routes ipv4 [**aspath-regex** *reg-exp*] [**community** *comm-id*] [**brief**] [**all**]

routes ipv4 [**aspath-regex** *reg-exp*] **hunt** [**community** *comm-id*] [**brief**] [**all**]

routes ipv4 [**detail** | **longer**] [**aspath-regex** *reg-exp*] [**community** *comm-id*] [**all**]

routes ipv6 [**aspath-regex** *reg-exp*] [**community** *comm-id*] [**brief**] [**all**]

routes ipv6 [**aspath-regex** *reg-exp*] **hunt** [**community** *comm-id*] [**brief**] [**all**]

routes ipv6 [**detail** | **longer**] [**aspath-regex** *reg-exp*] [**community** *comm-id*] [**all**]

routes label-ipv4 [**aspath-regex** *reg-exp*] [**community** *comm-id*] [**brief**] [**all**]

routes label-ipv4 [**aspath-regex** *reg-exp*] **hunt** [**community** *comm-id*] [**brief**] [**all**]

routes label-ipv4 [**detail** | **longer**] [**aspath-regex** *reg-exp*] [**community** *comm-id*] [**all**]

routes longer

routes mvpn-ipv4 [**aspath-regex** *reg-exp*] [**community** *comm-id*] [**rd** *rd*] [**brief**] [**type** *mvpn-type*]

[**originator-ip** *ip-address*] [**source-ip** *ipv4 address* | *ipv6 address*] [**group-ip** *ipv4 address* | *ipv6 address*]
[**source-as** *as-number*]

routes mvpn-ipv4 [**aspath-regex** *reg-exp*] **hunt** [**community** *comm-id*] [**rd** *rd*] [**brief**] [**type** *mvpn-type*]

[**originator-ip** *ip-address*] [**source-ip** *ipv4 address* | *ipv6 address*] [**group-ip** *ipv4 address* | *ipv6 address*]
[**source-as** *as-number*]

```

routes mvpn-ipv4 [detail | longer] [aspath-regex reg-exp] [community comm-id] [rd rd] [type mvpn-type]
[originator-ip ip-address] [source-ip ipv4 address | ipv6 address] [group-ip ipv4 address | ipv6 address]
[source-as as-number]

routes route-target [source-as as-number] [brief] [aspath-regex reg-exp] [community comm-id]

routes route-target [rtc-prefix rtc-prefix] [hunt] [brief] [aspath-regex reg-exp] [community comm-id]

routes route-target rtc-prefix rtc-prefix [aspath-regex reg-exp] [community comm-id]

routes route-target [rtc-prefix rtc-prefix] [detail | longer] [aspath-regex reg-exp] [community comm-id]

routes vpn-ipv4 [aspath-regex reg-exp] [community comm-id] [rd rd] [brief]

routes vpn-ipv4 [aspath-regex reg-exp] hunt [community comm-id] [rd rd] [brief]

routes vpn-ipv4 [detail | longer] [aspath-regex reg-exp] [community comm-id] [rd rd]

routes vpn-ipv6 [aspath-regex reg-exp] [community comm-id] [rd rd] [brief]

routes vpn-ipv6 [aspath-regex reg-exp] hunt [community comm-id] [rd rd] [brief]

routes vpn-ipv6 [detail | longer] [aspath-regex reg-exp] [community comm-id] [rd rd]

```

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.



Note: To apply a family filter to the route output of the command, the family name must be specified before all other filtering parameters except for the IP prefix/mask or IP address, which, if present, must be placed before the family name in the command.

Parameters

ip-prefix/mask | *ip-address*

displays parameters that match the specified IPv4 prefix and mask length or IPv6 address

aspath-regex *reg-exp*

displays all routes with an AS path matching the specified regular expression (80 characters maximum)

brief

provides a summarized display of the set of peers to which a BGP route is advertised; this option is only supported when no IP prefix/mask or IP address is specified

community

displays all routes with the specified BGP community; community matching is based only on RIB-In communities, not RIB-Out communities

comm-id

specifies community IDs, in the format *as-number1:comm-val1* | *ext-comm* | *well-known-comm*

Values	<i>as-number1</i>	0 to 65535
	<i>comm-val1</i>	0 to 65535
	<i>ext-comm</i>	<i>type</i> :{ <i>ip-address:comm-val1</i> <i>as-number1:comm-val2</i> <i>as-number2:comm-val1</i> <i>as-number1:val-in-mbps</i> } <i>ext:xyy:ovstate</i> where <i>type</i> : target origin bandwidth (keywords) <i>ip-address</i> : ipv4-address: a.b.c.d ipv6-address: x:x:x:x:x:x:x[- <i>interface</i>] x:x:x:x:x:d.d.d.d[- <i>interface</i>] <i>interface</i> - 32 chars max, mandatory for link local addresses x: [0 to FFFF]H d: [0 to 255]D
	<i>ext:xyy:ovstate</i>	<i>xx</i> : 43 <i>yy</i> : 0 <i>ovstate</i> : 0, 1, 2 (0 for valid, 1 for not-found, 2 for invalid)
	<i>well-known-comm</i>	null no-export no-export-subconfed no-advertise (keywords)

group-ip *ipv4 address* | *ipv6 address*

displays the group IPv4 or IPv6 address

originator-ip *ip-address*

displays the originating IPv4 or IPv6 address

source-ip *ipv4 address* | *ipv6 address*

displays the source IPv4 or IPv6 address

detail

displays a more detailed version of the output

source-as *as-number*

displays the source AS number

hunt

displays entries for the specified route in the RIB-In, RIB-Out, and RTM

ipv4

displays only those BGP peers that have the IPv4 family enabled

ipv6

displays only those BGP peers that have the IPv6 family enabled

label-ipv4

displays only those BGP peers for the IPv4 unicast (labeled) address family

longer

displays the specified route and subsets of the route

mvpn-ipv4

displays the BGP peers that are MVPN-IPv4 capable

mvpn-type

the specified MVPN type for which to display information

Values intra-ad | inter-ad | spmsi-ad | leaf-ad | source-ad | shared-join | source-join

route-target

displays a summary of route target constrained routes for this BGP peer

rtc-prefix *rtc-prefix*

displays route target constraint prefix information, in the format *source-as:{ip-addr:comm-val | 2byte-asnumber:ext-comm-val | 4byte-asnumber:comm-val}/prefix-length*

Values *ip-addr:*
ipv4-address: a.b.c.d
ipv6-address: x:x:x:x:x:x:x[-*interface*] x:x:x:x:x:d.d.d.d[-*interface*]
interface - 32 chars max, mandatory for link local addresses x: [0 to FFFF]H d: [0 to 255]D
comm-val: 0 to 65535
2byte-asnumber: 0 to 65535
ext-comm-val: 0 to 4294967295
4byte-asnumber: 0 to 4294967295
prefix-length: 0 to 96

rd *rd*

displays the route distinguisher value, in the format *ip-addr:comm-val | 2byte-asnumber:ext-comm-val | 4-byte-asnumber:comm-val*

Values *ip-addr:*
ipv4-address: a.b.c.d
ipv6-address: x:x:x:x:x:x:x[-*interface*] x:x:x:x:x:d.d.d.d[-*interface*]
interface - 32 chars max, mandatory for link local addresses x: [0 to FFFF]H d: [0 to 255]D
comm-val: 0 to 65535
2byte-asnumber: 0 to 65535
ext-comm-val: 0 to 4294967295
4byte-asnumber: 0 to 4294967295

vpn-ipv4

displays the BGP VPN IPv4 routes

vpn-ipv6

displays the BGP VPN IPv6 routes

originator-ip *ip-address*

filters BGP MVPN routes by the originating router IP address that is found in the intra-AD (auto-discovery) MVPN routes

source-ip *ip-address*

filters BGP MVPN routes by the source IP address that is found in the source-join, source-AD, or S-PMSI-AD MVPN routes

group-ip *ip-address*

filters BGP MVPN routes by the multicast group IP address that is found in the source-join, source-AD, or S-PMSI-AD MVPN routes

source-as *as-number*

filters BGP MVPN routes by source-AS (autonomous system) extended community attribute

Output

The following output is an example of BGP route information, and [Table 171: BGP routes field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B# show router bgp routes 10.10.10.5
=====
BGP Router ID:10.20.1.3      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP IPv4 Routes
=====
Flag  Network                LocalPref  MED
      Nexthop (Router)      Path-Id    Label
      As-Path
-----
u*>?  10.10.10.0/24           None       None
      10.20.1.4             None       -
      200 300
-----
Routes : 1
=====
*A:Sar18 Dut-B##

*A:Sar18 Dut-B# show>router>bgp# routes vpn-ipv4 10.10.10.6/32 rd 10.20.1.4:1 hunt
=====
BGP Router ID:10.20.1.3      AS:None      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
```

```

Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP VPN-IPv4 Routes
=====
-----
RIB In Entries
-----
Network       : 10.10.10.6/32
Nextthop      : 10.20.1.4
Route Dist.   : 10.20.1.4:1      VPN Label       : 131070
Path Id       : None
From          : 10.20.1.4
Res. Nextthop : n/a
Local Pref.   : 100
Aggregator AS : None             Interface Name : int_to_D
Atomic Aggr.  : Not Atomic       Aggregator    : None
AIGP Metric   : None            MED           : None
Connector     : None
Community     : target:100:100
Cluster       : No Cluster Members
Originator Id : None             Peer Router Id : 10.20.1.4
Fwd Class     : None            Priority       : None
Flags         : Used Valid Best Incomplete
Route Source  : Internal
AS-Path       : 106
VPRN Imported : 1
-----
RIB Out Entries
-----
-----
Routes : 1
=====
*A:Sar18 Dut-B#

*A:Sar18 Dut-B# show>router>bgp# routes 3FFE::606:609/128 vpn-ipv6 hunt
=====
BGP Router ID:10.20.1.3      AS:None      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
              l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP VPN-IPv6 Routes
=====
-----
RIB In Entries
-----
Network       : 3FFE::606:609/128
Nextthop      : ::FFFF:A14:104
Route Dist.   : 10.20.1.4:1      VPN Label       : 131070
Path Id       : None
From          : 10.20.1.4
Res. Nextthop : n/a
Local Pref.   : 100
Aggregator AS : None             Interface Name : int_to_D
Atomic Aggr.  : Not Atomic       Aggregator    : None
AIGP Metric   : None            MED           : None
Connector     : None
Community     : target:100:100
Cluster       : No Cluster Members
Originator Id : None             Peer Router Id : 10.20.1.4
Fwd Class     : None            Priority       : None

```

```

Flags      : Used Valid Best Incomplete
Route Source : Internal
AS-Path     : 106
VPRN Imported : 1
-----
RIB Out Entries
-----
Routes : 1
=====
*A: Sar18 Dut-B#

*A: Sar18 Dut-B# show>router>bgp# routes vpn-ipv6 3FFE::606:607 128 rd 10.20.1.4:1
hunt
=====
BGP Router ID:10.20.1.3      AS:None      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP VPN-IPv6 Routes
=====
RIB In Entries
-----
Network      : 3FFE::606:607/128
Nexthop      : ::FFFF:A14:104
Route Dist.  : 10.20.1.4:1      VPN Label    : 131070
Path Id      : None
From         : 10.20.1.4
Res. Nexthop : n/a
Local Pref.  : 100
Aggregator AS : None           Interface Name : int_to_D
Atomic Aggr. : Not Atomic      Aggregator    : None
AIGP Metric  : None           MED           : None
Connector    : None
Community    : target:100:100
Cluster      : No Cluster Members
Originator Id : None           Peer Router Id : 10.20.1.4
Fwd Class    : None           Priority       : None
Flags        : Used Valid Best Incomplete
Route Source : Internal
AS-Path      : 106
VPRN Imported : 1
-----
RIB Out Entries
-----
Routes : 1
=====
*A: Sar18 Dut-B#

*A: Sar18 Dut-B# show>router>bgp# routes vpn-ipv6 3FFE::606:607/128 rd 10.20.1.4:2
hunt
=====
BGP Router ID:10.20.1.3      AS:None      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge

```

```

Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP VPN-IPv6 Routes
=====
No Matching Entries Found
=====
*A:Sar18 Dut-B#

*A:Sar18 Dut-B# show router bgp routes hunt 10.10.10.1/32
=====
BGP Router ID:10.20.1.3      AS:5000      Local AS:5000
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP IPv4 Routes
=====
RIB In Entries
-----
Network       : 10.10.10.1/32
Nexthop       : 10.20.1.1
From          : 10.20.1.1
Res. Nexthop  : 10.20.1.1 (RSVP LSP: 1)
Local Pref.   : 100
Aggregator AS : None
Atomic Aggr.  : Not Atomic
Community     : No Community Members
Cluster       : No Cluster Members
Originator Id : None
Flags         : Used Valid Best Incomplete
AS-Path       : No As-Path
Peer Router Id : 10.20.1.1
Interface Name : ip-10.10.2.3
Aggregator    : None
MED           : None
-----
RIB Out Entries
-----
Routes : 1
=====
*A:Sar18 Dut-B#

*A:Sar18 Dut-B# show router bgp routes family ipv4
=====
BGP Router ID : 10.10.10.103      AS : 200      Local AS : 200
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Routes
=====
Flag  Network      Nexthop      LocalPref  MED
     VPN Label      As-Path
-----
No Matching Entries Found
=====
*A:Sar18 Dut-B#

*A:Sar18 Dut-B# show router bgp routes 10.10.10.0/24 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
                l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Routes
=====
Original Attributes
Network       : 10.10.10.0/24 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 :10.10.10.0/24 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: Sar18 Dut-B#

*A: Sar18 Dut-B# show router bgp routes 10.10.10.0/24 hunt
=====
BGP Router ID : 10.20.1.1 AS : 100Local AS : 100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
                l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Routes
=====
RIB In Entries
-----
Network       : 10.10.10.0/24
Nexthop       : 10.20.1.2
Route Dist.   : 10.20.1.2:1VPN Label: 131070
From          : 10.20.1.2
Res. Nexthop  : 10.10.1.2
Local Pref.   : 100Interface Name: to-sr7
Aggregator AS : noneAggregator: none
Atomic Aggr.  : Not AtomicMED: none
Community     : target:10.20.1.2:1
Cluster       : No Cluster Members
Originator Id : NonePeer 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:Sar18 Dut-B#

*A:Sar18 Dut-B# show router bgp routes mvpn-ipv4
=====
BGP Router ID:10.20.1.3      AS:200      Local AS:200
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP MVPN-IPv4 Routes
=====
Flag RouteType      OriginatorIP      LocalPref  MED      VPNLabel
      RD
      Nexthop
      As-Path
      SourceAS
      SourceIP
      GroupIP
-----
u*>i Intra-Ad        10.20.1.4        100        0
      1:1            -
      10.20.1.4      -
      No As-Path     -
u*>i Source-Ad       -                100        0
      1:1            -
      10.20.1.4      10.100.1.2
      No As-Path     10.0.0.0
u*>i Source-Join     -                100        0
      1:1            200
      10.20.1.4      10.100.1.2
      No As-Path     10.0.0.0
-----
Routes : 3
=====
*A:Sar18 Dut-B#

*A:Sar18 Dut-B# show router bgp routes mvpn-ipv4 brief
=====
BGP Router ID:10.20.1.3      AS:200      Local AS:200
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP MVPN-IPv4 Routes
=====
Flag RouteType      OriginatorIP      SourceIP
      RD
      SourceAS
      GroupIP
-----
u*>i Intra-Ad        10.20.1.4        -
      1:1            -
u*>i Source-Ad       -                10.100.1.2
      1:1            -                10.0.0.0
u* >i Source-Join     -                10.100.1.2
      1:1            200                10.0.0.0

```

```

-----
Routes : 3
=====
*A:Sar18 Dut-B#

*A:Sar18 Dut-B# show router bgp routes mvpn-ipv4 type source-join source-as 200
source-ip 10.100.1.2 group-ip 10.0.0.0 detail
=====
BGP Router ID:10.20.1.3      AS:200      Local AS:200
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP MVPN-IPv4 Routes
=====
Route Type      : Source-Join
Route Dist.     : 1:1
Source AS       : 200
Source IP       : 10.100.1.2
Group IP        : 10.0.0.0
Nexthop         : 10.20.1.4
From            : 10.20.1.4
Res. Nexthop    : 10.0.0.0
Local Pref.     : 100
Aggregator AS   : None
Atomic Aggr.    : Not Atomic
Community       : target:10.20.1.3:2
Cluster         : No Cluster Members
Originator Id   : None
Flags           : Used Valid Best IGP
AS-Path         : No As-Path
Interface Name  : NotAvailable
Aggregator      : None
MED             : 0
Peer Router Id  : 10.20.1.4
-----
Routes : 1
=====
*A:Sar18 Dut-B#

*A:Sar18 Dut-B# show router bgp routes ipv4 detail
=====
BGP Router ID:10.1.1.1      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
-----
Original Attributes
-----
Network        : 10.1.1.1/32
Nexthop        : 192.168.1.1
Path Id        : None
From           : 192.168.1.1
Res. Nexthop   : 192.168.1.1
Local Pref.    : n/a
Aggregator AS  : None
Atomic Aggr.   : Not Atomic
AIGP Metric    : 100
Interface Name : net
Aggregator     : None
MED            : 5000

```

```

Community      : None
Cluster        : No Cluster Members
Originator Id  : None
Fwd Class      : None
Flags          : Used Valid Best Incomplete
Route Source   : External
AS-Path        : 200 400 500

```

Modified Attributes

```

Network        : 10.1.1.1/32
Nexthop        : 192.168.1.1
Path Id        : None
From           : 192.168.1.1
Res. Nexthop   : 192.168.1.1
Local Pref.    : None
Aggregator AS  : None
Atomic Aggr.   : Not Atomic
AIGP Metric    : 110
Community      : None
Cluster        : No Cluster Members
Originator Id  : None
Fwd Class      : None
Flags          : Used Valid Best Incomplete
Route Source   : External
AS-Path        : 200 400 500

```

```

-----
Routes : 1
=====

```

```

*A:Sar18 Dut-B#

```

```

*A:Sar18 Dut-B# show router bgp routes 10.1.1.1/32 hunt

```

```

=====
BGP Router ID:1.1.1.1      AS:100      Local AS:100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====

```

BGP IPv4 Routes

RIB In Entries

```

-----
Network        : 10.1.1.1/32
Nexthop        : 192.168.1.1
Path Id        : None
From           : 192.168.1.1
Res. Nexthop   : 192.168.1.1
Local Pref.    : None
Aggregator AS  : None
Atomic Aggr.   : Not Atomic
AIGP Metric    : 110
Community      : None
Cluster        : No Cluster Members
Originator Id  : None
Fwd Class      : None
Flags          : Used Valid Best Incomplete
Route Source   : External
AS-Path        : 200 400 500
-----

```


RIB Out Entries

```

-----
Network       : 10.1.1.1/32
Nexthop       : 10.1.1.1
Path Id       : None
To            : 10.3.3.3
Res. Nexthop  : n/a
Local Pref.   : 100
Aggregator AS : None
Atomic Aggr.  : Not Atomic
AIGP Metric   : 150
Community     : None
Cluster       : No Cluster Members
Originator Id : None
Origin        : Incomplete
AS-Path       : 200 400 500
Interface Name : NotAvailable
Aggregator    : None
MED           : 5000
Peer Router Id : 10.3.3.3
-----

```

```

-----
Routes : 2
=====

```

```

*A:Sar18 Dut-B#

```

```

*A:Sar18 Dut-B# show router bgp routes

```

```

=====
BGP Router ID:10.20.1.1      AS:1      Local AS:1
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
=====
BGP IPv4 Routes
=====
Flag  Network      LocalPref  MED
      Nexthop      Path-Id    Label
      As-Path
-----
u*>i  10.0.0.1/32    100        2010
      10.20.1.2    None       131057
      2
ub*i  10.0.0.1/32    100        2010
      10.20.1.3    None       131067
      2
-----
Routes : 2
=====
*A:Sar18 Dut-B#

```

```

*A:Sar18 Dut-B# show router bgp routes vpn-ipv4 community target:10.100.100.100:2
hunt

```

```

=====
BGP Router ID:10.20.1.6      AS:62000      Local AS:62000
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP VPN-IPv4 Routes
=====
RIB In Entries

```

```

-----
Network      : 10.1.1.1/32
Nexthop      : 10.20.1.6
Route Dist.  : 10.100.100.100:2      VPN Label      : 131068
Path Id      : None
From         : 10.20.1.5
Res. Nexthop : n/a
Local Pref.  : 100
Aggregator AS : None                  Interface Name : system
Atomic Aggr. : Not Atomic             Aggregator     : None
AIGP Metric  : None                   MED            : 100
Connector    : None
Community    : target:10.100.100.100:2
Cluster      : 5.5.5.5
Originator Id : 10.20.1.6              Peer Router Id : 10.20.1.5
Flags        : Invalid IGP
Route Source  : Internal
AS-Path       : No As-Path
Route Tag     : 0
Neighbor-AS   : N/A
Add Paths Send : Default
Last Modified : 00h00m06s
VPRN Imported : None
-----

```

RIB Out Entries

```

-----
Network      : 10.1.1.1/32
Nexthop      : 10.20.1.6
Route Dist.  : 10.100.100.100:2      VPN Label      : 131068
Path Id      : None
To           : 10.20.1.5
Res. Nexthop : n/a
Local Pref.  : 100
Aggregator AS : None                  Interface Name : NotAvailable
Atomic Aggr. : Not Atomic             Aggregator     : None
AIGP Metric  : None                   MED            : None
Connector    : None
Community    : target:10.100.100.100:2
Cluster      : No Cluster Members
Originator Id : None                  Peer Router Id : 10.20.1.5
Origin       : IGP
AS-Path      : No As-Path
Route Tag     : 0
Neighbor-AS   : N/A
-----

```

Routes : 2

*A:Sar18 Dut-B#

*A:Sar18 Dut-B# show router bgp routes vpn-
ipv6 community target:10.100.100.100:2 hunt

```

=====
BGP Router ID:10.20.1.3      AS:61000      Local AS:61000
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
               l - leaked, x - stale, > - best, b - backup, p - purge
Origin codes  : i - IGP, e - EGP, ? - incomplete
=====
BGP VPN-IPv6 Routes
=====
-----
RIB In Entries

```

```

-----
Network      : 3ffe::100:0/109
Nexthop      : ::ffff:10.20.1.1
Route Dist.  : 1:1234                VPN Label      : 131067
Path Id      : None
From         : 10.20.1.2
Res. Nexthop : n/a
Local Pref.  : 100
Aggregator AS : None                Interface Name : toA
Atomic Aggr. : Not Atomic           Aggregator    : None
AIGP Metric  : None                MED           : 100
Connector    : None
Community    : 0:0 target:10.100.100.100:2 target:1:123456
              origin:10.100.100.100:2
Cluster      : 2.2.2.2
Originator Id : 10.20.1.1            Peer Router Id : 10.20.1.2
Flags        : Used Valid Best IGP
Route Source  : Internal
AS-Path       : No As-Path
Route Tag     : 0
Neighbor-AS   : N/A
Add Paths Send : Default
Last Modified : 00h03m41s
VPRN Imported : 3 1

Network      : 3ffe::101:100/120
Nexthop      : ::ffff:10.20.1.1
Route Dist.  : 1:1234                VPN Label      : 131067
Path Id      : None
From         : 10.20.1.2
Res. Nexthop : n/a
Local Pref.  : 100
Aggregator AS : None                Interface Name : toA
Atomic Aggr. : Not Atomic           Aggregator    : None
AIGP Metric  : None                MED           : 100
Connector    : None
Community    : 0:0 target:10.100.100.100:2 target:1:123456
              origin:10.100.100.100:2
Cluster      : 10.2.2.2
Originator Id : 10.20.1.1            Peer Router Id : 10.20.1.2
Flags        : Used Valid Best IGP
Route Source  : Internal
AS-Path       : No As-Path
Route Tag     : 0
Neighbor-AS   : N/A
Add Paths Send : Default
Last Modified : 00h03m41s
VPRN Imported : 3 1

Network      : 3ffe::101:100/123
Nexthop      : ::ffff:10.20.1.1
Route Dist.  : 1:1234                VPN Label      : 131067
Path Id      : None
From         : 10.20.1.2
Res. Nexthop : n/a
Local Pref.  : 100
Aggregator AS : None                Interface Name : toA
Atomic Aggr. : Not Atomic           Aggregator    : None
AIGP Metric  : None                MED           : 100
Connector    : None
Community    : 0:0 target:10.100.100.100:2 target:1:123456
              origin:10.100.100.100:2
Cluster      : 2.2.2.2
Originator Id : 10.20.1.1            Peer Router Id : 10.20.1.2

```

```

Flags      : Used Valid Best IGP
Route Source : Internal
AS-Path    : No As-Path
Route Tag   : 0
Neighbor-AS : N/A
Add Paths Send : Default
Last Modified : 00h03m41s
VPRN Imported : 3 1

Network    : 3ffe::300:0/104
Nexthop    : ::ffff:10.20.1.1
Route Dist. : 1:1234          VPN Label      : 131067
Path Id    : None
From       : 10.20.1.2
Res. Nexthop : n/a
Local Pref. : 100
Aggregator AS : None          Interface Name : toA
Atomic Aggr. : Not Atomic     Aggregator   : None
AIGP Metric  : None          MED           : 100
Connector   : None
Community   : 0:0 target:10.100.100.100:2 target:1:123456
              origin:123456:2
Cluster     : 10.2.2.2
Originator Id : 10.20.1.1      Peer Router Id : 10.20.1.2
Flags      : Used Valid Best IGP
Route Source : Internal
AS-Path    : No As-Path
Route Tag   : 0
Neighbor-AS : N/A
Add Paths Send : Default
Last Modified : 00h03m41s
VPRN Imported : 3 1

Network    : 3ffe::303:300/120
Nexthop    : ::ffff:10.20.1.1
Route Dist. : 1:1234          VPN Label      : 131067
Path Id    : None
From       : 10.20.1.2
Res. Nexthop : n/a
Local Pref. : 100
Aggregator AS : None          Interface Name : toA
Atomic Aggr. : Not Atomic     Aggregator   : None
AIGP Metric  : None          MED           : 100
Connector   : None
Community   : 0:0 target:100.100.100.100:2 target:1:123456
              origin:123456:2
Cluster     : 2.2.2.2
Originator Id : 10.20.1.1      Peer Router Id : 10.20.1.2
Flags      : Used Valid Best IGP
Route Source : Internal
AS-Path    : No As-Path
Route Tag   : 0
Neighbor-AS : N/A
Add Paths Send : Default
Last Modified : 00h03m41s
VPRN Imported : 3 1

Network    : ::/0
Nexthop    : ::ffff:10.20.1.1
Route Dist. : 10.100.100.100:2 VPN Label      : 131069
Path Id    : None
From       : 10.20.1.2
Res. Nexthop : n/a
Local Pref. : 100
Interface Name : toA

```

```

Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric : None
Connector : None
Community : target:10.100.100.100:2 origin:100.100.100.100:2
              origin:1:123456
Cluster : 2.2.2.2
Originator Id : 10.20.1.1
Flags : Used Valid Best IGP
Route Source : Internal
AS-Path : No As-Path
Route Tag : 0
Neighbor-AS : N/A
Add Paths Send : Default
Last Modified : 00h03m42s
VPRN Imported : 1

Network : 3ffe::100:0/104
Nexthop : ::ffff:10.20.1.1
Route Dist. : 10.100.100.100:2
Path Id : None
From : 10.20.1.2
Res. Nexthop : n/a
Local Pref. : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric : None
Connector : None
Community : 100:100 2:3 target:10.100.100.100:2
              origin:1:123456 origin:1:2
Cluster : 2.2.2.2
Originator Id : 10.20.1.1
Flags : Used Valid Best IGP
Route Source : Internal
AS-Path : No As-Path
Route Tag : 0
Neighbor-AS : N/A
Add Paths Send : Default
Last Modified : 00h03m42s
VPRN Imported : 1

Network : 3ffe::101:0/112
Nexthop : ::ffff:10.20.1.1
Route Dist. : 10.100.100.100:2
Path Id : None
From : 10.20.1.2
Res. Nexthop : n/a
Local Pref. : 100
Aggregator AS : None
Atomic Aggr. : Not Atomic
AIGP Metric : None
Connector : None
Community : 10:100 2:3 target:10.100.100.100:2
              origin:10.100.100.100:2 origin:1:123456 origin:1:2
Cluster : 2.2.2.2
Originator Id : 10.20.1.1
Flags : Used Valid Best IGP
Route Source : Internal
AS-Path : No As-Path
Route Tag : 0
Neighbor-AS : N/A
Add Paths Send : Default
Last Modified : 00h03m41s
VPRN Imported : 1

```

```

Network      : 3ffe::202:202/128
Nexthop      : ::ffff:10.20.1.1
Route Dist.  : 10.100.100.100:2      VPN Label    : 131069
Path Id      : None
From         : 10.20.1.2
Res. Nexthop : n/a
Local Pref.  : 100
Aggregator AS : None                  Interface Name : toA
Atomic Aggr. : Not Atomic             Aggregator    : None
AIGP Metric  : None                  MED           : 100
Connector    : None
Community    : 0:0 target:10.100.100.100:2
              origin:10.100.100.100:2 origin:1:123456
Cluster      : 2.2.2.2
Originator Id : 10.20.1.1             Peer Router Id : 10.20.1.2
Flags        : Used Valid Best IGP
Route Source  : Internal
AS-Path       : No As-Path
Route Tag     : 0
Neighbor-AS   : N/A
Add Paths Send : Default
Last Modified : 00h03m41s
VPRN Imported : 1

Network      : ::/0
Nexthop      : ::ffff:10.20.1.1
Route Dist.  : 123456:2              VPN Label    : 131068
Path Id      : None
From         : 10.20.1.2
Res. Nexthop : n/a
Local Pref.  : 100
Aggregator AS : None                  Interface Name : toA
Atomic Aggr.  : Not Atomic             Aggregator    : None
AIGP Metric  : None                  MED           : 100
Connector    : None
Community    : target:10.100.100.100:2 origin:10.100.100.100:2
              origin:1:123456
Cluster      : 2.2.2.2
Originator Id : 10.20.1.1             Peer Router Id : 10.20.1.2
Flags        : Used Valid Best IGP
Route Source  : Internal
AS-Path       : No As-Path
Route Tag     : 0
Neighbor-AS   : N/A
Add Paths Send : Default
Last Modified : 00h03m41s
VPRN Imported : 1

Network      : 3ffe::100:0/104
Nexthop      : ::ffff:10.20.1.1
Route Dist.  : 123456:2              VPN Label    : 131068
Path Id      : None
From         : 10.20.1.2
Res. Nexthop : n/a
Local Pref.  : 100
Aggregator AS : None                  Interface Name : toA
Atomic Aggr.  : Not Atomic             Aggregator    : None
AIGP Metric  : None                  MED           : 100
Connector    : None
Community    : 10:100 2:3 target:10.100.100.100:2
              origin:1:123456 origin:1:2
Cluster      : 2.2.2.2
Originator Id : 10.20.1.1             Peer Router Id : 10.20.1.2

```

```

Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : N/A
Add Paths Send : Default
Last Modified  : 00h03m41s
VPRN Imported  : 1

Network        : 3ffe::101:0/112
Nextthop       : ::ffff:10.20.1.1
Route Dist.    : 123456:2          VPN Label      : 131068
Path Id        : None
From           : 10.20.1.2
Res. Nextthop  : n/a
Local Pref.    : 100
Aggregator AS  : None              Interface Name : toA
Atomic Aggr.   : Not Atomic        Aggregator     : None
AIGP Metric    : None              MED            : 100
Connector      : None
Community      : 10:100 2:3 target:10.100.100.100:2
                  origin:10.100.100.100:2 origin:1:123456 origin:1:2
Cluster        : 2.2.2.2
Originator Id  : 10.20.1.1          Peer Router Id : 10.20.1.2
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : N/A
Add Paths Send : Default
Last Modified  : 00h03m41s
VPRN Imported  : 1

Network        : 3ffe::202:202/128
Nextthop       : ::ffff:10.20.1.1
Route Dist.    : 123456:2          VPN Label      : 131068
Path Id        : None
From           : 10.20.1.2
Res. Nextthop  : n/a
Local Pref.    : 100
Aggregator AS  : None              Interface Name : toA
Atomic Aggr.   : Not Atomic        Aggregator     : None
AIGP Metric    : None              MED            : 100
Connector      : None
Community      : 0:0 target:100.100.100.100:2
                  origin:100.100.100.100:2 origin:1:123456
Cluster        : 2.2.2.2
Originator Id  : 10.20.1.1          Peer Router Id : 10.20.1.2
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : N/A
Add Paths Send : Default
Last Modified  : 00h03m41s
VPRN Imported  : 1

-----
RIB Out Entries
-----
-----
Routes : 13
=====
*A: Sar18 Dut-B#

```

Table 171: BGP routes field descriptions

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, the value is the same as the AS.
Flag/Flags	<p>Legend:</p> <p>Status codes:</p> <ul style="list-style-type: none"> u - used s - suppressed h - history d - decayed * - valid <p>If an * is not present, the status is invalid</p> <ul style="list-style-type: none"> l - leaked x - stale > - best b - backup p - purge <p>Origin codes:</p> <ul style="list-style-type: none"> 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 not present
VPN Label	The label generated by the PE label manager

Label	Description
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 IP address
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
TieBreakReason	<p>The step in the BGP decision process where a BGP route lost the tiebreaker with the next BGP route for the same prefix</p> <p>LocalPref – this route is not the best because the next better route has a higher LOCAL_PREF</p> <p>AIGP – this route is not the best because the next better route has a lower derived AIGP metric value</p> <p>ASPathLen – this route is not the best because the next better route has a shorter AS PATH length</p> <p>Origin – this route is not the best because the next better route has a lower origin value</p> <p>MED – this route is not the best because the next better route has a lower MED, and MED comparison of the routes was allowed</p>

Label	Description
	<p>IBGP – this IBGP route is not the best because the next better route is an EBGP route</p> <p>NHCost – this route is not the best because the next better route has a lower metric value to reach the BGP NEXT HOP</p> <p>BGPID – this route is not the best because the next better route has a lower originator ID or BGP identifier</p> <p>ClusterLen – this route is not the best because the next better route has a shorter cluster list length</p> <p>PeerIP – this route is not the best because the next better route has a lower neighbor IP address</p>
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

all

displays BGP peers in all instances

family

the type of routing information to be distributed by the BGP instance

Values ipv4, vpn-ipv4, ipv6, vpn-ipv6, mvpn-ipv4, label-ipv4

ipv4, ipv6 – displays only those BGP peers that have the IPv4 or IPv6 family enabled

label-ipv4 – displays only those BGP peers that have the IPv4 unicast (labeled) family enabled

vpn-ipv4, vpn-ipv6 – displays the BGP peers that are IPv4-VPN capable or IPv6-VPN capable

mvpn-ipv4 – displays the BGP peers that are MVPN-IPv4 capable

ip-address

clears damping information for entries received from the BGP neighbor

Output

The following output is an example of BGP summary information, and [Table 172: BGP summary field descriptions](#) describes the fields.

Output example

```
*A:ALU-12# show router 6 bgp summary
=====
BGP Router ID : 10.55.55.1      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 Suppressed 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
-----
10.44.10.12
      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
-----
10.44.10.12
Def. Instance  65000      662    0 04h14m52s 600/563/569 (IPv4)
```

```

564      0      8656/8656/8672 (VpnIPv4)
=====
*A:ALU-12#

*A:7705_ALU-2>show>router>bgp# summary neighbor 10.44.10.12
=====
BGP Router ID : 10.44.10.1      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 Suppressed 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
-----
10.44.10.12
      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 : 10.44.10.1      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 Suppressed 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
-----
10.44.10.12
      65000      679      581      0      0 04h23m36s 600/563/569
=====
*A:ALU-12#

```

Table 172: BGP summary field descriptions

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 Suppressed 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

Label	Description
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
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

servers

Syntax

servers [all]

Context

show>router>dhcp

show>router>dhcp6

Description

This command lists the local DHCP or DHCPv6 servers.

Parameters

all

displays DHCP or DHCPv6 servers in all instances

Output

The following output is an example of DHCP server information, and [Table 173: DHCP or DHCPv6 server field descriptions](#) describes the fields.

Output example

```
*A:ALU-1# show router dhcp servers
=====
Overview of DHCP Servers
=====
Active Leases:      1
Maximum Leases:    4096

Router              Server                      Admin State
-----
Router: Base        dhcpServer1                inService
Service: 102        vprnServer                  inService
=====
*A:ALU-1#
```

Table 173: DHCP or DHCPv6 server field descriptions

Label	Description
Active Leases	The number of active leases
Maximum Leases	The maximum number of leases available
Router	The name of the router
Server	The name of the DHCP or DHCPv6 server
Admin State	The administrative state of the DHCP or DHCPv6 server

statistics

Syntax

statistics [**interface** *ip-int-name* | *ip-address*]

Context

show>router>dhcp
show>router>dhcp6

Description

This command displays statistics for DHCP relay and DHCPv6 relay.
If no interface name or IP address is specified, then all configured interfaces are displayed. If the **statistics** command is used in the **dhcp6** context, the interface name or IP address cannot be specified.

Parameters

ip-int-name | *ip-address*
displays statistics for the specified IP interface

Output

The following outputs are examples of DHCP or DHCPv6 statistics information:

- DHCP statistics ([Output example](#), [Table 174: DHCP statistics field descriptions](#))
- DHCPv6 statistics ([Output example](#), [Table 175: DHCPv6 statistics field descriptions](#))

Output example

```
*A:ALU-1# show router dhcp statistics
=====
DHCP Global Statistics (Router: Base)
=====
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 174: DHCP statistics field descriptions

Label	Description
DHCP Global Statistics (Router: Base)	
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

Output example

```

*A:ALU-1# show router dhcp6 statistics
=====
DHCP6 statistics (Router: Base)
=====
Msg-type           Rx           Tx           Dropped
-----
1 SOLICIT           0            0            0
2 ADVERTISE          0            0            0
3 REQUEST            0            0            0
4 CONFIRM            0            0            0
5 RENEW              0            0            0
6 REBIND             0            0            0
7 REPLY              0            0            0
8 RELEASE            0            0            0
9 DECLINE            0            0            0
10 RECONFIGURE        0            0            0
11 INFO_REQUEST       0            0            0
12 RELAY_FORW         0            0            0
13 RELAY_REPLY        0            0            0

-----
Dhcp6 Drop Reason Counters :
-----
1 Dhcp6 oper state is not Up on src itf           0
2 Dhcp6 oper state is not Up on dst itf           0
3 Relay Reply Msg on Client Itf                   0
4 Hop Count Limit reached                         0
5 Missing Relay Msg option, or illegal msg type    0
6 Unable to determine destination client Itf        0
7 Out of Memory                                   0
8 No global Pfx on Client Itf                     0
9 Unable to determine src Ip Addr                  0
10 No route to server                             0
11 Subscr. Mgmt. Update failed                     0
12 Received Relay Forw Message                    0
13 Packet too small to contain valid dhcp6 msg     0
14 Server cannot respond to this message           0
15 No Server Id option in msg from server           0
16 Missing or illegal Client Id option in client msg 0
17 Server Id option in client msg                  0
18 Server DUID in client msg does not match our own 0
19 Client sent message to unicast while not allowed 0
20 Client sent message with illegal src Ip address 0
21 Client message type not supported in pfx delegation 0
22 Nbr of addrs or pfxs exceeds allowed max (128) in msg 0
23 Unable to resolve client's mac address          0
24 The Client was assigned an illegal address       0
25 Illegal msg encoding                           0
=====
*A:ALU-1#

```

Table 175: DHCPv6 statistics field descriptions

Label	Description
DHCP6 Statistics (Router: Base)	
Msg-type	The number of messages received, transmitted, or dropped by the router for each message type

Label	Description
Dhcp6 Drop Reason Counters	The number of times that a message was dropped for a particular reason

summary

Syntax

summary

Context

show>router>dhcp

show>router>dhcp6

Description

This command displays a summary of DHCP and DHCPv6 configuration.

Output

The following outputs are examples of DHCP or DHCPv6 summary information:

- DHCP summary ([Output example, Table 176: DHCP summary field descriptions](#))
- DHCPv6 summary ([Output example, Table 177: DHCPv6 summary field descriptions](#))

Output example

```
*A:ALU-48# show router dhcp summary
=====
DHCP Summary (Router: Base)
=====
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 176: DHCP summary field descriptions

Label	Description
DHCP Summary (Router: Base)	
Interface Name SapId/Sdp	The name of the interface or SAP/SDP identifier
Arp Populate	Specifies whether ARP populate is enabled or disabled

Label	Description
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	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
Interfaces	The total number of DHCP interfaces

Output example

```
*A:ALU-48# show router dhcp6 summary
=====
DHCP6 Summary (Router: Base)
=====
Interface Name      Nbr    Used/Max Relay   Admin  Oper Relay
SapId              Resol.  Used/Max Server  Admin  Oper Server
-----
iesSap              No      0/0              Down   Down
  sap:1/2/3:801      0/8000
iesintf              No      0/0              Down   Down
  sdp:spoke-5:9999    0/8000
-----
Interfaces: 2
=====
*A:ALU-48#
```

Table 177: DHCPv6 summary field descriptions

Label	Description
DHCP Summary (Router: Base)	
Interface Name SapId	The name of the interface or SAP/SDP identifier
Nbr Resol.	Yes – neighbor resolution (discovery) is enabled
	No – neighbor resolution (discovery) is disabled
Used/Max Relay:	Used – number of relay routes currently being used on the interface

Label	Description
	Max Relay – maximum number of relay routes on the interface
Used/Max Server	Used – number of server routes currently being used on the interface
	Max Server – maximum number of server routes currently being used on the interface
Admin	The administrative state
Oper Relay	The operating state of the relay routes
Oper Server	The operating state of the server routes
Interfaces	The total number of DHCPv6 interfaces

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 178: IP interface field descriptions](#) describes the fields.

Output example

```
*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 178: IP interface field descriptions

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

msdp

Syntax

msdp

Context

show>router

Description

This command enables the context to display MSDP information.

group

Syntax

group [*group-name*] [**detail**]

Context

show>router>msdp

Description

This command displays information about MSDP groups.

Parameters

- group-name*

displays information about the specified group. If no *group-name* is specified, information about all groups is displayed.
- detail**

displays detailed MSDP group information

Output

The following output is an example of MSDP group information, and [Table 179: MSDP group field descriptions](#) describes the fields.

Output example

```
*A:ALA-48>show router 100 msdp group
=====
MSDP Groups
=====
Group Name           Mode       Act Srcs   Local Address
-----
main                 Mesh-group None       None
loop1                Mesh-group None       None
loop2                Mesh-group None       None
loop3                Mesh-group None       None
loop4                Mesh-group None       None
loop5                Mesh-group None       None
-----
Groups : 6
=====
*A:ALA-48>show router 100 msdp

*A:ALA-48>show router 100 msdp group test
=====
MSDP Groups
=====
Group Name           Mode       Act Srcs   Local Address
-----
test                 Mesh-group 50000     10.10.10.103
-----
Groups : 1
=====
```

```

*A:ALA-48>showrouter 100 msdp#

*A:ALA-48>show>router 100 msdp# group test detail
=====
MSDP Groups
=====
Group Name          : test
-----
Local Address       : 10.10.10.103
Admin State         : Up                Receive Msg Rate   : None
Receive Msg Time    : None              Receive Msg Thd    : None
Mode                : Mesh-group        SA Limit          : 50000
Export Policy       : None Specified / Inherited
Import Policy       : None Specified / Inherited
-----
Groups : 1
=====
*A:ALA-48>show router 100 msdp

```

Table 179: MSDP group field descriptions

Label	Description
Group Name	Displays the MSDP group name
Mode	Displays the mode of peers in the group, either Mesh-group or Standard
Act Srcs	Displays the configured maximum number of SA messages that will be accepted by MSDP
Local Address	Displays the local end of an MSDP session
Admin State	Displays the administrative state
Receive Msg Rate	Displays the rate that the messages are read from the TCP session
Receive Msg Time	Displays the time interval in which the number of MSDP messages set by the receive-msdp-msg-rate <i>number</i> parameter are read from the TCP session
Receive Msg Thd	Displays the configured threshold for the number of MSDP messages that can be processed before the MSDP message rate-limiting function is activated
SA Limit	Displays the SA message limit
Export Policy	Displays whether an export policy is configured or inherited
Import Policy	Displays whether an import policy is configured or inherited

peer

Syntax

peer [*ip-address*] [*group group-name*] [*detail*]

Context

show>router>msdp

Description

This command displays information about an MSDP peer.

Parameters

- ip-address*
displays information about the peer with the specified IP address. If no IP address is specified, information about all MSDP peers is displayed.
- group-name*
displays information about peers in the specified group. If no *group-name* is specified, information about all MSDP peers display is displayed.
- detail**
displays detailed MSDP peer information

Output

The following output is an example of MSDP peer information, and [Table 180: MSDP peer field descriptions](#) describes the fields.

Output example

```
A:ALA-48# show router 100 msdp peer
=====
MSDP Peers
=====
Peer           Local Address   State           Last State Change   SA Learnt
-----
10.20.1.1      10.20.1.6      Established 08/30/2002 03:22:13   1008
-----
Peers : 1
=====
A:ALA-48#

A:ALA-48# show router 100 msdp peer detail
=====
MSDP Peers
-----
Peer Address      : 10.20.1.1
-----
Group Name        : None
Local Address     : 10.20.1.6
Last State Change : 08/30/2002 03:22:13 Last Act Src Limit : N/A
Peer Admin State  : Up           Default Peer      : No
Peer Connect Retry : 0             State            : Established
```



```

SA accepted      : 1008          SA received      : 709
State timer expires: 18          Peer time out   : 62
Active Source Limit: None        Receive Msg Rate : 0
Receive Msg Time  : 0            Receive Msg Thd  : 0
Auth Status       : Disabled      Auth Key         : None
Export Policy     : None Specified / Inherited
Import Policy     : None Specified / Inherited
-----
Peers : 1
=====
A:ALA-48#

```

Table 180: MSDP peer field descriptions

Label	Description
Peer	Displays the IP address of the peer
Local Address	Displays the local IP address
State	Displays the current state of the peer
Last State Change	Displays the date and time of the peer's last state change
SA Learnt	Displays the number of SAs learned through a peer

source

Syntax

source [*ip-address/mask*] [**type** {**configured** | **dynamic** | **both**}] [**detail**]

Context

show>router>msdp

Description

This command displays the discovery method for the specified multicast source. By default, all user-created sources are displayed.

Parameters

ip-address/mask

specifies the IP address and mask for a multicast source

configured

displays user-created sources

dynamic

displays dynamically created sources

both

displays both user-configured and dynamically created sources

detail

displays detailed MSDP source information

Output

The following output is an example of MSDP source information and [Table 181: MSDP source field descriptions](#) describes the fields.

Output example

```
*A:7705:Dut-C# show router 100 msdp source
=====
MSDP Sources
=====
Source           Type           SA Limit   Num Excd   Last Exceeded
-----
3.3.3.3/32       Configured     None        0           N/A
-----
Sources : 1
=====
*A:7705:Dut-C#
```

Table 181: MSDP source field descriptions

Label	Description
Source	Displays the IP address of the MSDP peer
Type	Displays the type of peer
SA Limit	Displays the local IP address
Num Excd	Displays the number of times the global active source limit has been exceeded
Last Exceeded	Displays the date and time of the last state change of the peer

source-active

Syntax

source-active [{**group** *ip-address* | **local** | **originator** *ip-address* | **peer** *ip-address* | **source** *ip-address* | **group** *ip-address* **source** *ip-address*}] [**detail**]

Context

show>router>msdp

Description

This command displays source-active (SA) messages accepted by MSDP.

Parameters

group *ip-address*

displays information about the specified group IP address

local

displays information about local SA messages

originator *ip-address*

displays information about the specified originator IP address

peer *ip-address*

displays information about the specified peer IP address

source *ip-address*

displays information about the specified source IP address

detail

displays detailed MSDP SA information

Output

The following output is an example of accepted MSDP SA messages information, and [Table 182: MSDP source-active field descriptions](#) describes the fields.

Output example

```
A:ALA-48# show router 100 msdp source-active
=====
MSDP Source Active Info
=====
Grp Address      Src Address      Origin RP        Peer Address     State Timer
-----
10.100.0.0       10.112.1.2      10.20.1.1       10.20.1.1       69
10.100.0.1       10.112.1.2      10.20.1.1       10.20.1.1       69
10.100.0.2       10.112.1.2      10.20.1.1       10.20.1.1       69
10.100.0.3       10.112.1.2      10.20.1.1       10.20.1.1       69
10.100.0.4       10.112.1.2      10.20.1.1       10.20.1.1       69
10.100.0.5       10.112.1.2      10.20.1.1       10.20.1.1       69
10.100.0.6       10.112.1.2      10.20.1.1       10.20.1.1       69
10.100.0.7       10.112.1.2      10.20.1.1       10.20.1.1       69
10.100.0.8       10.112.1.2      10.20.1.1       10.20.1.1       69
10.100.0.9       10.112.1.2      10.20.1.1       10.20.1.1       69
-----
MSDP Source Active : 10
=====
A:ALA-48#

A:ALA-48# show router 100 msdp source-active detail
=====
MSDP Source Active
=====
Group Address    : 10.100.0.0      Source Address    : 10.112.1.2
Origin RP       : 10.20.1.1      Peer Address     : 10.20.1.1
State Timer     : 64             Up Time         : 3d 01:44:25
Group Address    : 10.100.0.1      Source Address    : 10.112.1.2
Origin RP       : 10.20.1.1      Peer Address     : 10.20.1.1
State Timer     : 64             Up Time         : 48d 18:22:29
Group Address    : 10.100.0.2      Source Address    : 10.112.1.2
Origin RP       : 10.20.1.1      Peer Address     : 10.20.1.1
```

```

State Timer      : 64           Up Time      : 48d 18:22:29
Group Address    : 228.100.0.3  Source Address : 10.112.1.2
Origin RP        : 10.20.1.1   Peer Address  : 10.20.1.1
State Timer      : 64           Up Time      : 48d 18:22:29
Group Address    : 10.100.0.4   Source Address : 10.112.1.2
Origin RP        : 10.20.1.1   Peer Address  : 10.20.1.1
State Timer      : 64           Up Time      : 48d 18:22:29
Group Address    : 10.100.0.5   Source Address : 10.112.1.2
Origin RP        : 10.20.1.1   Peer Address  : 10.20.1.1
State Timer      : 64           Up Time      : 48d 18:22:29
Group Address    : 10.100.0.6   Source Address : 10.112.1.2
Origin RP        : 10.20.1.1   Peer Address  : 10.20.1.1
State Timer      : 64           Up Time      : 48d 18:22:29
Group Address    : 10.100.0.7   Source Address : 10.112.1.2
Origin RP        : 10.20.1.1   Peer Address  : 10.20.1.1
State Timer      : 64           Up Time      : 48d 18:22:29
Group Address    : 10.100.0.8   Source Address : 10.112.1.2
Origin RP        : 10.20.1.1   Peer Address  : 10.20.1.1
State Timer      : 64           Up Time      : 48d 18:22:29
Group Address    : 10.100.0.9   Source Address : 10.112.1.2
Origin RP        : 10.20.1.1   Peer Address  : 10.20.1.1
State Timer      : 64           Up Time      : 48d 18:22:29
-----
MSDP Source Active : 10
=====
A:ALA-48#

```

Table 182: MSDP source-active field descriptions

Label	Description
Grp Address	Displays the IP address of the group
Src Address	Displays the IP address of the source
Origin RP	Displays the originating rendezvous point (RP) address
Peer Address	Displays the IP address of the peer
State Timer	Displays the state timeout value. If the value reaches 0, the SA entry is removed.

source-active-rejected

Syntax

source-active-rejected [**peer-group** *name*] [**group** *ip-address*] [**source** *ip-address*] [**originator** *ip-address*] [**peer** *ip-address*]

Context

show>router>msdp

Description

This command displays SA messages rejected by MSDP.

Parameters

name

displays information about rejected SA messages for the specified peer group

group ip-address

displays information about the specified group IP address

source ip-address

displays information about the source address of the source-active entry that is rejected

originator ip-address

displays information about the specified originator IP address

peer ip-address

displays information about the peer from which this rejected source-active entry was last received

Output

The following output is an example of rejected MSDP SA messages information, and [Table 183: MSDP source-active-rejected field descriptions](#) describes the fields.

Output example

```
*A:ALA-48# show router 100 msdp source-active-rejected
=====
MSDP Source Active Rejected Info
=====
Grp Address      Src Address      Origin RP        Peer Address     Reject Reason
-----
10.100.0.1       10.0.0.1         10.20.0.1       10.0.0.1        Import Policy
10.100.0.2       10.0.0.2         10.20.0.2       10.0.0.2        Export Policy
10.100.0.3       10.0.0.3         10.20.0.3       10.0.0.3        RPF Failure
10.100.0.4       10.0.0.4         10.20.0.4       10.0.0.4        Limit Exceeded
10.100.0.5       10.0.0.5         10.20.0.5       10.0.0.5        Limit Exceeded
10.100.0.6       10.0.0.6         10.20.0.6       10.0.0.6        Limit Exceeded
10.100.0.7       10.0.0.7         10.20.0.7       10.0.0.7        Limit Exceeded
-----
SA Rejected Entries : 7
=====
*A:ALA-48#
```

Table 183: MSDP source-active-rejected field descriptions

Label	Description
Grp Address	Displays the IP address of the group
Src Address	Displays the IP address of the source
Origin RP	Displays the originating rendezvous point (RP) address
Peer Address	Displays the address of the peer
Reject Reason	Displays the reason why this SA entry is rejected

statistics

Syntax

statistics [**peer** *ip-address*]

Context

show>router>msdp

Description

This command displays statistics information related to an MSDP peer.

Parameters

ip-address
displays statistics for the peer with the specified IP address

Output

The following output is an example of MSDP statistics information, and [Table 184: MSDP statistics field descriptions](#) describes the fields.

Output example

```
A:ALA-48# show router 100 msdp statistics
=====
MSDP Statistics
=====
Glo ActSrc Lim Excd: 0
-----
Peer Address      : 10.20.1.1
-----
Last State Change : 0d 11:33:16      Last message Peer : 0d 00:00:17
RPF Failures      : 0                Remote Closes    : 0
SA Msgs Sent      : 0                SA Msgs Recvd     : 709
SA req. Msgs Sent : 0                SA req. Msgs Recvd : 0
SA res. Msgs Sent : 0                SA res. Msgs Recvd : 0
KeepAlive Msgs Sent: 694             KeepAlive Msgs Recd: 694
Unknown Msgs Sent : 0                Error Msgs Recvd  : 0
-----
Peers : 1
=====
A:ALA-48#
```

Table 184: MSDP statistics field descriptions

Label	Description
Glo ActSrc Lim Excd	Displays the number of global active source messages that exceed the configured limit
Peer Address	Displays the address of the MSDP peer
Last State Change	Displays the date and time the peer state changed

Label	Description
Last message Peer	Displays the time the last message was received from the peer
RPF Failures	Displays the number of reverse path forwarding (RPF) failures
Remote Closes	Displays the number of times the remote peer closed
SA Msgs Sent	Displays the number of SA messages sent
SA Msgs Recvd	Displays the number of SA messages received
SA req. Msgs Sent	Displays the number of SA request messages sent
SA req. Msgs Recvd	Displays the number of SA request messages received
SA res. Msgs Sent	Displays the number of SA response messages sent
SA res. Msgs Recvd	Displays the number of SA response messages received
KeepAlive Msgs Sent	Displays the number of keepalive messages sent
KeepAlive Msgs Recd	Displays the number of keepalive messages received
Unknown Msgs Sent	Displays the number of unknown messages sent
Error Msgs Recvd	Displays the number of error messages received

status

Syntax

status

Context

show>router>msdp

Description

This command displays MSDP status information.

Output

The following output is an example of MSDP status information, and [Table 185: MSDP status field descriptions](#) describes the fields.

Output example

```
A:ALA-48# show router 100 msdp status
=====
MSDP Status
=====
```

```

Admin State           : Up
Local Address         : None
Global Statistics
Active Src Limit      : None
Act Src Lim Excd      : 0
Num. Peers            : 1
Num. Peers Estab      : 1
Num. Source Active    : 10
Policies              : None
Data Encapsulation    : Enabled
Receive Msg Rate      :
Rate                  : 0
Time                  : 0
Threshold             : 0
Last Msdp Enabled     : 08/30/2002 03:21:43
=====

```

```

A:ALA-48#

```

Table 185: MSDP status field descriptions

Label	Description
Admin State	Displays the administrative state
Local Address	Displays the local IP address
Global Statistics	Displays global MSDP statistics
Active Src Limit	Displays the active source limit
Act Src Lim Excd	Displays the number of times that the active source limit was exceeded
Num. Peers	Displays the number of peers
Num. Peers Estab	Displays the number of peers established
Num. Source Active	Displays the number of active sources
Policies	Specifies the policy used to export the SA state from the SA list into MSDP
Data Encapsulation	Specifies whether the rendezvous point (RP) encapsulates multicast data received in MSDP register messages inside forwarded MSDP SA messages
Rate	The receive message rate
Time	The receive message interval
Threshold	The number of MSDP messages that can be processed before the MSDP message rate-limiting function is activated
Last Msdp Enabled	The time the last MSDP was triggered

route-table

Syntax

route-table [*family*] [*ip-prefix*[/*prefix-length*]] [**longer** | **exact** | **protocol** *protocol-name*] [**all**]] [**next-hop-type** *type*] [**alternative**]

route-table [*family*] **summary**

route-table [*family*] [*ip-prefix*[/*prefix-length*]] [**longer** | **exact** | **protocol** *protocol-name*] extensive [**all**]

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.

The following adapter cards and platforms support the full IPv6 subnet range for IPv6 static routes:

- 6-port Ethernet 10Gbps Adapter card
- 8-port Gigabit Ethernet Adapter card, version 2 and version 3
- 2-port 10GigE (Ethernet) Adapter card (on the v-port)
- 10-port 1GigE/1-port 10GigE X-Adapter card
- 7705 SAR-X

For these cards and platforms, the supported route range for statically provisioned or dynamically learned routes is from /1 to /128.

For all other cards, modules, and ports (including the v-port on the 2-port 10GigE (Ethernet) module), the supported range for statically provisioned or dynamically learned routes is from /1 to /64 or is /128 (indicating a host route).

Parameters

family

specifies the type of routing information to be distributed by this peer group

Values	ipv4 – displays the routes that have the IPv4 family enabled, excluding IP-VPN routes ipv6 – displays the routes that are IPv6-capable, including IPv6 static routes mcast-ipv4 – displays the routes that are IPv4 multicast-capable mcast-ipv6 – displays the routes that are IPv6 multicast-capable
---------------	---

ip-prefix/prefix-length

displays only those entries matching the specified IP prefix and prefix length

longer

displays routes matching the *ip-prefix/prefix-length* and routes with longer masks

- exact**
displays the exact route matching the *ip-prefix/prefix-length* masks
- protocol-name*
displays routes learned from the specified protocol
Values bgp, bgp-vpn, isis, local, ospf, rip, static, aggregate, vpn-leak, managed
- all**
displays all routes, including inactive routes
- type*
displays tunneled next-hop information
- alternative**
displays LFA and backup route details
- extensive**
displays extensive route table information
- summary**
displays route table summary information

Output

The following outputs are examples of routing table information:

- standard route table information ([Output example, Table 186: Standard route table field descriptions](#))
- extensive route table information ([Output example, Table 187: Route table extensive field descriptions](#))
- LFA and backup route table information ([Output example, Table 188: Route table alternative field descriptions](#))

Output example

```
*A:ALU# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags]                Type   Proto   Age           Pref
  Next Hop[Interface Name]                Metric
-----
0.0.0.0/0                          Remote Static  00h00m03s    5
    upLink                          1
10.1.1.1/32                        Local  Local   35d08h00m    0
    system                          0
-----
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
      Backup = BGP backup route
      LFA = Loop-Free Alternate nexthop
      S = Sticky ECMP requested
=====
*A:Sar18 Dut-B>show>router#

*A:ALU-A# show router route-table protocol ospf
=====
Route Table (Router: Base)
=====
```

```

Dest Prefix[Flags]
Next Hop[Interface Name]
-----
10.10.0.1/32
10.10.13.1
-----
Type      Proto  Age    Pref
Metric
-----
Remote   OSPF   65844   10
0
-----
Flags: n = Number of times nexthop is repeated
Backup = BGP backup route
LFA = Loop-Free Alternate nexthop
S = Sticky ECMP requested
=====
*A: Sar18 Dut-B>show>router#

```

Table 186: Standard route table field descriptions

Label	Description
Dest Prefix	The route destination address and mask
[Flags]	n – number of times nexthop is repeated Backup – BGP backup route LFA – loop-free alternate next hop S – sticky ECMP requested
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
Protocol	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
[Flags]	n – number of times nexthop is repeated Backup – BGP backup route LFA – loop-free alternate next hop S – sticky ECMP requested

Output example

```

*A:7705:Dut-C# show router 10 route-table 200.0.0.1 32 extensive
=====
Route Table (Service: 10)
=====
Dest Prefix      : 200.0.0.1/32
Protocol        : OSPF
Age             : 00h01m53s

```

```

Preference      : 150
Next-Hop        : 10.1.1.1
Interface       : iftoA-1
QoS             : Priority=n/c, FC=n/c
Source-Class    : 0
Dest-Class      : 0
Metric          : 1001
ECMP-Weight     : 1
Next-Hop        : 10.1.2.1
Interface       : iftoA-2
QoS             : Priority=n/c, FC=n/c
Source-Class    : 0
Dest-Class      : 0
Metric          : 1001
ECMP-Weight     : 1
Next-Hop        : 10.1.3.1
Interface       : iftoA-3
QoS             : Priority=n/c, FC=n/c
Source-Class    : 0
Dest-Class      : 0
Metric          : 1001
ECMP-Weight     : 5
-----
No. of Destinations: 1
=====
*A:7705:Dut-C#

```

Table 187: Route table extensive field descriptions

Label	Description
Dest Prefix	The destination prefix of the LSP tunnel
Protocol	The routing protocol used by the route table
Age	The age of the LSP tunnel
Preference	The route preference value
Next-Hop	The next-hop address for the route destination
Interface	The next-hop interface name
QoS	The next-hop QoS type
Source-Class	The next-hop source class value
Dest-Class	The next-hop destination class value
Metric	The next-hop metric value
ECMP-Weight	The next-hop ECMP weight value
No. of Destinations	The total number of next-hop destinations

Output example

```

*A:ALU# show router route-table alternative
=====

```

```

Route Table (Router: Base)
=====
Dest Prefix[Flags]      Type  Proto  Age      Metric  Pref
  Next Hop[Interface Name]
    Alt-NextHop
    Alt-
    Metric
-----
10.10.1.0/24            Local  Local  00h07m52s  0
    ip-10.10.1.1
10.10.2.0/24            Local  Local  00h07m48s  0
    ip-10.10.2.1
10.10.4.0/24            Remote  ISIS   00h07m38s  15
    10.10.1.2
10.10.5.0/24            Remote  ISIS   00h07m38s  15
    10.10.2.3
10.10.9.0/24            Remote  ISIS   00h07m28s  15
    10.10.1.2
    10.20.1.5 (LFA) (tunneled:RSVP:3)
10.10.10.0/24           Remote  ISIS   00h04m40s  15
    10.20.1.5 (tunneled:RSVP:3)
10.20.1.1/32            Local  Local  00h07m55s  0
    system
10.20.1.2/32            Remote  ISIS   00h07m47s  15
    10.10.1.2
10.20.1.3/32            Remote  ISIS   00h07m38s  15
    10.10.2.3
10.20.1.4/32            Remote  ISIS   00h07m38s  15
    10.10.1.2
    10.20.1.5 (LFA) (tunneled:RSVP:3)
10.20.1.5/32            Remote  ISIS   00h04m40s  15
    10.20.1.5 (tunneled:RSVP:3)
10.20.1.6/32            Remote  ISIS   00h07m28s  15
    10.10.1.2
    10.10.2.3 (LFA)
-----
No. of Routes: 12
Flags: n = Number of times nexthop is repeated
      Backup = BGP backup route
      LFA = Loop-Free Alternate nexthop
      S = Sticky ECMP requested
=====
*A:ALU-A#

```

Table 188: Route table alternative field descriptions

Label	Description
Dest Prefix[Flags]	The route destination address and mask, and flags (if applicable)
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

Label	Description
Pref	The route preference value for the route
No. of Routes	The number of routes displayed in the list
Alt-NextHop	The backup next hop
Alt-Metric	The metric of the backup route

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, cflowd, dhcp, dns, ftp, icmp, igmp, isis, ldp, mld, ndis, ntp, ospf, pim, ptp, radius, rip, rsvp, snmp, snmp-notification, ssh, syslog, tacplus, telnet, tftp, traceroute, vrrp



- 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 189: Application QoS field descriptions](#) describes the fields.

Output example

```
*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
igmp	none	none
ldp	none	none
mld	none	none
ndis	none	none
ntp	none	none
ospf	none	none
pim	none	none
ptp	none	none
radius	none	none
rip	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
vrrp	none	none
=====		
Dot1p Application Values		
=====		
Application	Dot1p Value	Default Dot1p Value

arp	none	none
isis	none	none
=====		
*A:ALU-1>show#		

Table 189: Application QoS field descriptions

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 190: DSCP-to-FC mapping field descriptions](#) describes the fields.

Output example

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 190: DSCP-to-FC mapping field descriptions

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 191: Static ARP table field descriptions](#) describes the fields.

Output example

```
*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
=====
```

Table 191: Static ARP table field descriptions

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 [*family*] [*ip-prefix/prefix-length*] | **preference** *preference* | **next-hop** *ip-address* | **tag** *tag* [**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.

The following adapter cards and platforms support the full IPv6 subnet range for IPv6 static routes:

- 6-port Ethernet 10Gbps Adapter card
- 8-port Gigabit Ethernet Adapter card, version 2 and version 3
- 2-port 10GigE (Ethernet) Adapter card (on the v-port)
- 10-port 1GigE/1-port 10GigE X-Adapter card
- 7705 SAR-X

For these cards and platforms, the supported route range for statically provisioned or dynamically learned routes is from /1 to /128.

For all other cards, modules, and ports (including the v-port on the 2-port 10GigE (Ethernet) module), the supported range for statically provisioned or dynamically learned routes is from /1 to /64 or is /128 (indicating a host route).

Parameters

- family

displays only the static routes with the specified family

Values ipv4, ipv6, mcast-ipv4, or mcast-ipv6
- ip-prefix/prefix-length

displays only the static routes matching the specified IP prefix and prefix length
- preference

displays static routes with the specified route preference

Values 0 to 65535
- ip-address

only displays static routes with the specified next-hop IP address
- tag

displays the 32-bit integer tag added to the static route. The tag is used in route policies to control distribution of the route into other protocols.

Values 0 to 4294967295
- detail

displays detailed information about the static route

Output

The following output is an example of static route information, and [Table 192: Static route field descriptions](#) describes the fields.

Output example

```
*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	n/a				
No. of Static Routes: 5					
*A:ALU-12#					

*A:Ser18 Dut-B>show>router# static-route detail

Static Route Table (Service: 12) Family: IPv4

Prefix	: 10.1.1.1/32		
Nexthop	: 10.1.1.1		
Type	: Nexthop		
Interface	: toA	Active	: Y
Prefix List	: n/a	Prefix List Type	: n/a
Metric	: 1	Preference	: 5
Admin State	: Up	Tag	: 0
Creation Origin	: manual		
BFD	: disabled		
CPE-check	: disabled		
LDP Sync	: disabled		
Prefix	: 10.1.1.3/32		
Nexthop	: 10.1.1.2		
Type	: Indirect		
Interface	: n/a	Active	: N
Prefix List	: n/a	Prefix List Type	: n/a
Metric	: 44	Preference	: 22
Admin State	: Down	Tag	: 55
Creation Origin	: manual		
BFD	: disabled		
CPE-check	: disabled		
Tunnel Resolution	: disabled	Disallow-IGP	: disabled
RSVP-TE Tunnels	: disabled	LDP Tunnels	: disabled
SR-ISIS Tunnels	: disabled	SR-OSPF Tunnels	: disabled
SR-TE Tunnels	: disabled		
Inactive Reason	: VPRN admin disabled		
Prefix	: 10.1.1.3/32		
Nexthop	: n/a		
Type	: Blackhole		
Interface	: n/a	Active	: N
Prefix List	: n/a	Prefix List Type	: n/a
Metric	: 1	Preference	: 5
Admin State	: Down	Tag	: 0
Creation Origin	: manual		
BFD	: disabled		
CPE-check	: disabled		
Inactive Reason	: VPRN admin disabled		
Prefix	: 10.1.1.3/32		
Nexthop	: 10.1.1.1		
Type	: Indirect		
Interface	: n/a	Active	: N
Prefix List	: n/a	Prefix List Type	: n/a
Metric	: 1	Preference	: 5
Admin State	: Down	Tag	: 0
Creation Origin	: manual		

```

BFD : disabled
CPE-check : disabled
Tunnel Resolution: disabled
RSVP-TE Tunnels : disabled
SR-ISIS Tunnels : disabled
SR-TE Tunnels : disabled
Inactive Reason : VPRN admin disabled
-----
Prefix : 10.1.1.3/32
Nexthop : n/a
Type : Nexthop
Interface : toSar8
Prefix List : n/a
Metric : 1
Admin State : Down
Creation Origin : manual
BFD : disabled
CPE-check : disabled
LDP Sync : disabled
Inactive Reason : VPRN admin disabled
-----
No. of Static Routes: 5
=====
Disallow-IGP : disabled
LDP Tunnels : disabled
SR-OSPF Tunnels : disabled
Active : N
Prefix List Type : n/a
Preference : 5
Tag : 0
=====
*A: Sar18 Dut-B>show>router#

```

Table 192: Static route field descriptions

Label	Description
Prefix	The static route destination address and mask
Next Hop Nexthop	The next hop for the static route destination
Tag	The 32-bit integer tag added to the static route
Met Metric	The route metric value for the static route
Pref Preference	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 GRT – the route is a static route for the GRT next hop IPSec – the route is a static route for the IPSec tunnel next hop
Act Active	N – the static route is inactive; for example, the static route is disabled or the next-hop IP interface is down

Label	Description
	Y – the static route is active
Interface	The egress IP interface name for the static route n/a – indicates there is no current egress interface because the static route is inactive or a blackhole route
Prefix List	Identifies the prefix list used for this static route
Prefix List Type	Identifies the type of prefix list used for this static route
Admin State	The administrative state for this static route
Creation Origin	The method by which the static route was created: manual or automatic (dynamic)
BFD	The BFD state for this static route (enabled or disabled)
CPE-check	The configured state of CPE check for this static route (enabled or disabled)
Tunnel Resolution	n/a
Disallow-IGP	n/a
RSVP-TE Tunnels	n/a
LDP Tunnels	n/a
SR-ISIS Tunnels	n/a
SR-OSPF Tunnels	n/a
SR-TE Tunnels	n/a
Inactive Reason	Indicates the reason for the static route being inactive
No. of Static Routes:	The number of static routes displayed in the list

tunnel-table

Syntax

tunnel-table summary [ipv4 | ipv6]

tunnel-table [protocol protocol] {ipv4 | ipv6}

tunnel-table [ip-prefix[/mask]] [alternative] [ipv4 | ipv6] detail

tunnel-table [ip-prefix[/mask]] [alternative]

tunnel-table [ip-prefix[/mask]] protocol protocol [detail]

tunnel-table [ip-prefix[/mask]] sdp sdp-id

Context

show>router

Description

This command displays tunnel table information.

If the **auto-bind-tunnel** command is used when configuring a VPRN service, it means that the MP-BGP next-hop resolution is referring to the core routing instance for IP reachability. For a VPRN service, the next hop specifies the lookup to be used by the routing instance if no SDP to the destination exists.

Parameters

ip-prefix[/mask]
displays the specified tunnel table destination IP address and mask

protocol
displays protocol information

Values bgp, ldp, rsvp, sdp, ospf, isis, sr-te, fpe

sdp-id
displays information pertaining to the specified SDP

Values 1 to 17407

summary
displays summary tunnel table information

detail
displays detailed tunnel table information

alternative
displays backup route details

ipv4
displays information for IPv4 entries only

ipv6
displays information for IPv6 entries only

Output

The following output is an example of tunnel table information, and [Table 193: Tunnel table field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B>show>router# tunnel-table summary
```

Tunnel Table Summary (Router: Base)		
	Active	Available
LDP	1	1
SDP	1	1
RSVP	0	0


```

BGP                                0                                0
MPLS-TP                           0                                0
ISIS                              0                                0
OSPF                              0                                0
SR-TE                             0                                0
FPE                               0                                0
-----
Total                             2                                2
=====
*A: Sarl8 Dut-B>show>router#

```

```

A: Sarl8 Dut-B>show>router# tunnel-table
=====
IPv4 Tunnel Table (Router: Base)
=====
Destination      Owner      Encap TunnelId  Pref    Nexthop      Metric
-----
10.1.1.1/32      sdp        MPLS  1000        5       10.1.1.1      0
10.1.1.1/32      ldp        MPLS  65537       9       10.1.1.1      1
-----
Flags: B = BGP backup route available
      E = inactive best-external BGP route
=====
*A: Sarl8 Dut-B>show>router#

```

```

*A: Sarl8 Dut-B>show>router# tunnel-table detail
=====
Tunnel Table (Router: Base)
=====
Destination      : 10.1.1.1/32
NextHop          : 10.1.1.1
Tunnel Flags     : (Not Specified)
Age              : 26d21h16m
CBF Classes      : (Not Specified)
Owner            : sdp                      Encap           : MPLS
Tunnel ID        : 1000                    Preference      : 5
Tunnel Label     : -                      Tunnel Metric    : 0
Tunnel MTU       : 1546                   Max Label Stack : 1
-----
Destination      : 10.1.1.1/32
NextHop          : 10.1.1.1
Tunnel Flags     : (Not Specified)
Age              : 26d21h16m
CBF Classes      : (Not Specified)
Owner            : ldp                      Encap           : MPLS
Tunnel ID        : 65537                   Preference      : 9
Tunnel Label     : 131071                  Tunnel Metric    : 1
Tunnel MTU       : 1550                   Max Label Stack : 1
-----
Number of tunnel-table entries      : 2
Number of tunnel-table entries with LFA : 0
=====
*A: Sarl8 Dut-B>show>router#

```

```

*A: Sarl8 Dut-B>show>router# tunnel-table ipv6 protocol isis
=====
IPv6 Tunnel Table (Router: Base)
=====
Destination      Owner      Encap TunnelId  Pref
Nexthop                               Metric
-----
No Matching Entries.

```

```
-----
Flags: B = BGP backup route available
      E = inactive best-external BGP route
=====
```

```
*A: Sar18 Dut-B>show>router#
```

Table 193: Tunnel table field descriptions

Label	Description
Destination	The route destination address and mask
Owner	The tunnel owner (protocol)
Encap	The tunnel encapsulation type
Tunnel ID	The tunnel (SDP) identifier
Pref Preference	The route preference for routes learned from the configured peers
Nexthop	The next hop for the route destination
Metric	The route metric value for the route
CBF Classes	n/a
Tunnel Flags	The tunnel flags
Tunnel Label	The tunnel label
Tunnel Metric	The tunnel metric
Tunnel MTU	The tunnel MTU
Max Label Stack	The maximum label stack depth
Age	The tunnel age (that is, how long the tunnel has been operational)

7.8.2.4 Clear service commands

msdp

Syntax

msdp

Context

clear>router

Description

This command enables the context to clear and reset Multicast Source Discovery Protocol (MSDP) entities and statistics.

cache

Syntax

cache [**peer** *ip-address*] [**group** *ip-address*] [**source** *ip-address*] [**originrp** *ip-address*]

Context

clear>router>msdp

Description

This command clears IP addresses from the MSDP cache.

Parameters

peer *ip-address*

clears the specified peer address

group *ip-address*

clears the specified group address

source *ip-address*

clears the specified source address

originrp *ip-address*

clears the specified originating rendezvous point (RP) address

statistics

Syntax

statistics [**peer** *ip-address*]

Context

clear>router>msdp

Description

This command clears IP address statistics for the peer to which MSDP SA requests for groups matching this entry's group range were sent.

Parameters

ip-address

clears the MSDP statistics for the specified IP address

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 44: SAP ID configurations](#) 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

ip-transport

Syntax

ip-transport *ipt-id*

Context

clear>service>id

Description

This command clears configured information pertaining to a specified IP transport subservice.

If no port identifier is specified, information is cleared for all IP transport subservices.

Parameters

ipt-id

the IP transport subservice physical port identifier, in the format *slot/mda/port.channel*

remote-host

Syntax

remote-host *host-id*

Context

clear>service>id>ip-transport

Description

This command clears configured information pertaining to a specified remote host assigned to this IP transport subservice.

Parameters

host-id

the remote host identifier

Values 1 to 2147483647 or a name string up to 64 characters long

statistics**Syntax**

statistics

Context

clear>service>id>ip-transport

clear>service>id>ip-transport>remote-host

Description

This command clears statistics-related information pertaining to all configured IP transport subservices or to all configured remote hosts for a specified IP transport subservice.

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
 - mac**
clears MFIB statistics for the specified MAC address
 - grp-address*
clears MFIB statistics for the specified group IP or MAC address

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 44: SAP ID configurations](#) 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 *svclId-id:groupName* {**all** | **counters**}

Context

clear>service>statistics

Description

This command clears statistics for a specified SAP aggregation group.

Parameters

svclId-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

7.8.2.5 Debug MSDP commands

msdp

Syntax

[no] msdp

Context

debug>router

Description

This command enables debugging for Multicast Source Discovery Protocol (MSDP).

The **no** form of the command disables MSDP debugging.

packet

Syntax

packet [*pkt-type*] [**peer** *ip-address*]

Context

debug>router>msdp

Description

This command enables debugging for MSDP packets.

The **no** form of the command disables MSDP packet debugging.

Parameters

pkt-type

debugs information associated with the specified packet type

Values keep-alive, source-active, sa-request, sa-response

ip-address

debugs information associated with the specified peer IP address

pim

Syntax

pim [*grp-address*]

no pim

Context

debug>router>msdp

Description

This command enables debugging for MSDP PIM.

The **no** form of the command disables MSDP PIM debugging.

Parameters

grp-address

debugs the IP multicast group address for which this entry contains information

rtm

Syntax

rtm [*rp-address*]

no rtm

Context

debug>router>msdp

Description

This command enables debugging for MSDP route table manager (RTM).

The **no** form of the command disables MSDP RTM debugging.

Parameters

rp-address

debugs the IP multicast address for which this entry contains information

sa-db

Syntax

sa-db [**group** *grpAddr*] [**source** *srcAddr*] [**rp** *rpAddr*]

no sa-db

Context

debug>router>msdp

Description

This command enables debugging for MSDP source-active (SA) requests.
The **no** form of the command disables the MSDP SA database debugging.

Parameters

- grpAddr*
debugs the IP address of the group
- srcAddr*
debugs the source IP address
- rpAddr*
debugs the specified rendezvous point RP address

7.8.2.6 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

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 44: SAP ID configurations](#) 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 44: SAP ID configurations](#) 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

8 IPSec

This chapter provides information to configure security parameters.

Topics in this chapter include:

- [IPSec overview](#)
- [Public key infrastructure](#)
- [IPSec best practices recommendations](#)
- [Configuration notes](#)
- [Configuring IPSec with CLI](#)
- [IPSec command reference](#)

8.1 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](#)
- [OCSP](#)
- [Applications](#)
- [NAT-traversal for IKEv1/v2 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](#)
- [OSPFv3 packet authentication with IPv6 IPSec](#)
- [Network security with IPv6 IPSec](#)
- [IPSec over r-VPLS on a public-side service](#)

- [Statistics](#)
- [Security support](#)

8.1.1 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](#)
- [IKEv2 fragmentation](#)
- [Tunnel group](#)
- [Tunnel interfaces and SAPs](#)
- [IPSec tunnel configuration](#)
- [IPSec over MPLS with public-side IES](#)
- [IPSec transport tunnels with public-side VPRN](#)
- [GRE-encapsulated VLLs/VPLS over IPSec VPNs](#)
- [GRE-encapsulated VLLs/VPLS over IPSec over MPLS](#)

8.1.1.1 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 or VPRN for the public side of the tunnel. In the following figure, a public service instance (IES, VPRN, or network) 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.

Figure 128: IPSec implementation architecture



24019

In the figure, all ingress customer traffic from the trusted network is aggregated into the private VPNR service, where a VPNR 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, VPNR, or router interface that is configured on an encryption-capable adapter card; it gets routed to its destination via a standard FIB lookup.

IPSec traffic ingressing a public-side VPNR that is not configured for IPSec is dropped.

If the traffic passes all security checks, it is decrypted and the customer traffic is routed through the associated VPNR. Any traffic that does not match the tunnel security configuration is dropped.

8.1.1.2 Hardware support

The 7705 SAR supports IPSec on the following nodes and adapter cards:

- 7705 SAR-8 Shelf V2 or 7705 SAR-18 with one of the following:
 - 2-port 10GigE (Ethernet) Adapter card
 - 6-port Ethernet 10Gbps Adapter card
 - 8-port Gigabit Ethernet Adapter card, version 3
 - 10-port 1GigE/1-port 10GigE X-Adapter card, version 2 (supported on the 7705 SAR-18 only)
- 7705 SAR-Ax

- 7705 SAR-H
- 7705 SAR-Hc
- 7705 SAR-Wx
- 7705 SAR-X



Note: On the 7705 SAR-X, each Ethernet port has its own encryption engine.

8.1.1.3 IPSec encryption features

IPSec provides a variety of encryption features required to establish bidirectional IPSec tunnels, including:

- control plane:
 - manual keying
 - dynamic keying: Internet key exchange version 1, version 2 (IKEv1, IKEv2)
 - IKEv1 mode: main or aggressive
 - authentication: pre-shared-key (PSK)
 - perfect forward secrecy (PFS)
 - dead peer detection (DPD)
 - NAT-traversal (NAT-T)
 - security policy
- data plane:
 - encapsulating security payload (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, auth-encryption
 - encryption algorithm: NULL, DES, 3DES, AES128, AES192, AES256, AES128-GCM8, AES128-GCM12, AES128-GCM16, AES192-GCM8, AES192-GCM12, AES192-GCM16, AES256-GCM8, AES256-GCM12, AES256-GCM16
 - IPSec IKE policy (NULL is not supported):
 - authentication algorithm: MD5, SHA1, SHA256, SHA384, SHA512, auth-encryption
 - encryption algorithm: DES, 3DES, AES128, AES192, AES256, AES128-GCM8, AES128-GCM16, AES256-GCM8, AES256-GCM16
 - PRF algorithm: MD5, SHA1, SHA256, SHA384, SHA512, AES-XCBC, same-as-auth
 - 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).
- When using AES-GCM encryption algorithms, use either the default value for the **isakmp-lifetime** command or a value that is lower than the default. This ensures that the CHILD_SA lifetime is refreshed for AES-GCM at least once a day.

The 7705 SAR supports the use of IPSec and segment routing with entropy label for:

- IPSec over BGP 3107 over segment routing with entropy label
- IPSec over static route over segment routing with entropy label
- VLL over GRE over IPSec over BGP 3107 over segment routing with entropy label
- VLL over GRE over IPSec over static route over segment routing with entropy label

8.1.1.4 SHA2 support

The 7705 SAR supports RFC 4868. For data origin authentication and integrity verification functions in the IKEv1 or IKEv2 and 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 IKEv1 or 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

8.1.1.5 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.

8.1.1.6 IKEv2 fragmentation

IKEv2 uses UDP as the transport protocol for its messages. Most IKEv2 messages are relatively small. In some cases, though, an IKEv2 message can be large; for example, an IKE_AUTH message with a certificate payload. If the IKEv2 message size exceeds the network path MTU, it gets fragmented at the IP level into smaller IP fragments. However, some devices (such as firewalls) do not allow IP fragments to pass through the network. If the fragments do not pass through, IKE negotiation fails.

To address this problem, the 7705 SAR supports IKEv2 fragmentation, as specified in RFC 7383. With IKEv2 fragmentation, IKEv2 messages are fragmented at the IKEv2 protocol level into smaller messages. The resulting IP packets are smaller than the network path MTU and are therefore not fragmented through the network and can traverse network devices that do not allow IP fragments to pass through.

IKEv2 fragmentation is enabled in the **ike-policy** context by configuring the **ikev2-fragment** command with an MTU. The MTU specified is the maximum size of the IKEv2 packet.

The system enables IKEv2 fragmentation for a tunnel only if the **ikev2-fragment** command is configured and if the peer also announces its support by sending an IKEV2_FRAGMENTATION_SUPPORTED notification.

8.1.1.7 Tunnel group

A tunnel group is a collection of IPSec tunnels. The 7705 SAR supports one tunnel group that always uses tunnel ID 1.

8.1.1.8 Tunnel interfaces and SAPs

There are two types of tunnel interfaces and associated SAPs:

- public tunnel interface: configured on the public-side IES or public-side VPRN 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 on the private-side VPRN service
 - private tunnel SAP: associated with the private tunnel interface, logically linked to the public tunnel SAP

8.1.1.8.1 Public tunnel SAPs

An IES or VPRN 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)
- IKEv1/v2 (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.

8.1.1.8.2 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.

8.1.1.8.3 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 before 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).

8.1.1.9 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
- local gateway address
- dynamic keying
- IKE policy
- pre-shared key
- transform

The local gateway address must belong to the same subnet as the delivery-service (IES or VPRN) 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.

8.1.1.10 IPSec over MPLS with public-side IES

IPSec messages can be routed over MPLS tunneled routes. The 7705 SAR supports resolution of IPSec routes to the secure gateway address by using either BGP 3107 label routes or IGP shortcuts. When BGP learns IPv4 addresses as 3107 label routes, BGP resolves the next hops for these routes with an LDP or RSVP-TE tunnel. These BGP routes create BGP tunnels that can be used to resolve an IPSec secure gateway address. When an IGP shortcut is enabled on the 7705 SAR by using the **config>router>ospf>rsvp-shortcut** command, OSPF installs an OSPF route in the RIB, with an RSVP-TE LSP as the next hop. If this OSPF route is determined as the overall best route, then the next hop is an RSVP-TE tunnel. For information about setting up BGP 3107 label routes or IGP shortcuts to resolve IPSec routes, see [Configuring IPSec over MPLS](#).

8.1.1.11 IPSec transport tunnels with public-side VPRN

Configuring VPRN as the public-side service of an IPSec tunnel ensures that IPSec traffic from different customers arriving at the 7705 SAR is kept separate. Keeping the traffic from different customers separated gives service providers another layer of security because a specific VPRN is assigned to a

specific customer, and only the IPSec tunnels encrypted by that customer arrive on that VPRN. In contrast, when IES is used for the public-side service of an IPSec tunnel, all IPSec traffic arrives in the GRT and fans out to its corresponding private service through the IPSec public gateway.


When the public-side service of an IPSec tunnel is a VPRN, IPSec traffic is transported over MPLS or GRE. The 7705 SAR supports the transport of GRE-encapsulated VLLs (Cpipes and Epipes) over IPSec tunnels that use VPRN as the public-side service of the IPSec tunnel.

MP-BGP services can be provisioned using autobind tunnels or SDPs to push IPSec packets over MPLS or GRE transport tunnels.

An MP-BGP VC label is pushed on top of the IPSec packet and a transport tunnel label is pushed next. The transport tunnel can include BGP-labeled unicast (BGP-LU) routes (3107 label routes).

The following figure shows the concept of a VPRN as the public-side service of an IPSec tunnel.

Figure 129: VPRN public service IPSec transport tunnels



27826

8.1.1.12 GRE-encapsulated VLLs/VPLS over IPSec VPNs

The 7705 SAR can provide secure transport of Cpipe, Epipe, and VPLS traffic by routing it as GRE-encapsulated traffic over IPSec VPNs. This is achieved by enabling route processing in the GRT FIB through a GRT lookup at ingress to the VRF and GRT leaking at egress from the VRF. The 7705 SAR leaks only IPSec tunnels into the GRT as the available next hop; no other tunnel type is leaked from the VRF into the GRT as a next hop. This route processing is enabled with the **config>service>vprn>grt-lookup>enable-grt** command.

When a packet arrives at the VRF and the **grt-lookup>enable-grt** command is configured, the following sequence occurs:

- The packet undergoes a route lookup in the VRF FIB to determine the next hop.
- If there is no matching route found in the VRF, a route lookup is then performed in the GRT FIB.
- If there is a match for the route in the GRT FIB and the packet is:
 - an IP packet with a local address, it is extracted to the CSM and processed as a management packet
 - a GRE packet with a local address, it is processed as a service packet
 - an IP or GRE packet with no local address, it is routed to the available next hop as found in the GRT FIB

In order for a packet to leave the VRF, the route that needs to be resolved—the destination prefix—must be leaked to the GRT. The destination prefix is configured in a route policy using the **config>router>policy-options>prefix-list** command; that policy is then leaked into the GRT by referencing it in the **config>service>vpn>grt-lookup>export-grt** command. Packets with the matching route found in the GRT FIB are routed via the IPSec tunnel configured within the VPN.

The 7705 SAR supports the following packet types for GRT lookup in the VRF FIB:

- self-generated GRE packets for Cpipe, Epipe, and VPLS traffic
- self-generated IP packets for management, BGP, and T-LDP traffic
- transiting GRE packets over access and network interfaces
- transiting IP packets over access or network interfaces

The 7705 SAR supports the following packets types arriving in the VRF via the IPSec tunnel:

- IP packets, including management packets, with a local address
- GRE packets with a local address
- transiting IP packets via the GRT lookup
- transiting GRE packets via the GRT lookup


The 7705 SAR terminates GRE packets at a system IP address or any local interface IP address. When GRE-encapsulated packets are transported over an IPSec VPN, the IPSec tunnel can terminate on a 7705 SAR or 7750 SR and the GRE packets are processed on the secure gateway.

When configuring **grt-lookup**, the **config>service>vpn>static-route-entry>grt** command must be configured for the local IP address that will be looked up in the GRT. The next hop is set to the desired local IP address and is used on ingress to force a second route lookup in the GRT. If that lookup is successful, and the packet is a GRE packet destined to a local interface, it is forwarded for PW processing. If the packet is destined to a system IP address and is not GRE packet, it is forwarded for management processing. A second static route is required in the IPSec VPN to point the far-end IP address to the IPSec tunnel. It is used on egress and is configured for the far-end IP address with a next hop set to the IPSec tunnel.

By enabling GRT lookup, signaling packets such as T-LDP can be routed securely between two system IP addresses by using the IPSec tunnel. However, IGP protocols such as OSPF or IS-IS, which use a multicast destination, cannot use the IPSec tunnel.

The following figure shows an example of route leaking configured to resolve the far-end system IP address 10.0.0.2 using an IPSec VPN. Although the example shows GRE to a system IP address on the 7705 SAR for illustrative purposes, GRE to any other local IP address on the 7705 SAR can be substituted.

Figure 130: Routing GRE-encapsulated packets over IPSec



26201

Based on the figure, the CLI example below shows the configuration of a routing policy that is used to leak the far-end system IP address into the GRT. The command **config>router>policy-options>prefix-list>prefix** creates a prefix entry with the system IP address of the far-end node (10.0.0.2/32) in the route policy prefix list. The policy option is configured using the prefix specified above (10.0.0.32) with the action **accept**. The policy **preference** should be set so that it is lower than the IGP advertised preference.

```
#-----
*A:7705:Dut-A>config>router>policy-options# info
#-----
    prefix-list "grt"
        prefix 10.0.0.2/32 exact
    exit
    policy-statement "grt"
        entry 1
            from
                prefix-list "grt"
            exit
            action accept
                local-preference 3
                preference 3
                metric-set 1
            exit
        exit
    exit
```

The far-end system IP address 10.0.0.2 is resolved using a static route configured with an IPSec tunnel next hop of "tunnel2". The GRT lookup at the ingress VRF is enabled using the **config>service>vprn>grt-lookup>enable-grt** command and a second static route is configured to enable lookup at ingress for the local system IP address 10.0.0.14/32, as shown in the CLI example below:

```
#-----
*A:7705:Dut-A>config>sevice>vprn# info
#-----
    static-route-entry 10.10.0.2
        ipsec-tunnel "tunnel2"
        no shutdown
    exit
    exit
    static-route-entry 10.0.0.14/32
        grt
        no shutdown
    exit
```

```

exit
grt-lookup
  enable-grt
  export-grt grt
  exit
exit
exit

```

The far-end system IP address with a next hop IPSec tunnel ("tunnel2") is leaked into the GRT using the command **config>service>vprn>grt-lookup>export-grt**, referencing the routing policy configured above ("grt").

A GRE-encapsulated SDP to the far-end system IP address is configured using the commands **config>service>sdp sdp-id gre create** and **config>service>sdp>far-end**. A Cpipe, Epipe, or VPLS is then configured using that SDP. For information about configuring a Cpipe or Epipe, see [Configuring a VLL service with CLI](#) and [Configuring VLL components](#) in this guide. For information about configuring a VPLS, see [Configuring a VPLS service with CLI](#) in this guide.

For information about GRT lookup for management traffic, see [In-band management using a VPRN](#).

8.1.1.13 GRE-encapsulated VLLs/VPLS over IPSec over MPLS

The 7705 SAR can route Cpipe, Epipe, or VPLS traffic over IPSec using either BGP 3107 label routes or RSVP-TE IGP shortcuts.

When GRE-encapsulated Cpipe, Epipe, or VPLS traffic is routed over IPSec, the GRE packets and T-LDP packets can be routed to the far-end system IP address using the IPSec tunnel. However, there must be special consideration for the MPLS tunnel, in particular for BGP 3107 label routes using IBGP, because MPLS signaling packets cannot use an IPSec tunnel.

8.1.1.13.1 VLLs/VPLS over IPSec over MPLS (using BGP 3107 label routes) solution 1: changing BGP signaling to loopback interface


When resolving an IPSec route to the secure gateway address with a BGP 3107 label route, BGP can be set up to use either IBGP or EBGP.

The IBGP 3107 label route is usually set up to the system IP address. However, the Layer 2 services to be protected by IPSec are also set up to the system IP address. In this case routing the IBGP 3107 label routes via the IPSec tunnel creates some complexities.

To avoid routing both Layer 2 services and the IBGP (used for 3107) over the IPSec tunnel, IBGP should be setup to a loopback interface instead of to the system IP address. Also, all MPLS LSPs that resolve these 3107 labels should be setup to the same loopback interface. For RSVP-TE this means the far end has to be the loopback interface and for LDP this means an originate-fec needs to be configured to the loopback interface.

The following figure shows an example of a BGP 3107 label route configured with a loopback interface.

Figure 131: VLL/VPLS over IPsec over MPLS using BGP 3107 label routes



Based on the figure, the CLI example below shows BGP configured to a neighbor loopback address and the BGP local address configured as a local loopback address. BGP for 3107 label route advertisement is enabled using the **advertise-label ipv4** keyword.

```
#-----
*A:7705:Dut-A>config>router>bgp# info
#-----
    local-as 10
    group "2"
    peer-as 10
    local-address 10.10.0.14
    neighbor 10.10.0.2
        family vpn-ipv4
        export "gw"
        advertise-label ipv4
    exit
exit
no shutdown
```

In addition, IGP must be configured to make the loopback IP addresses reachable for BGP.

The BGP 3107 label route can be resolved with either an LDP or an RSVP-TE tunnel. To use an LDP tunnel, an LDP FEC must be configured to advertise the local loopback IP address to the neighbor, as shown in the CLI example below:

```
#-----
*A:7705:Dut-A>config>router>ldp# info
#-----
    fec-originate 10.10.0.14/32 advertised-label 32 pop
```

To use an RSVP-TE tunnel, an LSP is created to the neighbor loopback IP address, as shown in the CLI example below:

```
#-----
*A:7705:Dut-A>config>router>mpls# info
#-----
    lsp "to-14-loop"
    to 10.10.0.2
    cspf
    no shutdown
exit
```


With EBGp, BGP communicates between the local and neighbor IP interface so IGP is not required to resolve the BGP 3107 label routes. In the example shown in the figure, the configuration would use Layer 3 interfaces instead of loopback addresses, as shown in the CLI example:

```
#-----
*A:7705:Dut-A>config>router>bgp# info
#-----
    local-as 10
    group "2"
    peer-as 10
    neighbor 10.14.1.2
        family vpn-ipv4
        export "gw"
        advertise-label ipv4
    exit
exit
no shutdown
```

To use an LDP tunnel, an LDP FEC is configured to advertise the local interface IP address:

```
#-----
*A:7705:Dut-A>config>router>ldp# info
#-----
    fec-originate 10.14.2.2/32 advertised-label 32 pop
```


To use an RSVP-TE tunnel, an LSP is created to the neighbor interface IP address:

```
#-----
*A:7705:Dut-A>config>router>mpls# info
#-----
    lsp "to-14-loop"
        to 10.14.1.2
        cspf
        no shutdown
    exit
```

8.1.1.13.2 VLLs/VPLS over IPSec over MPLS (using BGP 3107 label routes) solution 2: GRE to local interface on 7705 SAR

When transporting a GRE-encapsulated VLL or VPLS over IPSec over MPLS, the 7705 SAR can terminate the GRE tunnel to a loopback address. In this scenario, the VLL or VPLS is created between the loopback interfaces and the BGP 3107 label route uses the system IP address to resolve the IPSec gateway. This relies on GRT lookup and leaking, but the far-end loopback IP address is used instead of the system IP address in the routing policy. The following figure shows an example of VLL/VPLS over IPSec over MPLS using a loopback address, followed by CLI configuration examples.

Figure 132: VLL/VPLS over IPsec over MPLS using a loopback address



Based on the figure, a local loopback interface is configured on the 7705 SAR:

```
#-----
*A:7705:Dut-A>config>router# info
#-----
    interface "loop1"
        address 10.10.0.14
        loopback
        no shutdown
    exit
```

T-LDP signaling is configured to the far-end loopback IP address:

```
#-----
*A:7705:Dut-A>config>router>ldp# info
#-----
    targeted-session
        peer 10.10.0.2
        local-lsr-id "loop1"
        no shutdown
    exit
exit
```

A routing policy is configured on the 7705 SAR using the loopback address of the secure gateway:

```
#-----
*A:7705:Dut-A>config>router>policy-options# info
#-----
    prefix-list "loop"
        prefix 10.10.0.2/32 exact
    exit
    policy-statement "loop"
        entry 1
            from
                prefix-list "loop"
            exit
            action accept
                local-preference 3
                preference 3
                metric-set 1
            exit
        exit
    exit
```

The routing policy is then used to enable GRT lookup and leaking in the VPRN:

```
#-----
*A:7705:Dut-A>config>sevice>vprn# info
#-----
static-route-entry 10.10.0.2
    ipsec-tunnel "tunnel2"
    no shutdown
exit
static-route-entry 10.0.0.14/32
    grt
    no shutdown
exit
grt-lookup
    enable-grt
    export-grt loop
    exit
exit
```

The GRE SDP is then created to the far-end loopback IP address:

```
#-----
*A:7705:Dut-A>config>sevice# info
#-----
sdp 200 gre create
    far-end 10.10.0.2
    keep-alive
    shutdown
exit
no shutdown
exit
```

8.1.1.13.3 VLLs/VPLS over IPSec over MPLS (using IGP shortcuts)

BGP can route VLL or VPLS traffic over an IPSec VPN using an IGP shortcut to resolve the secure gateway address, as described in [Configuring IPSec over MPLS](#). Although the VLL or VPLS traffic destined for the far-end system IP address are routed using an IPSec tunnel, the IGP packets themselves are destined for a multicast address and are not resolved over the IPSec tunnel.

8.1.2 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 because of 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 considered 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.

8.1.2.1 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

8.1.2.2 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.

8.1.2.3 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**.

8.1.2.4 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 allow certificate signing

8.1.2.5 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:

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**

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.

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`

4. Because the key is imported, remove the key file generated in Step 1 for security reasons.
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.

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](#).

8.1.2.6 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**.

This behavior can be changed using the **revocation-check crl-optional** command under the **ca-profile** context, for the CA certificate only. As an example, if the IPSec application needs to use the CRL of a specific **ca-profile** to check the revocation status of an end entity certificate and the CRL is nonexistent for any reason, this is treated as the system being unable to get an answer from the CRL, and the system falls back to either the **secondary** method or the **default-result** configured under the **status-verify** context.

If the system needs to check the revocation status of a CA certificate in a certificate chain and the CRL is nonexistent for any reason, the system skips checking the revocation status of the CA certificate. For example, certificate CA1 is issued by certificate CA2. If CA2's **revocation-check** is **crl-optional** and CA2's CRL is nonexistent, the system does not check the revocation status of certificate CA1 and considers it to be good.

For details about OCSP, see [OCSP](#).

8.1.2.7 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.

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.

8.1.3 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.

8.1.4 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.

8.1.5 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, 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**.

8.1.6 Certificate Management Protocol version 2

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 and certificate update.

The 7705 SAR supports the following CMPv2 operations:

- initial registration – the process by which the 7705 SAR enrolls a certificate with a specific CA for the first time, where:
 - a public/private key pair must be preprovisioned 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. That means that no new requests are allowed if a pending request is present.

8.1.7 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 about the revocation status of a certificate.

The 7705 SAR supports both the CRL and OCSP as the certificate revocation status checking method. For a specific 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, 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.

8.1.8 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)

8.1.8.1 Metrocell deployment

As shown in the following figure, 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 133: Typical metrocell deployment




8.1.8.2 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, additional IPSec tunnels can be established. See the following figure.

Figure 134: Typical small business deployment



8.1.9 NAT-traversal for IKEv1/v2 and IPSec

The 7705 SAR supports network address translation traversal (NAT-T) for IKEv1 and IKEv2. NAT-T is functionality belonging to IPSec and IKEv1/v2. It is not functionality belonging to the NAT device.

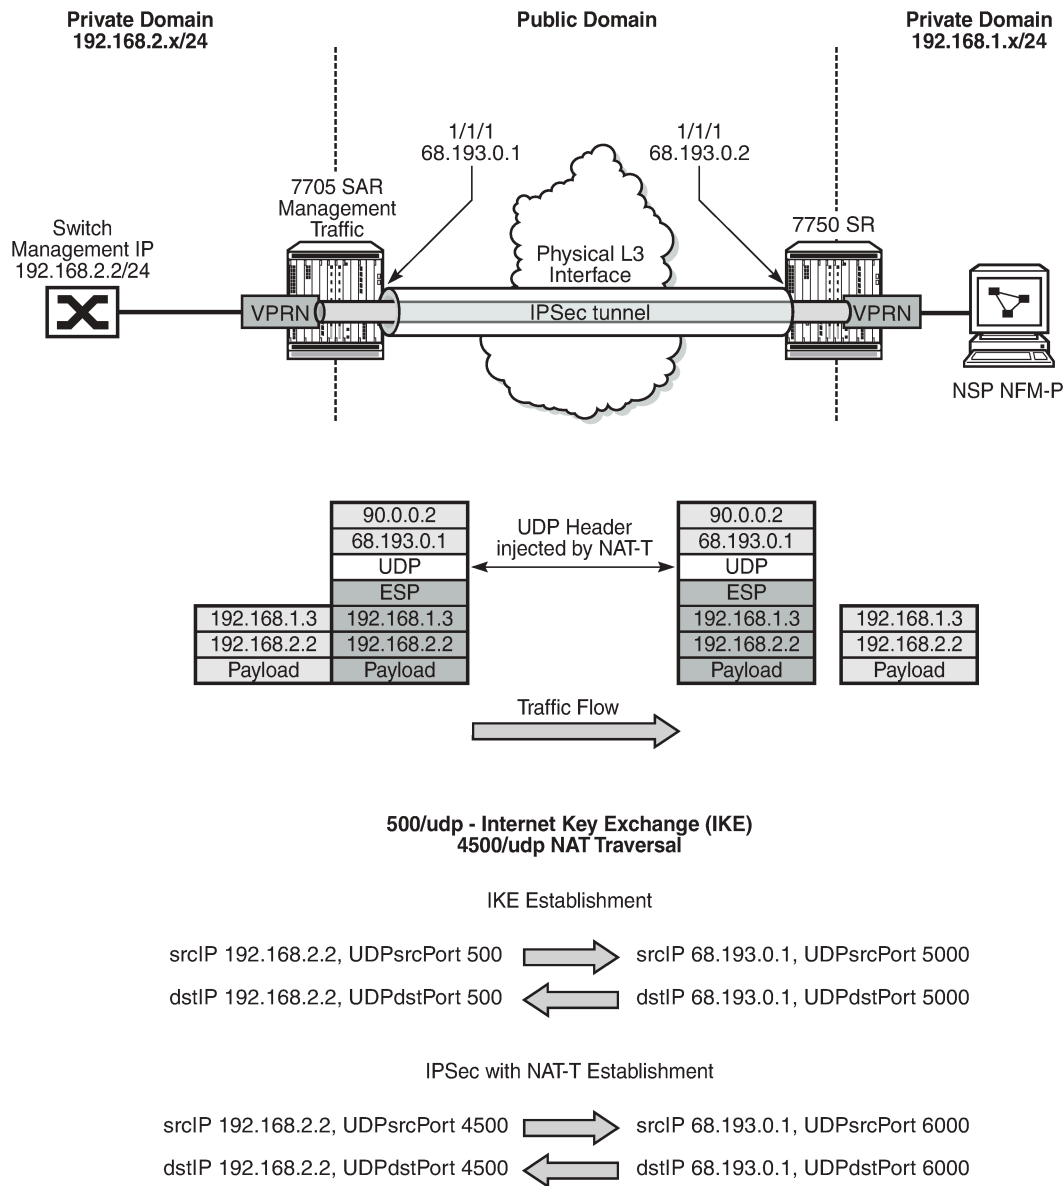
In a private network where the entire network is hidden behind a single public IP address, NAT-T for IPSec is used to support the fan-out of multiple IPSec tunnels in the private network.

IPSec is an IP protocol and therefore does not use ports. The following figure 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; therefore, 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 IKEv1/v2 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 the figure, 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 135: UDP header injected by a NAT-T-enabled IPSec tunnel



24016

8.1.10 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)

8.1.11 QoS for IPSec

This section contains information about the following topics:

- [Network and access ingress QoS \(decryption QoS\)](#)
- [Network and access egress QoS \(encryption QoS\)](#)

8.1.11.1 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 (see the 7705 SAR 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.

8.1.11.1.1 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 instead of on the DSCP marking of the IPSec outer IP header. For more information about the **ler-use-dscp** command, see the 7705 SAR Quality of Service Guide, "QoS policy network commands".

8.1.11.2 Network and access egress QoS (encryption QoS)

Customer packets arriving on access ingress in the VPRN are classified based on the SAP ingress policy (see the 7705 SAR 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.

8.1.12 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.

8.1.12.1 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>vpn>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.

8.1.12.2 Reassembly

The 7705 SAR does not support reassembly of fragmented IPSec packets.

8.1.13 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 the features supported include r-VPLS, VRRP and VRRPv3, ECMP, and LAG. See [VPRN services](#) for information on these features.

8.1.14 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.

8.1.15 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-Ax, 7705 SAR-H, 7705 SAR-Hc, 7705 SAR-Wx, and 7705 SAR-X).

To make use of 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.

8.1.16 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.

8.1.17 PBR and MFC

Both policy-based routing (PBR) and multi-field classification (MFC) are part of the ingress ACL configuration on the 7705 SAR. Both PBR and MFC are supported by IPSec on the 7705 SAR, as described in the following sections:

- [PBR](#)
- [MFC](#)

8.1.17.1 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. 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, the VPRN 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 about PBR, see the "Policy-based routing" section in the 7705 SAR 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 may 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 does not match the SeGW ESP configuration and the packet is dropped. Thus, it is the responsibility of the network administrator to ensure that the PBR configuration is correct and meets the network needs.

8.1.17.2 MFC

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 about MFC, see the "Multi-field classification" section in the 7705 SAR Router Configuration Guide.

8.1.18 OSPFv3 packet authentication with IPv6 IPSec

The 7705 SAR supports the use of IPv6 IPSec to authenticate OSPFv3 packets. The following features are supported:

- two types of encryption and authentication protocols: authentication header (AH) and IP encapsulating security payload (ESP)
- IPSec transport mode to authenticate the IP payload
- manually keyed IPSec security associations (SA)
- the MD5 and SHA1 authentication algorithms

To be authenticated, OSPFv3 peers must be configured with matching inbound and outbound SAs using the same parameters (for example, SPIs and encryption keys). One SA can be used for both inbound and outbound directions.

Authentication of OSPFv3 packets is supported on VPRN interfaces, network interfaces, and virtual links.

The 7705 SAR supports the rekeying procedure defined in RFC 4552, *Authentication/Confidentiality for OSPFv3*:

- For every router on the link, create an additional inbound SA for the interface being rekeyed using a new SPI and the new key.
- The SA replacement operation is atomic, meaning that no OSPFv3 packets are sent on the link until the replacement operation is complete. This ensures that no packets are sent without authentication or encryption.
- For every router on the link, remove the original inbound SA.

The key rollover procedure automatically starts when the operator changes the configuration of the inbound static security association or bidirectional static security association under an interface or virtual link. Within the KeyRolloverInterval time period, OSPFv3 accepts packets with both the previous inbound static SA and the new inbound static SA; the previous outbound static SA continues to be used. When the timer expires, OSPFv3 only accepts packets with the new inbound static SA. For outgoing OSPFv3 packets, the new outbound static SA is used instead.

8.1.19 Network security with IPv6 IPSec

A 7705 SAR system that supports encryption allows the use of IPv6 IPSec to provide network security over an IPv6 IPSec tunnel as defined in RFC 4301, *Security Architecture for the Internet Protocol*. An IPv6 IPSec tunnel is used to encrypt data from the access network to an endpoint.

Network security with IPv6 IPSec supports the following capabilities described in this guide:

- all IPv6 and 6VPE access capabilities
- static LAN-to-LAN
- PSK
- PKI
- VLL over GRE over IPSec
- static security association (SA) keying
- IKEv2 dynamic keying (IKEv1 is not supported on IPv6 IPSec)
- IKEv2 fragmentation as defined in RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*
- reassembly of IKE packets
- node management of IPv6 and IPv4 traffic over an IPv6 IPSec tunnel using GRT leaking

- dual-stack access interfaces
- BGPv6 signaling and static route configuration to direct traffic to the IPv6 IPSec tunnel
- BGPv6 support for IPv4 and IPv6 routes
- IPv4 packet fragmentation
- application of NAT and firewall security functions to IPv4 packets arriving on the private VPRN before the packets are routed to the IPv6 IPSec tunnel

Network security with IPv6 IPSec is also supported on a public VPRN and IES. On a public VPRN, the IPSec packets are treated as any other IP packets when 6VPE is configured on the VPRN. The 6VPE functionality is supported for IPSec packets over MPLS (MP-BGP) transport. The transport tunnels can be MPLS (SR, LDP/RSVP-TE) or GRE IPv4.

8.1.20 IPSec over r-VPLS on a public-side service

A 7705 SAR system that supports encryption allows the forwarding of IPSec IPv4 and IPv6 packets over an r-VPLS next-hop interface on a VPRN service. The r-VPLS next-hop interface is resolved using whatever IGP the customer uses (such as RIP, OSPF, IS-IS, static routes, or BGP). The r-VPLS interface bound to a VPLS service can then transport the IPSec tunneled packets over Layer 2 VPN to a SAP or a spoke or mesh SDP.

8.1.21 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 Shelf V2 or 7705 SAR-18), statistics do not synchronize over to the inactive CSM and at the time of a CSM switchover, all statistics are lost. Also, in the case of statistics rollover in the datapath, the newly retrieved statistics start from 0 (zero) again.

8.1.22 Security support

IPSec on the 7705 SAR requires a public-side service (IES or VPRN) and a private-side service (VPRN). All IPSec traffic on the public service is encrypted. By the time the traffic is routed to the private service, it has been decrypted. NAT can be applied to traffic traversing the IPSec public interface.

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, see the 7705 SAR Router Configuration Guide, "Configuring security parameters".

8.2 Public key infrastructure

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)

8.2.1 CA role in PKI

The role of a CA in PKI includes the following:

- 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.

8.2.2 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.

8.2.2.1 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

8.2.2.1.1 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.
 - a. The CA runs the hash over the X.509v3 certificate.
 - b. The result of the hash is encrypted via the CA public key (digital signature).
 - c. The digital signature is appended to the certificate and, consequently, the signed vendor certificate.

8.2.2.1.2 Vendor certificate authentication by a peer

While the IKE is being initiated, phase 1 of the IKE can be authenticated via PKI. Therefore, 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.

8.2.2.1.3 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.

8.2.2.1.4 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.

8.2.3 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. An example of the 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-R00TCA-CERT
09/22/2015  11:41a             458 CMS1-R00TCA-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 File(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.

8.2.4 CMPv2 certificate management

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**) – the process that the 7705 SAR uses to enroll a certificate with a particular CA for the first time. A public/private key pair must be preprovisioned before enrollment by means of local generation or another method.
- certificate update (**cr**) – the process whereby an initialized 7705 SAR obtains additional certificates
- key pair update (**kur**) – the process where the 7705 SAR updates an existing certificate for 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 those cases, the 7705 SAR supports polling requests and responses, as described in Section 5.3.22, *Polling Request and Response*, in RFC 4210.

8.2.4.1 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 preprovisioned with the operator CA certificate.

The 7705 SAR public/private key pair is always preprovisioned before enrollment by means of local generation or another method.

The 7705 SAR uses the CMPv2 initial registration process to enroll its preprovisioned 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

8.2.4.2 Key update

When a key pair is about 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 for 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, communication between the EE and the CA is 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.

8.2.4.3 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.

8.2.5 OCSP

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 Beans Certificate Authority (EJBICA) server contains (by default) an internal OCSP responder and 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.

8.2.6 Certificate or CRL expiration warning

The system can optionally generate a warning message before a certificate or a CRL expires. The amount of time before expiration is configurable with the **certificate-expiration-warning** and **crl-expiration-warning** commands. The warning messages can also be repeated at configured intervals.

If a configured EE certificate expires, the system does not bring down an established IPSec-tunnel; however, future certificate authentication fails.

If a CA certificate expires, the system brings the CA profile operationally down. This does not affect established tunnels; however, future certificate authentication that uses the CA profile fails.

8.2.7 Automatic CRL update

With the automatic CRL update feature, the 7705 SAR can be scheduled to automatically connect to a list of configured HTTP URLs to download a new CRL file. If a CRL file is successfully downloaded and qualified, it replaces the existing CRL file. A CRL file is considered qualified if it is a valid CRL signed by the CA and is more recent than the existing CRL. To determine if a downloaded CRL is more recent than the existing CRL, the system first compares the This Update field of both CRL files, which indicates the issue date of the CRL. If the dates are the same, the system compares the CRL number extension, if present; a higher number indicates a more recent CRL.

This feature supports two types of CRL update schedules:

- **periodic** – the system initiates a CRL update periodically, at the intervals specified by the **periodic-update-interval** command. For example, if the periodic update interval is 24 hours, the system checks the configured URLs for a new CRL file to download every 24 hours. The minimum periodic update interval is 1 hour.
- **next-update-based** – the system initiates a CRL update at the date and time specified in the Next Update field of the existing CRL file, minus the time configured with the **pre-update-time** command. For example, if the existing CRL Next Update is 2022-06-30 06:00 and the pre-update time is 1 hour, the system begins the CRL update process at 2022-06-30, 05:00.

Up to eight URL entries can be configured under each CA profile. The configured URLs must point to a DER-encoded CRL file. When a CRL update is initiated, the system accesses each URL in order, and the first successfully downloaded and qualified CRL is used to update the existing CRL. If the download fails or the downloaded CRL is not qualified, the system moves to the next URL in the list. If no CRL can be downloaded or qualified, the system attempts to contact each URL again at the next scheduled update time (when the schedule type is **periodic**) or after the time configured with the **retry-interval** command (when the schedule type is **next-update-based**).

HTTP transport can be over IPv4 or IPv6. Automatic CRL update supports base, management, or VPRN routing instances. If VPRN is used, the HTTP server port can only be 80 or 8080.

A CRL update is initiated immediately if **auto-crl-update** is enabled and the system detects that the configured CRL file does not exist, or is invalid or expired, or if the schedule type is configured as **next-update-based** and the scheduled update time has already passed.

A CRL update can be initiated manually with the **admin>certificate>crl-update** command, but automatic CRL update must first be shut down.

8.3 IPSec best practices recommendations

To prevent 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 two or three 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 three times larger than the CHILD_SA lifetime.

The IKE protocol is the control plane of IPSec; therefore, 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.

8.4 Configuration notes

This section describes operational conditions and IPSec configuration guidelines and restrictions:

- 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 does 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 does 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 require a configuration change. If the subnet has been changed, the tunnels are 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 or VPRN 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.

8.5 Configuring IPSec with CLI

This section provides information to configure IPSec using the CLI.

Topics in this section include:

- [Basic configuration overview](#)

- [Common configuration tasks](#)
- [Configuring IPSec over MPLS](#)
- [Service management tasks](#)

8.6 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 a 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 an IES or VPRN service.
 - Create an interface and public tunnel SAP.
- Configure a VPRN static route for the IPSec tunnel.

8.7 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 private service](#)
- [Configuring X.509v3 certificate parameters](#)
- [Configuring CMPv2](#)
- [Configuring OCSP](#)

8.7.1 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#

```

8.7.2 Configuring router interfaces for IPSec

An IPSec tunnel requires the following three interfaces:

- public tunnel interface (under IES or VPRN)
- private tunnel interface (under VPRN)
- physical untrusted network/Internet-facing interface: IES, VPRN, 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>router# info
-----
...
  router
    interface "internet"
      address 10.10.7.118/11
      port 1/1/1
    exit
    interface "system"
      address 10.20.1.118/12
    exit
    autonomous-system 123
  exit
...
-----
*A:ALU-49>config>router#

*A:7705:Dut-A>config>service>ies# info
-----
      description "ies interface toward internet"
      interface "public" create
        address 10.1.1.1/1
        sap 1/1/8 create
          description "sap-100-10.1.1.1"
        exit
      exit
      no shutdown
-----

```

8.7.3 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
  ike-version 2
  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
  no ikev2-fragment
  dpd interval 25
exit
ipsec-transform 2 create
  esp-auth-algorithm md5
  esp-encryption-algorithm 3des
exit
exit
#-----

```

8.7.4 Configuring IPSec and IPSec tunnels in services

IPSec is configured under IES and VPRN services.

For the private-side IPSec tunnel interface and SAP, under the VPRN service context, configure IPSec security policies and create tunnel interfaces, private tunnel SAPs, IPSec tunnels, and IPSec tunnel parameters. The **tunnel** keyword must be used when creating an interface for a private tunnel SAP.

For a public-side IPSec tunnel interface and SAP, under the IES or VPRN service context, create an interface and public tunnel SAP. The **tunnel** keyword is not used when creating an interface for a public tunnel SAP.

Private-side and public-side tunnels function in pairs, where a pair is defined by the service ID and the interface subnet.

The local gateway address and delivery service configured using the VPRN **ipsec-tunnel>local-gateway-address** command correspond to the IES or VPRN interface address and service ID where the public-side 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 a private-side VPRN service and a public-side 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 198.51.100.0/24
        remote-ip 198.51.100.0/24
      exit
    exit
  security-policy 15 create
    entry 15 create
      no local-ip

```

```

        no remote-ip
    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 10.10.10.11
        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-entry 192.100.200.10/32
    ipsec-tunnel "ipsec_tunnel_tag1"
    no shutdown
exit
exit
-----
*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 10.10.10.1/8
        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#
```

8.7.5 Configuring IPSec IPv6 parameters for a VPRN private service

Use the following CLI syntax to configure IPSec IPv6 parameters for a VPRN private 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$
    2001:db8:a::123/64
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$
    2001:db8:a::222/64
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 2001:db8:a::123/64
                remote-v6-ip 2001:db8:a::222/64
            exit
        exit
    exit
```

8.7.6 Configuring X.509v3 certificate parameters

Perform the following steps to configure certificate enrollment:

1. Generate a key:

```
admin certificate gen-keypair cf3:/key_plain_rsa2048 size 2048 type rsa
```

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
```

3. Send the certificate request to CA-1 to sign and get the signed certificate.

4. Import the key:

```
admin certificate import type key input cf3:/key_plain_rsa2048 output key1_rsa2048 format der
```

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:

1. Import the CA certificate:

```
admin certificate import type cert input cf3:/CA_1_cert.pem output ca_cert format pem
```

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>vpn>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
```

8.7.7 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, because 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-for-ir**) to support interoperation with CMPv2 servers.

8.7.8 Configuring OCSP

OCSP server information is configured under the corresponding **ca-profile**:

CLI syntax:

```
config>system>security>pki>ca-profile>
  oosp
    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>
  oosp
    responder-url "http://ocsp.example.com/request"
    service 100
```

For a specified IPSec tunnel, the user can configure a primary method, a secondary method, and a default result.

CLI syntax:

```
config>service>vpn>if>sap>ipsec-tun>
  cert
    status-verify
    primary {ocsp | crl}
    secondary {ocsp | crl}
    default-result {revoked | good}
```

Example:

```
config>service>vpn>if>sap>ipsec-tun>
  cert
    status-verify
    primary ocsp
    secondary crl
```

8.8 Configuring IPSec over MPLS

On the 7705 SAR, IPSec routes to the secure gateway address can be resolved by using either a BGP 3107 label route or an IGP shortcut. When BGP learns IPv4 addressed as BGP 3107 label routes, BGP resolves the next hops for these routes with an LDP or RSVP-TE tunnel. These BGP routes create BGP tunnels that can be used to resolve an IPSec secure gateway address. When an IGP shortcut is enabled on the 7705 SAR by using the **config>router>ospf>rsvp-shortcut** command, OSPF installs an OSPF route in the RIB, with an RSVP-TE LSP as the next hop. If this OSPF route is determined as the overall best route, then the next hop is an RSVP-TE tunnel.

The IPSec implementation on the 7705 SAR is VPN-based. To configure IPSec, a private VPRN and a public IES or VPRN must both be configured; the encryption and decryption functions occur between these two services.

This section shows a configuration example of an IPSec route resolved by a BGP 3107 label route and a configuration example of an IPSec route resolved by an IGP shortcut.

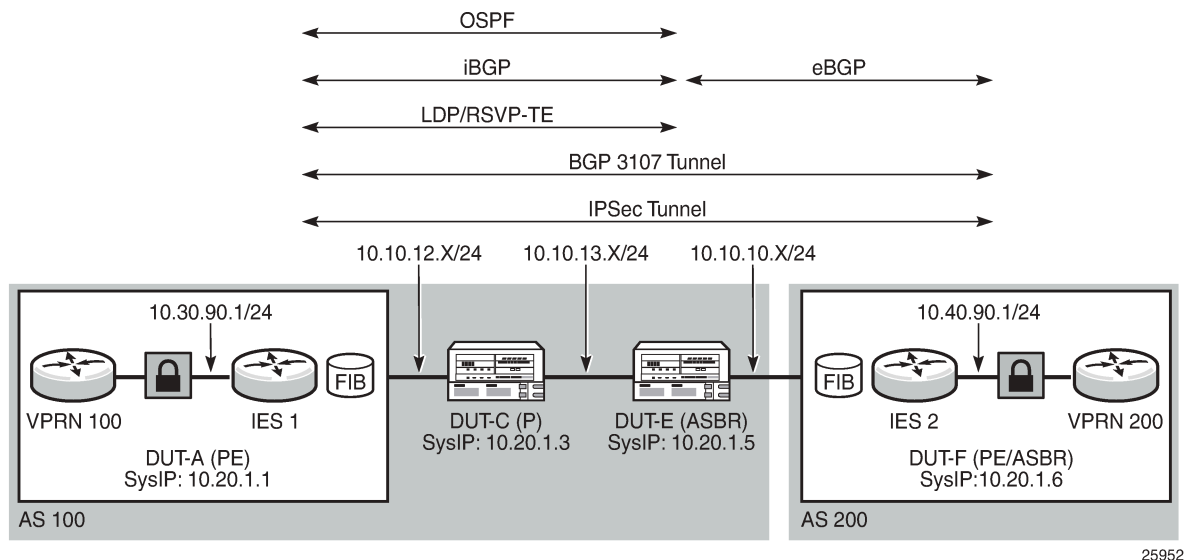
8.8.1 IPSec over BGP 3107 label routes

To route IPSec traffic using BGP 3107 label routes, the following components must be configured:

- a static LAN-to-LAN tunnel for IPSec traffic
- a policy option to advertise the IPSec gateway using BGP
- BGP with a BGP 3107 label route configured
- an LDP or RSVP-TE tunnel to resolve the BGP 3107 label route

The following figure shows a scenario where IPSec traffic is routed over a BGP 3107 label route. In this example, both the BGP 3107 tunnel and the IPSec tunnel are set up between Dut-A and Dut-F. The nature of BGP 3107 requires the LDP or RSVP-TE tunnel to be set up inside the autonomous system between Dut-A and Dut-E.

Figure 136: IPSec over BGP 3107 label route



8.8.1.1 Static LAN-to-LAN tunnel configuration

Setting up a static LAN-to-LAN tunnel for IPSec traffic involves configuring a number of elements, including:

- VPRN private-side service parameters, including the following:
 - BGP parameters
 - route distinguisher parameter
 - **auto-bind-tunnel** parameter or VPRN spoke SDP
 - VRF route-target associations or VRF import/export policies
 - OSPF parameters
 - a VPRN interface and its SAP parameters

- spoke-SDP parameters on the VPRN interface
- IES or VPRN public-side service parameters
- IPSec parameters

The CLI output below is an example of a static LAN-to-LAN tunnel configuration.

```
*A:7705:Dut-A>config>service>vprn# info
-----
description "Default Description For VPRN ID 90"
snmp-community "Ku/I.yvsMoQ" hash2 version both
ipsec
    security-policy 1 create
        entry 1 create
            local-ip any
            remote-ip any
        exit
    exit
exit
router-id 10.20.1.1
autonomous-system 900
route-distinguisher 10.20.1.1:90
auto-bind-tunnel
    resolution-filter
        ldp
    exit
    resolution filter
exit
vrf-target target:65000:90
interface "ies-90-192.168.90.1" create
    address 192.168.90.1/24
    sap 1/2/1:900 create
        description "sap-90-192.168.90.1"
    exit
exit
interface "ies-90-192.168.90.2" create
    address 192.168.90.2/24
    loopback
exit
interface "vprn-90-sap-tunnelPrivate-1" tunnel create
    sap tunnel-1.private:1 create
        description "sap-90-IPSEC"
    ipsec-tunnel "tunnelPrivateSide1" create
        security-policy 1
        local-gateway-address 10.30.90.1 peer 10.40.90.1 delivery-
service 9090
    dynamic-keying
        ike-policy 1
        pre-shared-key "SmS3kjoVVF8ovXf0fxudQJ/
tw3MPVYZp1x1v2z2KkYJ5xY0hdURJyU" hash2
        transform 1
    exit
    no shutdown
    exit
    exit
    static-route-entry 10.1.1.1/8
        ipsec-tunnel "tunnelPrivateSide1"
        no shutdown
    exit
    exit
    bgp
        min-route-advertisement 1
```

```

import "BgpVpn_to_Bgp"
export "BgpVpn_to_Bgp"
router-id 10.20.1.1
group "ce-peers"
    neighbor 10.1.1.4
        local-address 10.1.1.3
        peer-as 90000
    exit
    neighbor 10.1.1.5
        local-address 10.1.1.6
        med-out 100
        peer-as 9001
    exit
exit
no shutdown
exit
service-name "XYZ Vprn 90"
no shutdown
-----
*A:7705:Dut-A>config>service>vprn#
*A:7705:Dut-A>config>service>vprn# exit all
*A:7705:Dut-A# configure service ies 9090
*A:7705:Dut-A>config>service>ies# info
-----
description "Default Ies description for service id 9090"
interface "tunnelPublicSide1" create
    address 10.30.90.3/8
    sap tunnel-1.public:1 create
        description "sap-9090-10.30.90.3"
    exit
exit
service-name "XYZ Ies 9090"
no shutdown
-----
*A:7705:Dut-A>config>service>ies#

```

8.8.1.2 Policy option configuration

The CLI output below is an example of a policy option configuration.

```

#-----
*A:7705:Dut-A>config>router>policy-options# info
#-----
prefix-list "pe_sys_pref"
    prefix 10.30.90.0/8 longer
exit
policy-statement "pe_sys_to_bgp"
    entry 10
        from
            prefix-list "pe_sys_pref"
        exit
        to
            protocol bgp
        exit
        action accept
        exit
    exit
exit
commit
exit

```

8.8.1.3 BGP configuration with BGP 3107 label route advertisement

The CLI output below is an example of BGP enabled with label route advertisement.

```
#-----
*A:7705:Dut-A>config>router>bgp# info
#-----
    bgp
      connect-retry 5
      keepalive 5
      hold-time 15
      min-route-advertisement 2
      transport-tunnel mpls
      group "to_asbr_Dut-E"
        description "Group to ASBR - vpn label v4"
        peer-as 100
        neighbor 10.20.1.5
          family ipv4 vpn-ipv4 vpn-ipv6
          export "pe_sys_to_bgp"
          peer-as 100
          advertise-label ipv4
        exit
      exit
    no shutdown
  exit
```

8.8.1.4 LDP or RSVP-TE tunnel configuration

The CLI output below is an example of an LDP tunnel that is configured to resolve the next hop for the BGP 3107 label route. An RSVP-TE tunnel can also be configured to resolve the next hop.

```
*A:7705:Dut-A>config>router# info
#-----
echo "IP Configuration"
#-----
    interface "system"
      address 10.20.1.1/8
      no shutdown
    exit
    interface "to-P-Dut-C"
      address 10.10.12.1/8
      port 1/2/7:12
      no shutdown
    exit
    interface "to-P-Dut-D"
      address 10.10.3.1/8
      port 1/2/3:1
      no shutdown
    exit
    autonomous-system 100
#-----
echo "OSPFv2 Configuration"
#-----
    ospf
      traffic-engineering
      timers
        spf-wait 1000 1000 1000
    exit
```

```

        area 0.0.0.0
        interface "system"
            hello-interval 5
            dead-interval 15
            metric 100
            no shutdown
        exit
        interface "to-P-Dut-D"
            interface-type broadcast
            hello-interval 1
            dead-interval 4
            mtu 1518
            metric 100
            no shutdown
        exit
        interface "to-P-Dut-C"
            interface-type broadcast
            hello-interval 1
            dead-interval 4
            mtu 1518
            metric 100
            no shutdown
        exit
    exit
exit
#-----
echo "MPLS Configuration"
#-----
    mpls
        interface "system"
            no shutdown
        exit
        interface "to-P-Dut-D"
            no shutdown
        exit
        interface "to-P-Dut-C"
            no shutdown
        exit
    exit
#-----
echo "MPLS LSP Configuration"
#-----
    mpls
        path "to-Dut-E"
            hop 1 10.20.1.3 strict
            no shutdown
        exit
        lsp "lsp-to-Dut-E"
            to 10.20.1.5
            cspf
            fast-reroute facility
        exit
        retry-timer 20
        primary "to-Dut-E"
        exit
        no shutdown
    exit
    no shutdown
exit
#-----
echo "LDP Configuration"
#-----
    ldp
        interface-parameters

```

```

        interface "to-P-Dut-D"
        exit
        interface "to-P-Dut-C"
        exit
    exit
    targeted-session
    exit
    no shutdown
exit

```

8.8.2 IPSec over IGP shortcut

To route IPSec traffic over an IGP shortcut, the following must be configured:

- a static LAN-to-LAN tunnel
- an IGP shortcut (by creating an RSVP-TE tunnel in the OSPF context)
- an RSVP-TE LSP to the system IP address or loopback address, with CSPF enabled

8.8.2.1 Static LAN-to-LAN tunnel configuration

The CLI output below is an example of a static LAN-to-LAN tunnel configuration.

```

#-----
echo "IPsec Configuration"
#-----
ipsec
    ike-policy 1 create
        description "ikePolicy_1"
        own-auth-method psk
        dh-group 1
        auth-algorithm md5
        dpd interval 10 max-retries 2
    exit
    ipsec-transform 1 create
        esp-auth-algorithm sha512
        esp-encryption-algorithm aes256
    exit
exit
#-----
echo "Service Configuration"
#-----
service
    customer 1 create
        description "Default customer"
    exit
    ies 101 customer 1 create
        interface "tunnelPublicSide_1" create
    exit
    exit
    vprn 1001 customer 1 create
        interface "tunnelPrivateSide_1" tunnel create
    exit
    interface "toIxia_1" create
    exit
    exit
    ies 101 customer 1 create
        description "Default Ies description for service id 101"
        interface "tunnelPublicSide_1" create

```



```

        address 10.1.254.1/8
        sap tunnel-1.public:1 create
        description "sap-10-10.1.254.1"
    exit
exit
service-name "XYZ Ies 101"
no shutdown
exit
vprn 1001 customer 1 create
description "Default Description For VPRN ID 1001"
ipsec
    security-policy 1 create
    entry 1 create
        local-ip 10.10.10.0/8
        remote-ip 10.1.1.0/8
    exit
exit
exit
route-distinguisher 1.1.1.1:1001
interface "tunnelPrivateSide_1" tunnel create
    sap tunnel-1.private:1 create
    description "sap-1001-IPSEC"
    ipsec-tunnel "tunnelPrivateSide_1.1" create
        security-policy 1
        local-gateway-address 10.1.1.1 peer 10.2.2.2 delivery-
service 101
    dynamic-keying
        ike-policy 1
        pre-shared-
key ".7ZAfd0optpg.FzYqTSVYbfFgzc.GZYw7W98X2uDhnHy/VmhhkWqkP." hash2
        auto-establish
        transform 1
    exit
    no shutdown
exit
exit
interface "toIxia_1" create
    address 10.254.254.1/8
    sap 1/2/1:101 create
    exit
exit
static-route-entry 10.1.1.0/8
    ipsec-tunnel "tunnelPrivateSide_1.1"
    no shutdown
    exit
exit
service-name "XYZ Vprn 1001"
no shutdown
exit
exit
#-----
echo "OSPFv2 Configuration"
#-----
    ospf
        area 0.0.0.0
        interface "tunnelPublicSide_1"
            hello-interval 5
            dead-interval 15
            no shutdown
        exit
    exit
exit
#-----

```

8.8.2.2 IGP shortcut configuration

The CLI output below is an example of an IGP shortcut configuration. An IGP shortcut is created using the **rsvp-shortcut** command in the **ospf** context.

```
#-----
echo "OSPFv2 Configuration"
#-----
    ospf
        traffic-engineering
        timers
            spf-wait 1000 1000 1000
        exit
        rsvp-shortcut
        area 0.0.0.0
            interface "system"
                hello-interval 5
                dead-interval 15
                no shutdown
            exit
            interface "network"
                hello-interval 5
                dead-interval 15
                metric 100
                no shutdown
            exit
        exit
    exit
exit
#-----
```

8.8.2.3 RSVP-TE LSP configuration

The CLI output below is an example of an RSVP-TE LSP with CSPF enabled.

```
#-----
echo "MPLS Configuration"
#-----
    mpls
        interface "system"
            no shutdown
        exit
        interface "network"
            no shutdown
        exit
    exit
#-----
echo "RSVP Configuration"
#-----
    rsvp
        interface "system"
            no shutdown
        exit
        interface "network"
            no shutdown
        exit
    exit
#-----
echo "MPLS LSP Configuration"
```

```
#-----
mpls
  path "Path1AToC"
    no shutdown
  exit
  lsp "Lsp1AToC"
    to 10.10.20.1
    cspf
    retry-timer 20
    metric 100
    primary "Path1AToC"
    exit
    no shutdown
  exit
  no shutdown
exit
exit
#-----
```

8.9 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-side IPSec tunnel SAP and interface](#)
- [Deleting a private-side IPSec tunnel SAP and interface](#)
- [Deleting an IPSec security policy](#)
- [Deleting an IPSec tunnel](#)

8.9.1 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
```

8.9.2 Deleting a public-side IPSec tunnel SAP and interface

A public-side IPSec tunnel interface and SAP are created under an IES or VPRN service. The output below uses the CLI syntax and an example from the IES context to show how to delete a public-side 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
```

8.9.3 Deleting a private-side IPSec tunnel SAP and interface

A private-side IPSec tunnel interface and SAP are created under a VPRN service. To delete a private-side 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
```

8.9.4 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
```

8.9.5 Deleting an IPSec tunnel

IPSec tunnels are created under the VPRN service. Although an IPSec tunnel is created on the private side of the tunnel in the CLI, the configuration itself is general and can apply to either the public or private side of the tunnel. 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
```

8.10 IPSec command reference

8.10.1 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](#)
 - [Automatic CRL update commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

8.10.1.1 IPSec configuration commands

8.10.1.1.1 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
```

8.10.1.1.2 IPSec commands

```
config
- ipsec
- ike-policy ike-policy-id [create]
- no ike-policy ike-policy-id
- auth-algorithm {md5 | sha1 | sha256 | sha384 | sha512 | auth-encryption}
- 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 | aes128-gcm8
| aes128-gcm16 | aes256-gcm8 | aes256-gcm16}
- no encryption-algorithm
- ike-mode {main | aggressive}
- no ike-mode
- ike-version {1 | 2}
- no ike-version
- ikev2-fragment mtu octets reassembly-timeout seconds
- no ikev2-fragment
- ipsec-lifetime ipsec-lifetime
- no ipsec-lifetime
- isakmp-lifetime isakmp-lifetime
- no isakmp-lifetime
- [no] match-peer-id-to-cert
- 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
- prf-algorithm {md5 | sha1 | sha256 | sha384 | sha512 | aes-xcbc | same-as-auth}
- no prf-algorithm
- ipsec-transform transform-id [create]
- no ipsec-transform transform-id
- esp-auth-algorithm {null | md5 | sha1 | sha256 | sha384 | sha512 | auth-
encryption}
- no esp-auth-algorithm
- esp-encryption-algorithm {null | des | 3des | aes128 | aes192 | aes256 | aes128-
gcm8 | aes128-gcm12 | aes128-gcm16 | aes192-gcm8 | aes192-gcm12 | aes192-gcm16 | aes256-gcm8 |
aes256-gcm12 | aes256-gcm16}
- no esp-encryption-algorithm
- static-sa sa-name [create]
- no static-sa sa-name
- authentication auth-algorithm ascii-key ascii-string
- authentication auth-algorithm hex-key hex-string [hash | hash2]
- no authentication
- direction ipsec-direction
- no direction
- protocol ipsec-protocol
- no protocol
- spi spi
- no spi

```

8.10.1.1.3 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

```

8.10.1.1.4 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 b ip-
address
- no bfd-enable
- [no] clear-df-bit
- [no] copy-df-bit
- description description-string
- no description
- [no] dynamic-keying
- [no] auto-establish
- cert
- cert-profile profile
- no cert-profile
- remote-id type {ipv4 | ipv6 | fqdn | email} value value
- no remote-id
- 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
- ike-policy ike-policy-id
- no ike-policy
- local-id type {ipv4 | fqdn | ipv6} value value
- no local-id
- pre-shared-key key [hash | hash2]
- 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

```

8.10.1.1.5 Service static route commands

```

config
- service
- vprn service-id
- [no] static-route-entry ip-prefix/prefix-length
- [no] ipsec-tunnel ipsec-tunnel-name
- [no] description description-string
- [no] metric metric
- [no] preference preference
- [no] shutdown
- [no] tag tag

```

See [VPRN service configuration commands](#) for the command descriptions.

8.10.1.2 PKI configuration commands

8.10.1.2.1 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]

```

8.10.1.2.2 PKI infrastructure commands

```

config
- 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
- revocation-check {crl | crl-optional}
- [no] shutdown
- certificate-display-format {ascii | utf8}
- certificate-expiration-warning hours [repeat repeat-hours]
- no certificate-expiration-warning
- crl-expiration-warning hours [repeat repeat-hours]
- no crl-expiration-warning
- maximum-cert-chain-depth level
- no maximum-cert-chain-depth

```

8.10.1.2.3 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

```

8.10.1.2.4 Automatic CRL update commands

```

admin
- certificate
  - crl-update ca ca-profile-name
config
- system
  - file-transmission-profile name [create]
  - no file-transmission-profile name
    - ipv4-source-address ip-address
    - no ipv4-source-address
    - ipv6-source-address ipv6-address
    - no ipv6-source-address
    - redirection level
    - no redirection
    - retry count
    - no retry
    - router router-instance
    - router service vprn-service-name
    - timeout seconds
- security
  - pki
    - ca-profile name [create]
    - no ca-profile name
      - auto-crl-update [create]
      - no auto-crl-update
        - crl-urls
          - url-entry entry-id [create]
          - no url-entry entry-id
            - file-transmission-profile profile-name
            - no file-transmission-profile
            - url url
            - no url
          - periodic-update-interval [days days] [hrs hours] [min minutes]
        - pre-update-time [days days] [hrs hours] [min minutes] [sec seconds]
        - retry-interval seconds
        - no retry-interval
        - schedule-type schedule-type
        - [no] shutdown

```

8.10.1.3 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-8 Shelf
V2 and 7705 SAR-18)
- mda aggregate-statistics      (for 7705 SAR-Ax, 7705 SAR-H, 7705 SAR-Hc, and 7705 SAR-
Wx)

```

See the section "Show, Monitor, Clear, and Debug Command Reference" in the 7705 SAR Interface Configuration Guide for information about the **show>mda** commands.

```

show
- router
  - interface ip-int-name statistics

```

See the section "IP Router Command Reference" in the 7705 SAR Router Configuration Guide for information about the **show>router >interface statistics** command.

8.10.1.4 Clear commands

```

clear
- mda {slot/mda | all}
- mda all statistics
- mda slot/mda statistics security [encryption]

```

8.10.1.5 Debug commands

```

debug
- [no] cmpv2
  - [no] ca-profile profile-name
- ipsec
  - [no] certificate filename
  - tunnel [ipsec-tunnel-name] [detail]
  - no tunnel [ipsec-tunnel-name]

```

8.10.2 Command descriptions

- [IPSec configuration commands](#)
- [PKI configuration commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

8.10.2.1 IPSec configuration commands

- [Generic commands](#)
- [ISA tunnel commands](#)
- [Internet key exchange \(IKE\) and transform commands](#)
- [Service configuration commands](#)
- [Service interface tunnel commands](#)

8.10.2.1.1 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>vpn>interface

config>service>vpn>if>sap

config>service>vpn>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 (such as #, \$, or spaces), 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>vpn>interface
config>service>vpn>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.

8.10.2.1.2 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.

8.10.2.1.3 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 | auth-encryption}

no auth-algorithm

Context

config>ipsec>ike-policy

Description

This command specifies which hashing algorithm to use for the IKE authentication function.

The **auth-encryption** option must be specified when the encryption algorithm configured for the IKE session is an AES-GCM algorithm.

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 HMAC-SHA256 algorithm for authentication

sha384

specifies the HMAC-SHA384 algorithm for authentication

sha512

specifies the HMAC-SHA512 algorithm for authentication

auth-encryption

specifies an AES-GCM encryption 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 because 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 is used to test connectivity to the tunnel peer. If the peer initiates the connectivity check before the interval timer, it is 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 | aes128-gcm8 | aes128-gcm16 | aes256-gcm8 | aes256-gcm16}

no encryption-algorithm

Context

config>ipsec>ike-policy

Description

This command specifies the encryption algorithm to use for the IKE session.

When AES-GCM is configured as the encryption algorithm for an IKE policy:

- **auth-algorithm** must be set to **auth-encryption**
- manual keying is not possible

The **no** form of the command returns the algorithm to its default value.

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 3DES 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.

aes128-gcm8

configures the AES-GCM algorithm with a 128-bit key size and an 8-byte Integrity Check Value (ICV) for encryption and authentication

aes128-gcm16

configures the AES-GCM algorithm with a 128-bit key size and a 16-byte ICV for encryption and authentication

aes256-gcm8

configures the AES-GCM algorithm with a 256-bit key size and an 8-byte ICV for encryption and authentication

aes256-gcm16

configures the AES-GCM algorithm with a 256-bit key size and a 16-byte ICV for encryption and authentication

ike-mode

Syntax

ike-mode {main | aggressive}

no ike-mode

Context

config>ipsec>ike-policy

Description

This command specifies the mode of operation for IKEv1 phase 1, either main mode or aggressive mode. The difference between the modes is the number of messages used to establish the session. IKEv1 phase

1 main mode uses three pairs of messages (for a total of six messages) between IPSec peers. IKEv1 phase 1 aggressive mode has only three message exchanges.

This command does not apply to IKEv2.

The **no** form of the command removes the mode of operation.

Default

main

Parameters

main

specifies that IKEv1 phase 1 operates in main mode

aggressive

specifies that IKEv1 phase 1 operates in aggressive mode

ike-version

Syntax

ike-version {1 | 2}

no ike-version

Context

config>ipsec>ike-policy

Description

This command configures the version of the IKE protocol that the IKE policy uses.

The **no** form of the command removes the configured version.

Default

2

Parameters

1

specifies that the IKE policy uses IKEv1

2

specifies that the IKE policy uses IKEv2

ikev2-fragment

Syntax

ikev2-fragment mtu *octets* reassembly-timeout *seconds*

no ikev2-fragment

Context

config>ipsec>ike-policy

Description

This command enables IKEv2 protocol-level fragmentation (per RFC 7383). The MTU specified is the maximum size of the IKEv2 packet.

IKEv2 fragmentation is enabled for a tunnel only if this command is configured and if the peer also announces its support by sending an IKEV2_FRAGMENTATION_SUPPORTED notification.

Default

no ikev2-fragment

Parameters

octets

the MTU for IKEv2 messages

Values 512 to 9000

seconds

the time allowed for fragment reassembly before the fragments are discarded

Values 1 to 5

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

match-peer-id-to-cert

Syntax

[no] match-peer-id-to-cert

Context

config>ipsec>ike-policy

Description

This command enables a peer ID check during certificate authentication.

The certificate is authenticated if the Subject Alternative Name field matches the IKE identifier of the peer certificate.

When this command is configured, the [remote-id](#) command must be disabled because the configurations are mutually exclusive.

Default

no match-peer-id-to-cert

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
<i>keep-alive-interval</i>	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 to self-authenticate. This command (**own-auth-method**) applies only to IKEv2.

The default self-authentication method used by the 7705 SAR 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

prf-algorithm

Syntax

prf-algorithm {md5 | sha1 | sha256 | sha384 | sha512 | aes-xcbc | same-as-auth}

no prf-algorithm

Context

config>ipsec>ike-policy

Description

This command specifies the authentication algorithm to use in an IKE policy for the pseudorandom function (PRF).

If an AES-GCM authenticated encryption algorithm is used for IKE encryption, the **same-as-auth** keyword cannot be used for the PRF algorithm.

The **no** form of the command returns the command to the default setting.

Default

same-as-auth

Parameters

md5

specifies the HMAC-MD5 algorithm for PRF

sha1

specifies the HMAC-SHA1 algorithm for PRF

sha256

specifies the HMAC-SHA256 algorithm for PRF

sha384

specifies the HMAC-SHA384 algorithm for PRF

sha512

specifies the HMAC-SHA512 algorithm for PRF

aes-xcbc

specifies the AES128-XCBC algorithm for PRF

same-as-auth

specifies to use the same algorithm that is being used for the IKE session

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** | **auth-encryption**}
no esp-auth-algorithm

Context

config>ipsec>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 **auth-encryption** option must be specified when the ESP encryption algorithm configured for IPSec transform is an AES-GCM algorithm.

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 HMAC-SHA256 algorithm for authentication

sha384

configures ESP to use the HMAC-SHA384 algorithm for authentication

sha512

configures ESP to use the HMAC-SHA512 algorithm for authentication

auth-encryption

configures ESP to use an AES-GCM algorithm for authentication

esp-encryption-algorithm

Syntax

```
esp-encryption-algorithm {null | des | 3des | aes128 | aes192 | aes256 | aes128-gcm8 | aes128-gcm12 |
aes128-gcm16 | aes192-gcm8 | aes192-gcm12 | aes192-gcm16 | aes256-gcm8 | aes256-gcm12 |
aes256-gcm16}
no esp-encryption-algorithm
```

Context

```
config>ipsec>ipsec-transform
```

Description

This command specifies the encryption algorithm to use for the IPSec session. Encryption only applies to 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>dynamic-keying>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.

When AES-GCM is configured as the ESP encryption algorithm for IPSec transform:

- **esp-auth-algorithm** must be set to **auth-encryption**
- manual keying is not possible

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 3DES 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.

aes128-gcm8

configures ESP to use AES-GCM with a 128-bit key size and an 8-byte ICV for encryption and authentication

aes128-gcm12

configures ESP to use AES-GCM with a 128-bit key size and a 12-byte ICV for encryption and authentication

aes128-gcm16

configures ESP to use AES-GCM with a 128-bit key size and a 16-byte ICV for encryption and authentication

aes192-gcm8

configures ESP to use AES-GCM with a 192-bit key size and an 8-byte ICV for encryption and authentication

aes192-gcm12

configures ESP to use AES-GCM with a 192-bit key size and a 12-byte ICV for encryption and authentication

aes192-gcm16

configures ESP to use AES-GCM with a 192-bit key size and a 16-byte ICV for encryption and authentication

aes256-gcm8

configures ESP to use AES-GCM with a 256-bit key size and an 8-byte ICV for encryption and authentication

aes256-gcm12

configures ESP to use AES-GCM with a 256-bit key size and a 12-byte ICV for encryption and authentication

aes256-gcm16

configures ESP to use AES-GCM with a 256-bit key size and a 16-byte ICV for encryption and authentication

static-sa

Syntax

static-sa *sa-name* [**create**]

no static-sa *sa-name*

Context

config>ipsec

Description

This command configures an IPSec static security association (SA).

Default

no static-sa

Parameters

sa-name
specifies the name of the IPSec static SA, up to 32 characters

authentication

Syntax

authentication *auth-algorithm* **ascii-key** *ascii-string*
authentication *auth-algorithm* **hex-key** *hex-string* [**hash** | **hash2**]
no authentication

Context

config>ipsec>static-sa

Description

This command configures the authentication algorithm to use for the specified static SA.
The **no** form of the command resets to command to the default value.

Default

sha1

Parameters

auth-algorithm
specifies an authentication algorithm
Values md5 | sha1

ascii-string
specifies a string for an ASCII key
Values md5: must be 16 characters
 sha1: must be characters

hex-string
specifies a string for a hexadecimal key

Values md5: must be 2 hexadecimal nibbles
 sha1: must be 40 hexadecimal nibbles

hash

specifies that the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

specifies that the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

direction

Syntax

direction *ipsec-direction*

no direction

Context

config>ipsec>static-sa

Description

This command configures the direction for the specified static SA.

The **no** form of the command resets the command to the default value.

Default

bidirectional

Parameters

ipsec-direction

specifies the direction in which this static SA entry can be applied

Values inbound | outbound | bidirectional

protocol

Syntax

protocol *ipsec-protocol*

no protocol

Context

config>ipsec>static-sa

Description

This command configures the security protocol to use for the specified static SA. The **no** form of the command resets the command to the default value.

Default

esp

Parameters

ipsec-protocol

specifies the IPsec protocol used with this static SA

Values	ah – specifies the authentication header protocol
	esp – specifies the encapsulating security payload protocol

spi

Syntax

spi *spi*

no spi

Context

config>ipsec>static-sa

Description

This command configures the security parameter index (SPI) key value for the specified IPsec SA.

The SPI is used to look up the instruction to verify and decrypt the incoming IPsec packets when the value of the **direction** command is **inbound**.

The SPI value specifies the SPI that is used in the encoding of the outgoing packets when the value of the **direction** command is **outbound**. The remote node can use this SPI to look up the instruction to verify and decrypt the packet.

If no SPI is configured, the static SA cannot be used. The **no** form of the command removes the configured SPI.

Default

none

Parameters

spi

specifies the SPI for this SA

Values 256 to 16383

8.10.2.1.4 Service configuration commands

ipsec

Syntax

ipsec

Context

config>service>vpn

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>vpn>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>vpn>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/prefix-length* | *ip-prefix netmask* | **any**}

no local-ip

Context

config>service>vpn>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 evaluates 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 is 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/prefix-length | ip-prefix

the local IPv4 address or the IPv4 address and prefix length

netmask

the subnet mask in dotted-decimal notation

any

keyword to specify that it can be any address

local-v6-ip

Syntax

local-v6-ip {*ipv6-prefix/prefix-length | any*}

no local-v6-ip

Context

config>service>vpn>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 evaluates 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 is 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 and prefix length

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>vpn>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 evaluates 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 is 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/prefix-length | *ip-prefix*

the remote IP address or the IP address and prefix length

netmask

the subnet mask in dotted-decimal notation

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>vpn>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 evaluates 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 is 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/prefix-length

the remote IPv6 address and prefix length

any

keyword to specify that it can be any address

8.10.2.1.5 Service interface tunnel commands

interface

Syntax

interface *ip-int-name* [**tunnel**] [**create**]

no interface *ip-int-name*

Context

config>service>vpn

config>service>ies

Description

This command creates a logical IP routing interface.

When creating tunnel interfaces, the **tunnel** keyword must be used for private-side (VPRN) interfaces. The **tunnel** keyword is not used for public-side (IES or VPRN) interfaces.

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. An IES or VPRN public tunnel SAP is created when the *sap-id* includes the **tunnel** and **public** keywords. 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

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- <i>id</i> . <i>[private public]:tag</i>
	tunnel keyword
	<i>id</i> 1 to 16 (only the value 1 is allowed)
	private keyword
	public keyword
	<i>tag</i> 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>vpn>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>vpn>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 goes 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>vpn>if>sap>ipsec-tunnel

Description

This command assigns a BFD session to provide the heartbeat mechanism for the specified IPSec tunnel. There can be only one BFD session assigned to any specific 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 also brings 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-name

specifies the name of the interface used by the BFD session

Values 1 to 32 characters (must start with a letter)

ip-address

the IPv4 destination address to be used by the BFD session

clear-df-bit

Syntax

[no] clear-df-bit

Context

config>service>vpn>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>vpn>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>vpn>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 has 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>vpn>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

cert**Syntax**

cert

Context

config>service>vpn>if>sap>ipsec-tunnel>dynamic-keying

Description

This command enters the context to configure IPSec tunnel certificate parameters

Default

n/a

cert-profile**Syntax**

cert-profile *profile-name* [create]

no cert-profile *profile-name*

Context

```
config>service>vpn>if>sap>ipsec-tunnel>dynamic-keying>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

remote-id

Syntax

remote-id type {**ipv4** | **ipv6** | **fqdn** | **email**} **value** *value*

no remote-id

Context

```
config>service>vpn>if>sap>ipsec-tunnel>dynamic-keying
```

Description

This command configures a remote ID that is used to compare and authenticate an incoming X.509v3 certificate. The certificate is authenticated if the type and value in the Subject Alternative Name field of the incoming certificate match the configured remote ID type and value. If the fields do not match, the certificate is not processed.

When this command is configured, the [match-peer-id-to-cert](#) command must be disabled because the configurations are mutually exclusive.

Default

no remote-id

Parameters

type

specifies the type of remote ID payload

Values **ipv4**: specifies IPv4 as the remote ID type
 ipv6: specifies IPv6 as the remote ID type
 fqdn: specifies FQDN as the remote ID type

email: specifies an email address as the remote ID type

value

specifies an IPv4 or IPV6 address, an FQDN value, or an email address

- Values**
- ipv4-address*
 - ipv6-address*
 - fqdn*: a fully qualified domain name value (for example, "myhost.example.com"), up to 255 characters maximum
 - email*: an email address, up to 255 characters maximum

status-verify

Syntax

status-verify

Context

config>service>vpn>if>sap>ipsec-tunnel>dynamic-keying>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>vpn>if>sap>ipsec-tunnel>dynamic-keying>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 | ocsf}

no primary

Context

config>service>vpn>if>sap>ipsec-tunnel>dynamic-keying>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

ocsf

the OCSP server is configured in the corresponding CA profile

secondary**Syntax**

secondary {crl | ocsf}

no secondary

Context

config>service>vpn>if>sap>ipsec-tunnel>dynamic-keying>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>vpn>if>sap>ipsec-tunnel>dynamic-keying>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**

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 is 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 | ipv6} **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 IPv6 address, or a fully qualified domain name (FQDN) string.

If **local-id** is configured, the tunnel local ID is set to the explicit **type** and **value** specified by the **local-id** command. If **local-id** is not configured, the tunnel 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.

ipv6: specifies IPv6 as the local ID type. The default value is the local gateway IP address.

value

specifies an IPv4 or IPV6 address, or an FQDN value

Values *ipv4-address*

ipv6-address

fqdn: specifies a fully qualified domain name value (for example, "myhost.example.com"), up to 255 characters maximum

pre-shared-key**Syntax****pre-shared-key** *key* [**hash** | **hash2**]**no pre-shared-key****Context**

config>service>vpn>if>sap>ipsec-tunnel>dynamic-keying

Description

This command specifies the pre-shared key (PSK), or secret passphrase, that is used to initiate the tunnel IKE session. If the **hash** or **hash2** parameter is not used, the key is a clear text key; otherwise, the key text is encrypted. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

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 ASCII characters. If the string contains special characters (such as #, \$, or spaces), the entire string must be enclosed within parentheses.

hash

specifies that the key is entered in an encrypted form

hash2

specifies that the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted

transform

Syntax

transform *transform-id* [*transform-id*...(up to 4 max)]

no transform

Context

config>service>vpn>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>vpn>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 VPN access interface and destined for an IPSec uplink will be fragmented if the incoming packet is larger than:

- the VPN 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
 - IPv6 IPSec overhead if NAT-T is enabled or disabled (a UDP header is not inserted for IPv6 IPSec)
IPv6 IPSec overhead = outer IPSec (40) + ESP (24) + trailer (17) + ICV (32) = 113 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 9000

local-gateway-address

Syntax

local-gateway-address *ip-address* **peer** *ip-address* **delivery-service** *service-id*

no local-gateway-address

Context

```
config>service>vpn>if>sap>ipsec-tunnel
```

Description

This command specifies the local gateway address used by the IPSec 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.

The **no** form of the command removes the configured information.

Parameters

ip-address

the IPv4 or IPv6 address of the local and peer ends of the tunnel

service-id

specifies the ID of the IES or VPRN (front-door) delivery service of this IPSec 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 association (SA) 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>vpn>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 to 0xFFFFFFFF...(max 128 hex nibbles)

encryption-key

specifies the key used for the encryption algorithm

Values none or 0x0 to 0xFFFFFFFF...(max 64 hex nibbles)

spi

specifies the security parameter index (SPI) 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 is 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 is 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>vpn>if>sap>ipsec-tunnel

Description

This command identifies an IPSec security policy (defined under the **vpn>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 uses

Values 1 to 8192

8.10.2.2 PKI configuration commands

- [X.509 and certificate commands](#)
- [PKI infrastructure commands](#)
- [IPSec PKI commands](#)

- [Automatic CRL update commands](#)

8.10.2.2.1 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, 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

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 clears 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 uses 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**, the system uses 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 uses 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 plaintext. 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 characters maximum
 local-url : *cflash-id/file-path*
 cflash-id : cf1:, cf2:, 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

<i>url-string</i>	the name of the key file
Values	<i>url-string</i> : <i>local-url</i> , 99 characters maximum <i>local-url</i> : <i>cflash-id/file-path</i> <i>cflash-id</i> : cf1:, cf2:, 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

<i>url-string</i>	the name of the key file in cf3:\system-pki\key that is used to generate a certificate request
Values	<i>url-string</i> : <i>local-url</i> , 99 characters maximum <i>local-url</i> : <i>cflash-id/file-path</i> <i>cflash-id</i> : cf1:, cf2:, 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 is 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 characters maximum
 local-url : *cflash-id/file-path*
 cflash-id : cf1:, cf2:, 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 is 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 aborts.

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

8.10.2.2.2 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** does not affect the current up and running **ipsec-tunnel** associated with the **ca-profile**; however, subsequent authentication fails.

Executing a **no shutdown** command in this context causes 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, at least keyCertSign and cRLSign are asserted.
- If the certificate is not a self-signing certificate, the system looks 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, the system proceeds with a warning message.
- If the certificate is not a self-signing certificate, the system looks 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, the system proceeds with a warning message.

If any of above checks fails, the **no shutdown** command fails.

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 accepts 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 accepts 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 is resolved by the DNS server configured (if configured) in the corresponding router context.

If the *service-id* is 0 or omitted, then the system tries to resolve the FQDN using the DNS server configured in bof.cfg. After resolution, the system first connects 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 fails.

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 is only 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 fails.

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

revocation-check

Syntax

revocation-check {*crl* | *crl-optional*}

Context

config>system>security>pki>ca-profile

Description

This command specifies the method used to check the revocation status of a certificate issued by the CA.

By default, the system uses the configured certificate revocation list (CRL) to check the revocation status. If the **revocation-check** is configured to use the CRL but the CRL does not exist, the system does not use the configured **secondary** method or **default-result** to verify the revocation status.

If the **revocation-check** is configured as **crl-optional** and the CRL does not exist, the system skips the check and the status is assumed to be good.

The CA profile must be shut down before the **revocation-check** configuration can be changed.

Default

crl

Parameters

crl

specifies to use the configured CRL to check the revocation status

crl-optional

specifies to skip the revocation status check if the CRL does not exist

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 verifies the configured **cert-file** and **crl-file**. If the verification fails, then the **no shutdown** command fails.

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

certificate-expiration-warning

Syntax

certificate-expiration-warning *hours* [**repeat** *repeat-hours*]

no certificate-expiration-warning

Context

config>system>security>pki

Description

This command enables the system to issue two types of warning messages related to certificate expiration:

- BeforeExp – a warning message issued before a certificate expires
- AfterExp – a warning message issued when a certificate expires

The *hours* parameter configures how many hours before a certificate expiry the system issues a BeforeExp message. For example, with **certificate-expiration-warning 5**, the system issues a BeforeExp message 5 hours before the certificate expires. The optional **repeat** parameter causes the system to repeat the BeforeExp message at the configured hourly intervals until the certificate expires.

To receive only the AfterExp message after the certificate has expired, set the *hours* parameter to 0.

There are several ways to clear BeforeExp and AfterExp warning messages.

- If the certificate is reloaded with the **admin>certificate>reload** command and the reloaded certificate has not expired, the AfterExp message is cleared. If the reloaded certificate is outside of the configured warning window, the BeforeExp message is also cleared.
- If the CA profile is shut down, both the BeforeExp and AfterExp messages for the corresponding certificates are cleared.
- If the **no certificate-expiration-warning** command is issued, all existing BeforeExp and AfterExp messages are cleared.
- If the **certificate-expiration-warning** command is configured so that any certificates are no longer in the warning window, the BeforeExp messages for the corresponding certificates are cleared.
- If the system time changes and the new time causes any certificate to no longer be in the warning window, the corresponding BeforeExp message is cleared. If the new time causes an expired certificate to become unexpired, the AfterExp message is cleared.

Default

no certificate-expiration-warning

Parameters

hours

the number of hours before a certificate expires that the system issues a BeforeExp message

Values 0 to 8760

repeat-hours

specifies the intervals at which the system repeats the BeforeExp message

Values 0 to 8760

crl-expiration-warning

Syntax

crl-expiration-warning *hours* [**repeat** *repeat-hours*]

no crl-expiration-warning

Context

config>system>security>pki

Description

This command enables the system to issue two types of warning messages related to CRL expiration:

- BeforeExp – a warning message issued before a CRL expires
- AfterExp – a warning message issued when a CRL expires

The *hours* parameter configures how many hours before a CRL expiry the system issues a BeforeExp message. For example, with **crl-expiration-warning 5**, the system issues a BeforeExp message 5 hours before the CRL expires. The optional **repeat** parameter causes the system to repeat the BeforeExp message at the configured hourly intervals until the CRL expires.

To receive only the AfterExp message after the CRL has expired, set the *hours* parameter to 0.

There are several ways to clear BeforeExp and AfterExp warning messages.

- If the CRL is reloaded with the **admin>certificate>reload** command and the reloaded file has not expired, the AfterExp message is cleared. If the reloaded file is outside of the configured warning window, the BeforeExp message is also cleared.
- If the CA profile is shut down, both the BeforeExp and AfterExp messages for the corresponding CRLs are cleared.
- If the **no crl-expiration-warning** command is issued, all existing BeforeExp and AfterExp messages are cleared.
- If the **crl-expiration-warning** command is configured so that the CRL file is no longer in the warning window, the BeforeExp message for the corresponding file is cleared.

- If the system time changes and the new time causes the CRL to no longer be in the warning window, the corresponding BeforeExp message is cleared. If the new time causes an expired CRL to become unexpired, the AfterExp message is cleared.

Default

no crl-expiration-warning

Parameters

hours

the number of hours before a CRL expires that the system issues a BeforeExp message

Values 0 to 8760

repeat-hours

specifies the intervals at which the system repeats the BeforeExp message

Values 0 to 8760

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

8.10.2.2.3 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 overrides the **trust-anchor-profile** configuration in the **config>service>vpn>if>sap>ipsec-tunnel>cert** context.

Default

no trust-anchor-profile

Parameters

profile-name

the **trust-anchor-profile** name

8.10.2.2.4 Automatic CRL update commands

crl-update

Syntax

crl-update **ca** *ca-profile-name*

Context

admin>certificate

Description

This command manually initiates a CRL update for the specified CA profile.

Automatic CRL update must be **shutdown** before this command can be issued.

Default

n/a

Parameters

ca-profile-name

the name of the CA profile

file-transmission-profile

Syntax

file-transmission-profile *name* [**create**]

no file-transmission-profile *name*

Context

config>system

Description

This command creates a new file transmission profile. The profile can be configured with transport parameters for protocols such as HTTP and additional file transmission options.

Default

n/a

Parameters

name

the file transmission profile name, up to 32 characters

create

keyword required when first creating the configuration context. When the context is created, you can navigate into the context without the **create** keyword.

ipv4-source-address**Syntax**

ipv4-source-address *ip-address*

no ipv4-source-address

Context

config>system>file-trans-prof

Description

This command specifies the IPv4 source address used for the transport protocol. The address should be a local interface.

The **no** form of this command reverts to the default IPv4 source address, typically the address of the egress interface.

Default

no ipv4-source-address

Parameters

ip-address

the IPv4 source address

ipv6-source-address**Syntax**

ipv6-source-address *ipv6-address*

no ipv6-source-address

Context

config>system>file-trans-prof

Description

This command specifies the IPv6 source address used for the transport protocol. The address should be a local interface.

The **no** form of this command reverts to the default IPv6 source address, typically the address of the egress interface.

Default

no ipv6-source-address

Parameters

ipv6-address

the IPv6 source address

redirection**Syntax**

redirection *level*

no redirection

Context

config>system>file-trans-prof

Description

This command allows the system to accept HTTP redirection responses and configures the maximum level of redirection. The router can send a new request to another server if the CRL files are not available or are temporarily available to another server.

Default

no redirection

Parameters

level

the maximum level of HTTP redirection

Values 1 to 8

retry**Syntax**

retry *count*

no retry

Context

config>system>file-trans-prof

Description

This command specifies the number of times the system attempts to reconnect to a server that returns no data in the time configured with the **timeout** command.

The **no** form of this command disables any retry attempt.

Default

no retry

Parameters

count
the maximum number of retry attempts
Values 1 to 256

router

Syntax

router *router-instance*
router service *vprn-service-name*

Context

config>system>file-trans-prof

Description

This command specifies the routing instance that the transport protocol uses.

Default

Base

Parameters

router-instance
the router instance used to establish the file transmission connection

Values {*router-name* | *vprn-svc-id*}

<i>router-name</i> :	Base or Management
	<i>router-name</i> is an alias used for input only and is automatically replaced with an ID value by the 7705 SAR
<i>vprn-svc-id</i> :	1 to 2147483647

vprn-service-name
the VPRN service name

timeout

Syntax

timeout *seconds*

Context

config>system>file-trans-prof

Description

This command configures how long the system waits to receive any data from a server, such as an HTTP server. If no data is received before the timeout period expires, the system attempts to reconnect to the server if the file transmission profile is configured for one or more retries with the **retry** command.

Default

60 s

Parameters

seconds

the connection timeout for the file transmission

Values 1 to 3600

auto-crl-update

Syntax

auto-crl-update [create]

no auto-crl-update

Context

config>system>security>pki>ca-profile

Description

This command creates the context to configure automatic CRL update parameters.

When automatic CRL update is configured and enabled with the **no shutdown** command, the system downloads a CRL file from a list of configured HTTP URLs, either periodically or before an existing CRL expires. If the downloaded CRL is a valid CRL signed by the CA and is more recent than the existing CRL, the existing CRL is replaced.

The **no** form of this command deletes the automatic CRL update context and any configurations inside it.

Default

n/a

Parameters

create

keyword required when first creating the configuration context. When the context is created, you can navigate into the context without the **create** keyword.

crl-urls

Syntax

crl-urls

Context

config>system>security>pki>ca-prof>auto-crl-update

Description

This command enables the context to configure CRL URL parameters. Up to eight URL entries can be configured under each CA profile. The configured URLs must point to a DER-encoded CRL file.

When a CRL update is initiated, the system accesses each URL in order, and the first successfully downloaded and qualified CRL is used to update the existing CRL. If the download fails or the downloaded CRL is not qualified, the system moves to the next URL in the list. If no CRL file is successfully downloaded or qualified, the system attempts to contact each URL again at the next scheduled update time (when the schedule type is configured as **periodic**) or after the time configured with the **retry-interval** command (when the schedule type is configured as **next-update-based**).

The CRL download can be manually interrupted by issuing the **shutdown** command in the **auto-crl-update** context.

Default

n/a

url-entry

Syntax

url-entry *entry-id* [**create**]

no url-entry *entry-id*

Context

config>system>security>pki>ca-prof>auto-crl-update>crl-urls

Description

This command creates a new CRL URL entry or enters an existing URL entry configuration context.

The **no** form of this command removes the specified entry.

Default

n/a

Parameters*entry-id*

the URL entry identifier

Values 1 to 8**create**

keyword required when first creating the URL entry. When the URL entry is created, you can navigate into the context without the **create** keyword.

file-transmission-profile**Syntax****file-transmission-profile** *profile-name***no file-transmission-profile****Context**

config>system>security>pki>ca-prof>auto-crl-update>crl-urls>url-entry

Description

This command specifies an existing file transmission profile to use when the system downloads a CRL from the configured URL in this URL entry. The profile must already be configured with the **config>system>file-transmission-profile** command.

Automatic CRL update supports base, management, or VPRN routing instances. If VPRN is used, the HTTP server port can only be 80 or 8080.

The **no** form of this command removes the file transmission profile name from the URL entry.

Default

no file-transmission-profile

Parameters*profile-name*

the name of the file transmission profile to be used

url**Syntax****url** *url***no url**

Context

config>system>security>pki>ca-prof>auto-crl-update>crl-urls>url-entry

Description

This command specifies the HTTP URL of the CRL file for the URL entry. The system supports both IPv4 and IPv6 HTTP connections. The URL must point to a DER-encoded CRL.

The **no** form of this command removes the URL from the URL entry.

Default

no url

Parameters

url

specifies the location of a CRL to be downloaded

periodic-update-interval

Syntax

periodic-update-interval [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

Context

config>system>security>pki>ca-prof>auto-crl-update

Description

This command specifies the interval between automatic CRL updates when the **schedule-type** command is configured as **periodic**. The minimum interval is 1 hour. The maximum interval is 366 days.

Default

1 day

Parameters

days

specifies the number of days for periodic updates

Values 0 to 366

hours

specifies the number of hours for periodic updates

Values 0 to 23

minutes

specifies the number of minutes for periodic updates

Values 0 to 59

seconds

specifies the number of seconds for periodic updates

Values 0 to 59

pre-update-time

Syntax

pre-update-time [*days days*] [*hrs hours*] [*min minutes*] [*sec seconds*]

Context

config>system>security>pki>ca-prof>auto-crl-update

Description

This command specifies how much time before the next update time that the CRL is downloaded when the **schedule-type** command is configured as **next-update-based**.

Default

1 hr

Parameters

days

specifies how many days before the next CRL update that the CRL is downloaded

Values 0 to 366

hours

specifies how many hours before the next CRL update that the CRL is downloaded

Values 0 to 23

minutes

specifies how many minutes before the next CRL update that the CRL is downloaded

Values 0 to 59

seconds

specifies how many seconds before the next CRL update that the CRL is downloaded

Values 0 to 59

retry-interval

Syntax

retry-interval *seconds*

no retry-interval**Context**

```
config>system>security>pki>ca-prof>auto-crl-update
```

Description

This command specifies how long the system waits before retrying the configured URL entry list when the **schedule-type** is configured as **next-update-based** and no qualifying CRL could be downloaded during a CRL update.

The **no** form of this command causes the system to retry immediately.

Default

3600 s

Parameters

seconds

specifies the time before retrying to update the CRL

Values 1 to 31622400

schedule-type**Syntax**

schedule-type *schedule-type*

Context

```
config>system>security>pki>ca-prof>auto-crl-update
```

Description

This command configures the automatic CRL update schedule. The system supports two types:

- **periodic** – the system initiates a CRL update periodically, at the intervals specified by the **periodic-update-interval** command. The minimum periodic update interval is 1 hour.
- **next-update-based** – the system initiates a CRL update at the date and time specified in the Next Update field of the existing CRL file, minus the time configured with the **pre-update-time** command.

Default

next-update-based

Parameters

schedule-type

the schedule type for automatic CRL updates

Values periodic or next-update-based

shutdown

Syntax

[no] shutdown

Context

config>system>security>pki>ca-profile>auto-crl-update

Description

This command disables automatic CRL update.

The **no** form of this command enables automatic CRL update. If the **no shutdown** command is issued, the system immediately initiates a CRL update if the configured CRL file does not exist or is invalid or expired, or if the schedule type is configured as **next-update-based** and the scheduled update time has already passed.

Default

shutdown

8.10.2.3 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.

Output example

```
*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
Revoke Chk     : crl-optional 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.

Output example

```
*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.

Output example

```
*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.

Output example

```
*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 specified IKE policy. When an *ike-policy-id* is not specified, a summary display showing all IKE policies is displayed. When an *ike-policy-id* is specified, 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 194: IPSec IKE policy field descriptions](#) describes the fields.

Output example

```
*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
PRF Algorithm  : Same-As-Auth
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
Peer to Cert   : No-Match
IKEv2 Fragment : Disabled
```

Table 194: IPSec IKE policy field descriptions

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 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
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 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

Label	Description
PRF Algorithm	The authentication algorithm used in an IKE policy for the pseudorandom function (PRF)
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 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
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
Peer to Cert	Indicates whether the Subject Alternative Name field matches the IKE identifier of the peer certificate
IKEv2 Fragment	Indicates whether IKEv2 fragmentation is enabled

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 specified security policy.

Parameters

- service-id*

specifies the service ID or name of the tunnel delivery service

Values1 to 2147483690 or *service-name*
- security-policy-id*

specifies the IPSec security policy entry that this service uses

Values1 to 8192

Output

The following output is an example of IPSec security policy information, and [Table 195: IPSec security policy field descriptions](#) describes the fields.

Output example

```
*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       10.11.11.11/32  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 PlcyId	Policy ParamsId	LocalIp	RemoteIp
20	1	1	any	any
20	1	2	10.11.11.11/32	10.10.10.10/32
No. of IPsec Security Policy Param Entries: 2				
*A:7705custDoc:Sar18>show>ipsec#				

Table 195: IPSec security policy field descriptions

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
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 196: IPSec transform field descriptions](#) describes the fields.

Output example

```
*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 196: IPSec transform field descriptions

Label	Description
IPsec Transforms	
TransformId	The identifier of the IPSec transform policy
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 197: IPSec tunnel field descriptions](#) describes the fields.

Output example

```
*A:7705custDoc:Sar18>show>ipsec# tunnel
=====
IPsec Tunnels
=====
TunnelName      LocalAddress      SvcId      Admn   Keying
SapId           RemoteAddress     DlvrySvcId Oper    Sec
                               Plcy
-----
vprn_ipsec_tunnel      10.0.0.0          20         Down   Manual
tunnel-1.private:1     10.10.0.0         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    : 10.0.0.0
Remote Address   : 10.0.0.0
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   : 10.11.11.11
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
Remote ID        : www.nokia.com
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 197: IPSec tunnel field descriptions

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
DivrySvcId	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
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

Label	Description
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	
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

Label	Description
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
Remote ID	The remote ID value, if configured, with remote-id
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
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

8.10.2.4 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

8.10.2.5 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 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 the 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. Multiple IPSec tunnels can be debugged at the same time; up to 16 instances of this command can run concurrently.

Parameters

ipsec-tunnel-name

specifies an IPSec tunnel name up to 32 characters in length

detail

enables detailed debug information

9 Network group encryption

This chapter provides information to configure network group encryption (NGE).

Topics in this chapter include:

- [NGE overview](#)
- [Key groups](#)
- [Services encryption](#)
- [Router interface encryption](#)
- [Layer 2 encryption](#)
- [NGE packet overhead and MTU considerations](#)
- [1588v2 encryption with NGE](#)
- [QoS for NGE traffic](#)
- [Statistics](#)
- [Remote network monitoring support](#)
- [Configuration notes](#)
- [Configuring NGE with CLI](#)
- [NGE command reference](#)

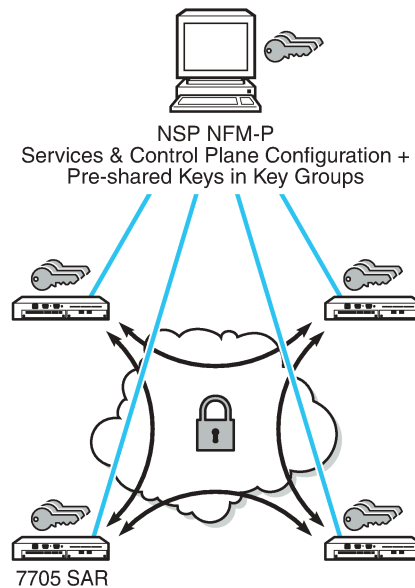
9.1 NGE overview

The network group encryption (NGE) feature enables end-to-end encryption of MPLS services, Layer 3 user traffic, and IP/MPLS control traffic. NGE is an encryption method that uses a group-based keying security architecture, which removes the need to configure individual encryption tunnels to achieve network-wide encryption.

NGE relies on the NSP NFM-P to manage the network and apply encryption to specific MPLS services, Layer 3 user traffic, or control plane traffic depending on the security requirements of the network. Operators designate traffic types that require added security and then apply NGE to those traffic types using the NSP NFM-P. The NSP NFM-P also acts as the network-wide NGE key manager, downloading encryption and authentication keys to nodes and performing hitless rekeying of the network at operator-defined intervals. For more information about managing NGE within a network, see the NSP NFM-P User Guide.

The following figure shows an NGE network with NSP NFM-P services, control plane configuration, and key management.

Figure 137: NGE network with NSP NFM-P management



26216

NGE is enabled on the 7705 SAR and provides three main types of encryption to secure an IP/MPLS network:

- SDP encryption – MPLS user plane encryption enabled on MPLS tunnels (SDPs) supporting VPRN or IES services using spoke SDPs, VPLS using spoke or mesh SDPs, routed VPLS into VPRN, Epipes, and Cpipes
- VPRN encryption
 - unicast VPRN – MP-BGP-based VPRN-level encryption using spoke SDPs (**spoke-sdp**) or autobind SDPs (**auto-bind-tunnel**) with LDP, GRE, RSVP-TE, or segment routing (SR-ISIS, SR-OSPF, or SR-TE) tunnels
 - multicast VPRN – NG-MVPN using mLDP with auto-discovery
- router interface and Ethernet port Layer 2 encryption – Layer 3 user plane and control plane encryption, Layer 2 encryption of IS-IS and LLDP

NGE is supported on the following adapter cards and platforms:

- 8-port Gigabit Ethernet Adapter card, version 3
- 6-port Ethernet 10Gbps Adapter card
- 10-port 1GigE/1-port 10GigE X-Adapter card, version 2
- 7705 SAR-Ax
- 7705 SAR-H (Ethernet and PPP/MLPPP network links)
- 7705 SAR-Hc
- 7705 SAR-Wx
- 7705 SAR-X (Ethernet and PPP/MLPPP network links)

This section contains information about the following topics:


- [NGE key groups and encryption partitions](#)
- [NGE domains](#)
- [Network Services Platform management](#)

9.1.1 NGE key groups 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, router interfaces, or Ethernet ports to use specific key groups depending on security policies for the service and network topology.

The following figure shows a typical application of NGE key group partitioning in which 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 may be considered less critical than the transmission or distribution substation network equipment. It may be ideal to ensure that nodes more at risk of a security breach do not contain more critical information than is necessary. Therefore, encryption keys for the sensitive portions of the network (such as control center traffic) should not be available on nodes that are at risk. The 7705 SAR NGE feature enables operators to partition and distribute encryption keys among different services, NGE domains, or nodal groups in a network. NGE partitions are enabled by configuring different key groups per security partition and applying those key groups as needed.

Figure 138: Key group partitioning



25080

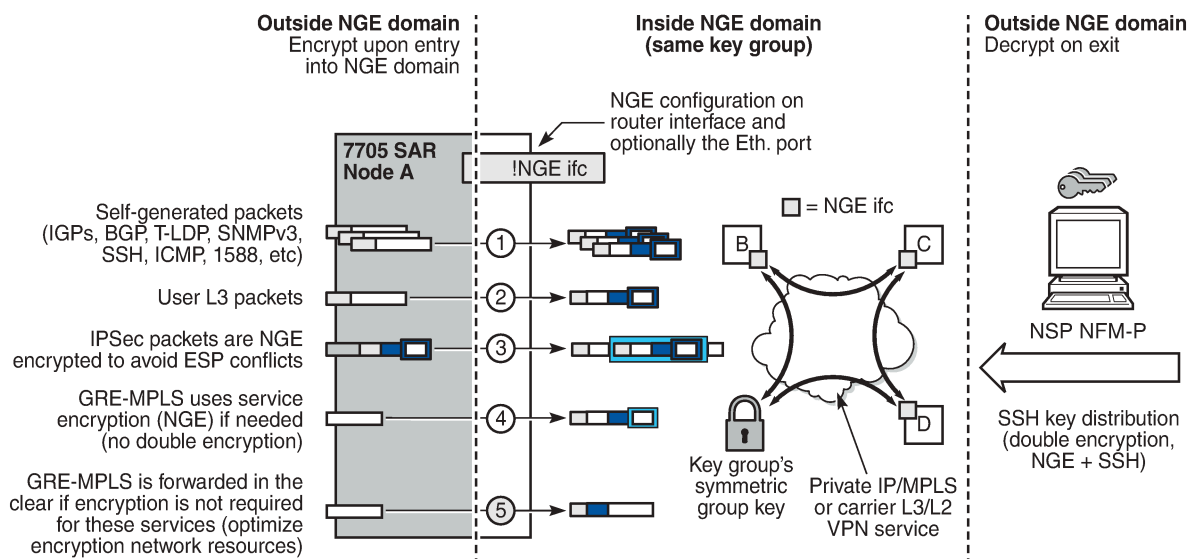
Another application of key group partitioning allows different parts of an organization to have their own method of end-to-end communication without the need to share encryption keys between each organization. If two partitions need to communicate between themselves, gateway nodes configured with both key groups allow inter-organization traffic flows between the key group partitions, as needed.

9.1.2 NGE domains

An NGE domain is a group of nodes and router interfaces forming a network that uses a single key group to create a security domain. NGE domains are created when router interface encryption is enabled on router interfaces that need to participate in the NGE domain. The NSP NFM-P assists operators in managing the nodes and interfaces that participate in the NGE domain. See the NSP NFM-P User Guide for more information.

The following figure shows various traffic types crossing an NGE domain.

Figure 139: NGE domain transit



26217

In the figure, nodes A, B, C, and D have router interfaces configured with router interface encryption enabled (and optionally Layer 2 NGE on the Ethernet port). Traffic is encrypted when entering the NGE domain using the key group configured on the router interface and is decrypted when exiting the NGE domain.

Traffic may traverse multiple hops before exiting the NGE domain, yet decryption only occurs on the final node when the traffic exits the NGE domain. Various traffic types are supported and encrypted when entering the NGE domain, as illustrated by the following items on node A in the figure:

- item 1: self-generated packets – these packets, which include all types of control plane and management packets such as OSPF, BGP, LDP, SNMPv3, SSH, ICMP, RSVP-TE, and 1588, are encrypted
- item 2: user Layer 3 packets – any Layer 3 user packets that are routed into the NGE domain from an interface outside the NGE domain are encrypted
- item 3: IPsec packets – IPsec packets are NGE-encrypted when entering the NGE domain to ensure that the IPsec packet's security association information does not conflict with the NGE domain

GRE-MPLS-based service traffic consists of Layer 3 packets, and router interface NGE is not applied to these types of packets. Instead, service-level NGE is used for encryption to avoid double-encrypting these packets and impacting throughput and latencies. The two types of GRE-MPLS packets that can enter the NGE domain are illustrated by items 4 and 5 in the figure:

- item 4: GRE-MPLS packets (SDP or VPRN) with service-level NGE enabled – these encrypted packets use the key group that is configured on the service. The services key group may be different from the key group configured on the router interface where the GRE-MPLS packet enters the NGE domain.
- item 5: GRE-MPLS packets (SDP or VPRN) with NGE disabled – these packets are not encrypted and can traverse the NGE domain in clear text. If these packets require encryption, SDP or VPRN encryption must be enabled.


Creating an NGE domain from the NSP NFM-P requires the operator to determine the type of NGE domain being managed. This will indicate whether NGE gateway nodes are required to manage the NGE domain, and other operational considerations. The two types of NGE domains are:

- [Private IP/MPLS network NGE domain](#)
- [Private over intermediary network NGE domain](#)

9.1.2.1 Private IP/MPLS network NGE domain

One type of NGE domain is a private IP/MPLS network, as shown in the following figure.

Figure 140: Private IP/MPLS network NGE domain



26215


In a private IP/MPLS network NGE domain, all interfaces are owned by the operator and there is no intermediary service provider needed to interconnect nodes. Each interface is a point-to-point private link between private nodes. When a new node is added to this type of NGE domain (node D in the figure), the links that connect node D to the existing nodes in the NGE domain (nodes A, B, and C) must be enabled with NGE router interface encryption. Links from the new node to the existing nodes are enabled one at a time. The NSP NFM-P provides tools that simplify adding nodes to the NGE domain and enabling NGE on their associated interfaces. In this type of NGE domain, each interface is a direct link between two nodes and is not used to communicate with multiple nodes over a broadcast medium offered by an intermediary network. Also, there are no NGE gateway nodes required between the NSP NFM-P and new nodes entering the NGE domain.

9.1.2.2 Private over intermediary network NGE domain

The other type of NGE domain is a private IP/MPLS network that traverses an intermediary network NGE domain; the intermediary network is used to interconnect nodes in the NGE domain using a multipoint-to-multipoint service. The intermediary network is typically a service provider network that provides a private IP VPN service or a private VPLS service used to interconnect a private network that does not mimic point-to-point links as described in [Private IP/MPLS network NGE domain](#).

This type of NGE domain is shown in the following figure.

Figure 141: Private over intermediary network NGE domain



26214

Private over intermediary network NGE domains have nodes with links that connect to a service provider network where a single link can communicate with multiple nodes over a Layer 3 service such as a VPRN or a Layer 2 service such as VPLS. In the figure, node A has NGE enabled on its interface with the service provider and uses that single interface to communicate with nodes B and C, and eventually with node D when node D has been added to the NGE domain. This type of NGE domain requires the recognition of NGE gateway nodes that allow the NSP NFM-P to reach new nodes that enter the domain. Node C is designated as a gateway node.

When node D is added to the NGE domain, it must first have the NGE domain key group downloaded to it from the NSP NFM-P. The NSP NFM-P creates an NGE exception ACL on the gateway node, C, to allow communication with node D using SNMPv3 and SSH through the NGE domain. After the key group is downloaded, the NSP NFM-P enables router interface encryption on node D's interface with the service provider and node D is now able to participate in the NGE domain. The NSP NFM-P automatically removes the IP exception ACL from node C when node D enters the NGE domain.

See [Router interface NGE domain concepts](#) for more information.

9.1.3 Network Services Platform management

The NGE feature is tightly integrated with the NSP NFM-P. The following functions are provided by the NSP NFM-P :

- managing and synchronizing encryption and authentication keys within key groups on a network-wide basis
- configuring NGE on MPLS services and managing associated key groups
- configuring NGE on router interfaces and Ethernet ports and managing associated key groups
- coordinating network-wide rekeying of key groups

The NSP NFM-P acts as the key manager for NGE-enabled nodes and allocates the keys in key groups that are used to perform encryption and authentication. The NSP NFM-P ensures that all nodes in a key group are kept in synchronization and that only the key groups that are relevant to the associated nodes are downloaded with key information.

The NSP NFM-P performs network-wide hitless rekeying for each key group at the rekeying interval specified by the operator. Different key groups can be rekeyed at different times as needed, or all key groups can be rekeyed network-wide at the same time.

For more information about NSP NFM-P management, see the “Service Management” section in the NSP NFM-P User Guide.

9.2 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:

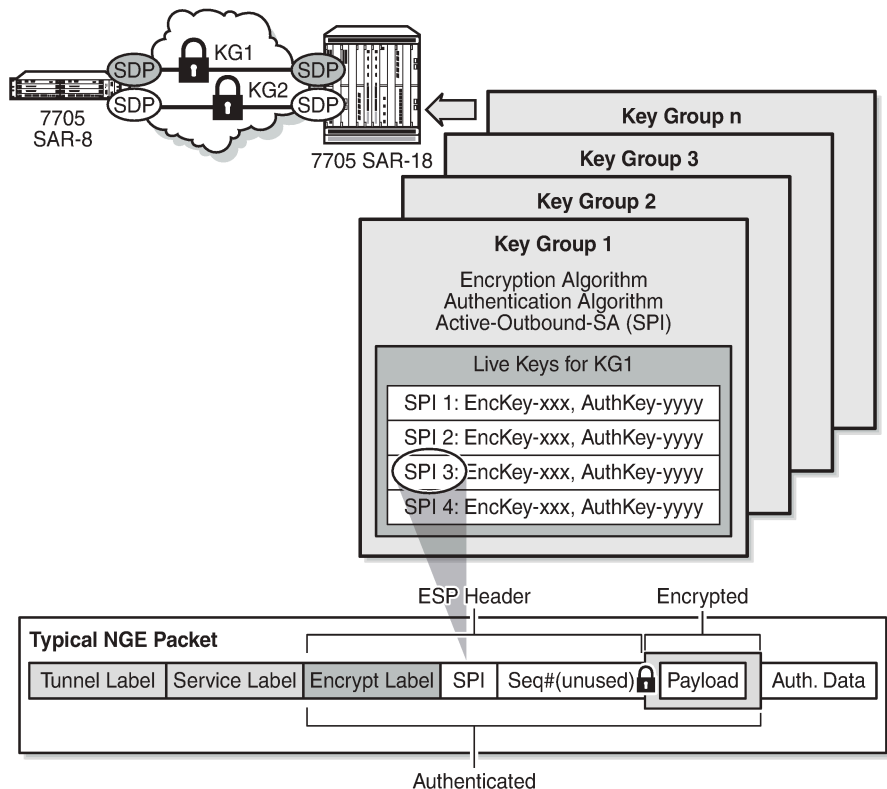
- an encryption algorithm – see [Key group algorithms](#)
- an authentication algorithm – see [Key group algorithms](#)
- a list of security associations (SAs) – see [Security associations](#)
- an active outbound SA – see [Active outbound SA](#)

The following figure illustrates the use of key groups (KGs), security associations (SAs), and security parameter indexes (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 142: Key groups and a typical NGE packet



25081

9.2.1 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. 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.

9.2.1.1 Encapsulating security payload

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)

9.2.2 Security associations

Each key group has a list of up to four security associations (SAs). An SA is a reference to a pair of encryption and authentication keys that are 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 if 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 142: Key groups and a typical NGE packet](#) illustrates these relationships.

As shown in the figure, each SA is referenced by an 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. When entered, they are never displayed in their original, clear text form. Keys are displayed in a 7705 SAR-encrypted form, which is indicated by the system-appended **crypto** keyword when an **info** command is run. The 7705 SAR also includes the **crypto** keyword with an **admin>save** operation so that the 7705 SAR 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).

9.2.2.1 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 ESP header of packets being encrypted and authenticated.

9.3 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
- NG-MVPN

For services that use SDPs, all tunnels may be either MPLS LSPs (RSVP-TE, LDP, or static LSP) or GRE tunnels. NGE is not supported for IP tunnels.

For MP-BGP services, resolving routes using spoke SDPs (**spoke-sdp**) or autobind SDPs (**auto-bind-tunnel**) is supported using LDP, GRE, RSVP-TE, or segment routing (SR-ISIS, SR-OSPF, or SR-TE).

In addition, the 7705 SAR supports VLL, VPLS, and VPRN NGE interactions that use segment routing with entropy labels.

This section contains information about 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 for NG-MVPN](#)

9.3.1 Services encryption overview

NGE adds a global encryption label to the label stack for encrypting MPLS services. The global encryption label must be a unique network-wide label; in other words, 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 those nodes.


The global encryption label is used to identify packets that have an NGE-encrypted payload 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.

When the global encryption label is set, it should not be changed. If the label must be changed without impacting traffic, all key groups in the system should first be deleted. Next, the label should be changed, and then all key groups should be reconfigured.

The NSP NFM-P helps to coordinate the distribution of the global encryption label and ensures that all nodes in the network are using the same global encryption label.

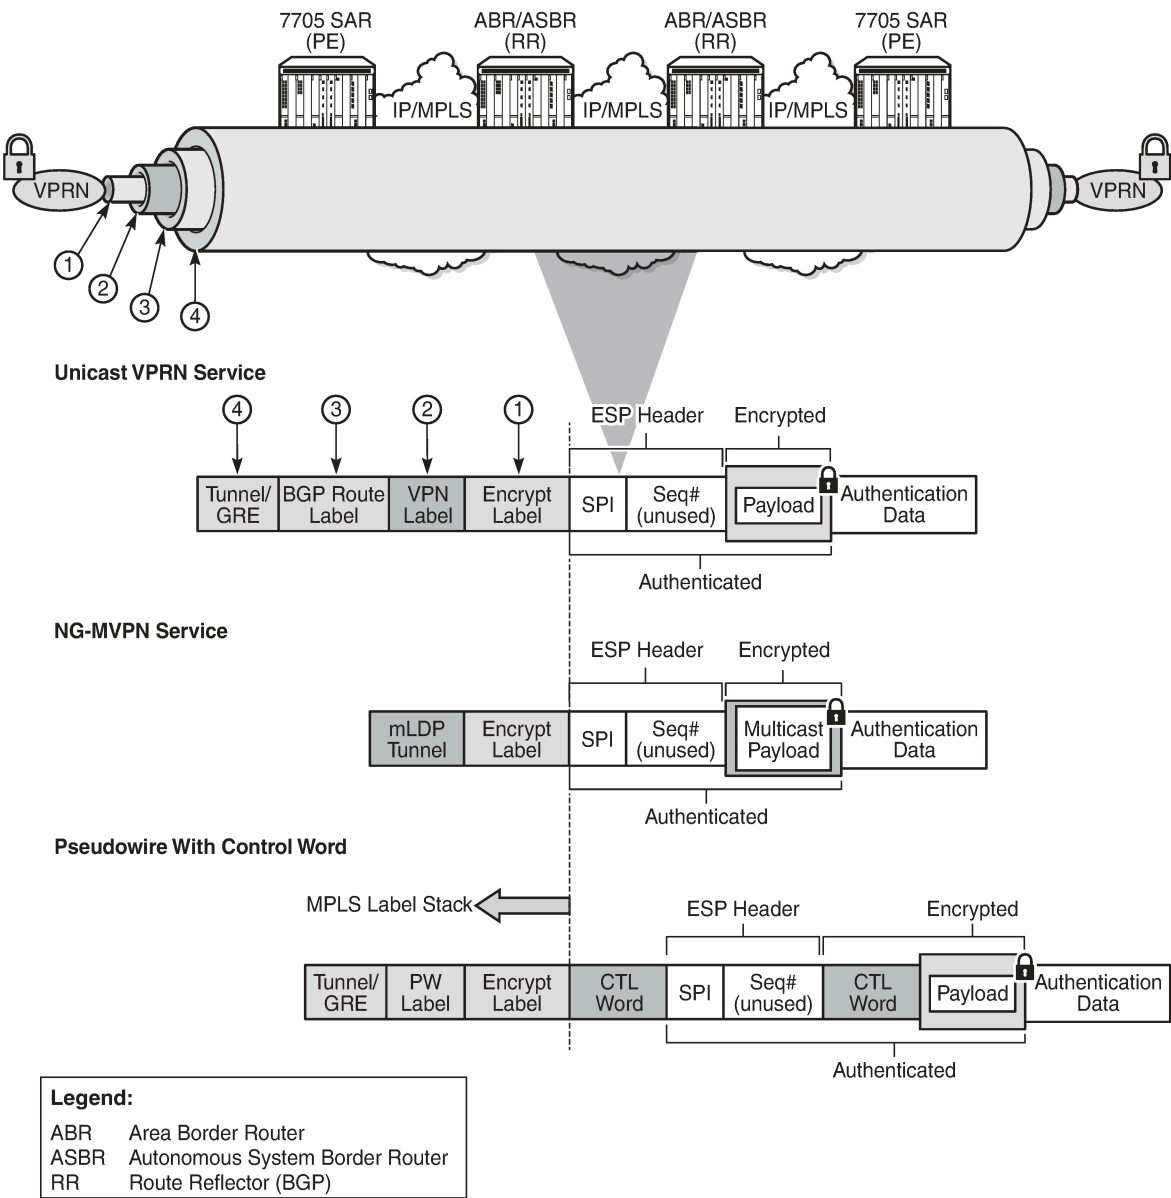
The following figure illustrates the NGE MPLS/GRE label stack.

Figure 143: NGE MPLS/GRE label stack



The following figure illustrates VPRN and PW (with control word) packet formats using NGE encryption.

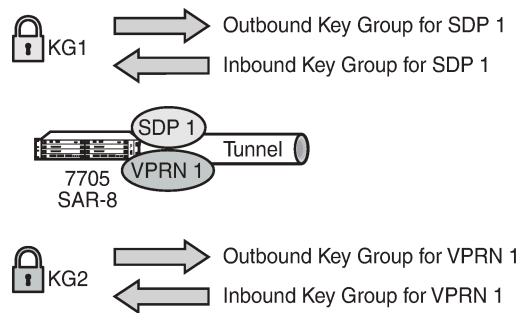
Figure 144: NGE encryption and packet formats



35843

9.3.2 Assigning key groups to services

Assigning key groups to services requires configuring an inbound and outbound key group for directional processing on a per-service basis, as shown in the following figure.

Figure 145: Inbound and outbound key group assignments

25082

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, the node ensures that packets are either unencrypted or are using one of the valid key groups configured in the system.

In most 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:

- gracefully enable and disable NGE for services
- move services from one key group domain to another domain without halting encryption

The NGE feature makes use of the NSP NFM-P to help manage the assignment of key groups to services on a network-wide basis. See the NSP NFM-P User Guide for more information.

9.3.3 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 switching 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 switching 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 switching 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. When traffic is switched from an unencrypted PW to an encrypted PW, the traffic is encrypted before it is sent.

See [Pseudowire switching](#) for more information.

9.3.4 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.

9.3.5 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 (either with spoke SDPs using **spoke-sdp** or autobind SDPs using **auto-bind-tunnel**) 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**.

9.3.6 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.

9.3.7 NGE for NG-MVPN

NGE is supported for NG-MVPN services with multicast configurations that include:

- I-PMSI
- S-PMSI
- C-multicast signaling
- mLDP multicast tunnel LSPs

For more information about MVPN, see [Multicast VPN \(MVPN\)](#).

When r-VPLS is configured for the VPRN, the node can act as a receiver for an NG-MVPN service that is NGE -encrypted. The node cannot act as a source of NG-MVPN traffic from a VPLS service, so the source of this NG-MVPN traffic must originate from another node that supports that capability.

When NGE is enabled in a VPRN with NG-MVPN-based services, transit nodes (LSRs) have no knowledge that NGE is being employed or that the NGE encryption label is being used with an ESP header after the NGE label. Features that inspect packet contents to make further decisions are not supported and must be disabled for mLDP multicast paths that need to carry NG-MVPN traffic that is NGE-encrypted. These features include:

- ingress multicast path management (7750 SR only)
- IP-based LSR hashing

The packet inspection restriction includes any third-party routing function that may inspect the packet contents after the mLDP transport label and is expecting a non-encrypted payload in order to make hashing decisions.

9.4 Router interface encryption


The 7705 SAR supports Layer 3 encryption on router interfaces for IPv4 traffic. NGE is not supported on dual-stack IPv4/IPv6 or IPv6-only interfaces.

NGE is enabled on a router interface by configuring the **group-encryption** command on the router interface. The interface is considered part of the NGE domain, and any received packets that are NGE-encrypted are decrypted if the key group is configured on the node. To encrypt packets egressing the interface, the outbound key group must be configured on the interface. All IP packets, such as self-generated traffic or packets forwarded from router interfaces that are not inside the NGE domain, are encrypted when egressing the interface. There are some exceptions to this general behavior, as described in the sections below; for example, GRE-MPLS packets are not encrypted when router interface encryption is enabled.

The outbound and inbound key groups configured on the router interface determine which keys to use to encrypt and decrypt traffic.

To perform encryption, router interface encryption reuses the IPSec transport mode packet format as shown in the following figure.

Figure 146: Router interface encryption packet format (IPSec transport mode)



The protocol field in the IP header of an NGE packet is always set to "ESP". Within an NGE domain, the SPI that is included in the ESP header is always an SPI for the key group configured on the router interface. Other fields in the IP header, such as the source and destination addresses, are not altered by NGE router interface encryption. Packets are routed through the NGE domain and decrypted when the packet leaves the NGE domain.

The group keys used on an NGE-enabled router interface provide encryption of broadcast and multicast packets within the GRT. For example, OSPF uses a broadcast address to establish adjacencies, which can be encrypted by NGE without the need to establish point-to-point encryption tunnels. Similarly, multicast packets are also encrypted without point-to-point encryption tunnels.

9.4.1 Router interface NGE domain concepts

An NGE domain is a group of nodes whose router interfaces in the base routing context (GRT) are enabled for router interface NGE or whose ports are enabled for encryption of some types of Layer 2 traffic. An interface without router interface NGE enabled is considered to be outside the NGE domain. NGE domains use only one key group when the domain is created; however, two key groups may be active at once if some links within the NGE domain are in transition from one key group to the other.

The following figure illustrates the NGE domain concept and the table describes the three configuration scenarios inside the NGE domain.

Figure 147: Inside and outside NGE domains




Table 198: Inside and outside NGE domains – configuration scenarios

Key	Description
1	NGE enabled, no inbound/outbound key group Outbound packets are sent without encrypting. Inbound packets can be NGE-encrypted or clear text.
2	Outbound key group, no inbound key group Outbound packets are encrypted using the interface key group if not already encrypted. Inbound packets can be NGE-encrypted or clear text.
3	Inbound and outbound key group Outbound packets are encrypted using the interface key group if not already encrypted. Inbound packets must be encrypted using the interface key group keys.
4	Outside the NGE domain, the interface is not configured for NGE. Any ESP packets are IPSec packets.

A router interface is considered to be inside the NGE domain when it has been configured with **group-encryption** on the interface. When **group-encryption** is configured on the interface, the router can receive unencrypted packets or NGE-encrypted packets from any configured key group on the router, but any other type of IPSec-formatted packet is not allowed. If an IPSec-formatted packet is received on an interface that has **group-encryption** enabled, it will not pass NGE authentication and will be dropped. Therefore, IPSec

packets cannot exist within the NGE domain without first being converted to NGE packets. This conversion requirement delineates the boundary of the NGE domain and other IPsec services.

When NGE router interface encryption is enabled and only an outbound key group is configured, the interface can receive unencrypted packets or NGE-encrypted packets from any configured key group on the router. All outbound packets are encrypted using the outbound key group if the packet was not already encrypted further upstream in the network.

When NGE router interface encryption has been configured with both an inbound and outbound key group, only NGE packets encrypted with the key group security association can be sent and received over the interface.

When there is no NGE router interface encryption, the interface is considered outside the NGE domain where NGE is not applied.

9.4.2 GRE-MPLS packets inside the NGE domain

NGE router interface encryption is never applied to GRE-MPLS packets that are used to transport MPLS services over Layer 3 networks; for example, GRE with the GRE protocol ID set to MPLS Unicast (0x8847) or Multicast (0x8848). GRE-MPLS packets that enter the NGE domain or transit the NGE domain are forwarded as is.

Because GRE-MPLS packets provide transport for MPLS-based services, they already use the NGE services-based encryption techniques for MPLS, such as SDP or VPRN-based encryption. To avoid double encryption, the packets are left in clear text when entering an NGE domain or crossing intermediate nodes in the NGE domain, and are simply forwarded as needed when exiting an NGE domain.

The payload of these GRE-MPLS packets can also be sent in clear text depending on the security requirements that are configured for the MPLS service.

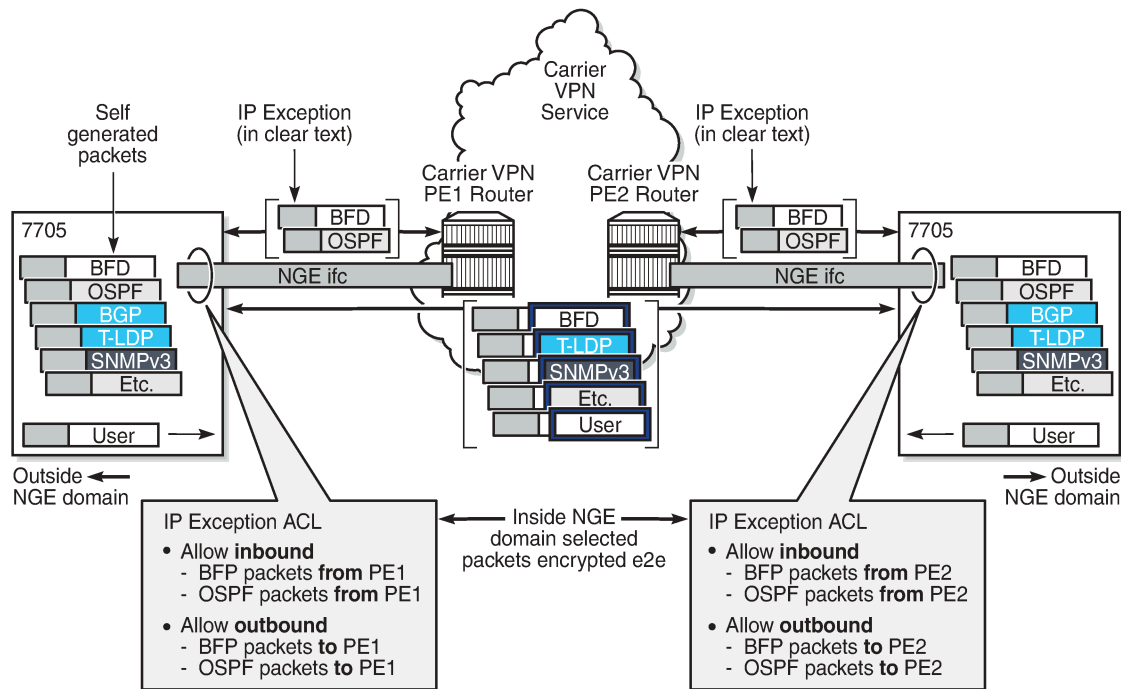
GRE packets with the protocol set to other protocols are encrypted using NGE router interface encryption.

9.4.3 Router encryption exceptions using ACLs

In some cases, Layer 3 packets may need to cross the NGE domain in clear text, such as when an NGE-enabled router needs to peer with a non-NGE-capable router to exchange routing information. This can be accomplished by using a router interface NGE exception filter applied on the router interface for the required direction, inbound or outbound.

The following figure shows the use of a router interface NGE exception filter.

Figure 148: Router interface NGE exception filter example



26277

The inbound or outbound exception filter is used to allow specific packet flows through the NGE domain in clear text, where there is an explicit inbound and outbound key group configured on the interface. The behavior of the exception filter for each router interface configuration is as follows:

- NGE enabled, no inbound/outbound key group – in this scenario, the router does not encrypt outbound traffic, and so the outbound exception filter is not applied. The router can still receive inbound NGE packets, so the exception filter is applied to inbound packets. If the filter detects a match, clear text packets can be received and forwarded by the router.
- outbound key group, no inbound key group – the outbound exception filter is applied to outbound traffic, and packets that match the filter are not encrypted on egress. The router can receive inbound NGE packets without an inbound key group set and applies the exception filter to inbound packets. If the filter detects a match, clear text packets can be received and forwarded by the router.
- inbound and outbound key group – the inbound and outbound exception filters are applied, and any packets that match are passed in clear text


9.4.4 IPSec packets crossing an NGE domain

IPSec packets can cross the NGE domain because they are still considered Layer 3 packets. To avoid confusion between the security association used in an IPSec packet and the one used in a router interface NGE packet, the router always applies NGE to any IPSec packet that traverses the NGE domain.

IPSec packets that originate from a router within the NGE domain are not allowed to enter the NGE domain. The only exception to this restriction is OSPFv3 packets.

The following figure shows how IPSec packets can transit an NGE domain.

Figure 149: IPsec packets transiting an NGE domain



26181

An IPsec packet enters the router from outside the NGE domain. When the router determines that the egress interface to route the packet is inside an NGE domain, it selects an NGE router interface with one of the following configurations:

- NGE enabled with no inbound or outbound key group configured – this link cannot forward the IPsec packet without adding the NGE ESP, but because nothing is configured for the outbound key group, the packet must be dropped
- NGE enabled with outbound key group configured and no inbound key group configured – the packet originates outside the NGE domain, so the router adds an ESP header over the existing ESP and encrypts the payload using the NGE domain keys for the configured outbound key group
- NGE enabled with both inbound and outbound key groups configured – the packet originates outside the NGE domain, so the router adds an ESP header over the existing ESP and encrypts the payload using the NGE domain keys for the configured outbound key group

OSPFv3 IPsec support also uses IPsec transport mode packets. These packets originate from the CSM, which is considered outside the NGE domain; however, the above rules for encapsulating the packets with an NGE ESP apply and allow these packets to successfully transit the NGE domain.

9.4.5 Multicast packets traversing the NGE domain


Multicast packets that traverse an NGE domain can be categorized into two main scenarios:

- scenario 1 – multicast packets that ingress the router on an interface that is outside the NGE domain. These packets can egress a variety of interfaces that are either inside or outside the NGE domain.

- scenario 2 – multicast packets that ingress the router on an interface that is inside the NGE domain. These packets can egress a variety of interfaces that are either inside or outside the NGE domain. This scenario has two cases:
 - scenario 2a – the ingress multicast packet is not yet NGE-encrypted
 - scenario 2b – the ingress multicast packet is NGE-encrypted

The following figure shows these scenarios.

Figure 150: Processing multicast packets



Multicast packets received from outside the NGE domain (scenario 1) are processed similarly to multicast packets received from inside the NGE domain (scenarios 2a and 2b).

The processing rule is that multicast packets are always forwarded as clear text over the fabric. This means that for scenario 2b, when a multicast packet is received on an encryption-capable interface and is NGE-encrypted, the packet is always decrypted first so that it can be processed in the same way as packets in scenarios 1 and 2a.

On egress, the following scenarios apply:

- egressing an interface outside the NGE domain – packets are processed in the same way as any multicast packets forwarded out a non-NGE interface
- egressing an NGE router interface and no inbound or outbound key group is configured – the router forwards these packets out from the egress interface without encrypting them, because there is no outbound key group configured. This behavior also applies to unicast packets in the same scenario.
- egressing an NGE router interface with the outbound key group configured – the router encrypts the multicast packet using the SPI keys of the outgoing SA configured in the key group. This behavior also applies to unicast packets in the same scenario.

9.4.6 Assigning key groups to router interfaces

Assigning key groups to router interfaces involves the following three steps:

1. Enable NGE with the **group-encryption** command.
2. Configure the outbound key group.
3. Configure the inbound key group.

Step 1 is required so that the router can initialize and differentiate the interface for NGE traffic before accepting or sending NGE packets. This assigns the interface to an NGE domain, and IPSec packets will no longer be accepted on the interface.

Assigning key groups to a router interface in steps 2 and 3 is similar to assigning key groups to SDPs or VPRN-based services. An outbound key group cannot be configured for a router interface without first enabling **group-encryption**.


When group-encryption is enabled and no inbound key group is configured, the router accepts NGE Layer 3 packets that were encrypted using keys from any security association configured in any key group on the system. If the packet specifies a security association that is not configured in any key group on the node, the packet is dropped.

The outbound key group references the key group to use when traffic egresses the router on the router interface. The inbound key group is used to make sure ingress traffic is using the correct key group on the router interface. If ingress traffic is not using the correct key group, the router counts these packets as errors.

9.4.7 Router interface NGE firewall considerations

An interface that has been configured to exist inside the NGE domain cannot be added to a security zone; firewalls are not supported on NGE router interfaces, as shown in the following figure.

Figure 151: Firewall considerations



26179

It is recommended that firewall security policies be enabled on interfaces outside the NGE domain before traffic is allowed to enter the NGE domain.

9.4.8 NGE and BFD support

When NGE is enabled on a router interface, BFD packets that originate from the network processor on the adapter card or from the system are encrypted in the same way as BFD packets that are generated by the CSM.

9.4.9 NGE and ACL interactions

When NGE is enabled on a router interface, the ACL function is applied as follows:

- on ingress – normal ACLs are applied to traffic received on the interface that can be either NGE-encrypted or clear text. For NGE-encrypted packets, this implies that only the source, destination, and IP options are available to filter on ingress because the protocol is ESP and the packet is encrypted. If an IP exception ACL is also configured on the interface, the IP exception ACL is applied first to allow any clear text packets to ingress as needed. After the IP exception ACL is applied and if another filter or ACL is configured on the interface, the other filter processes the remaining packet stream (NGE-encrypted and IP exception ACL packets), and other ACL functions such as PBR or Layer 4 information filtering could be applied to any clear text packets that passed the exception ACL.
- on egress – ACLs are applied to packets before they are NGE-encrypted as per normal operation without NGE enabled

9.4.10 Router interface NGE and ICMP interactions over the NGE domain

Typically, ICMP works as expected over an NGE domain when all routers participating in the NGE domain are NGE-capable; this includes running an NGE domain over a private IP/MPLS network and over a Layer 2 service provider where dedicated point-to-point circuits are used to create single-hop links between NGE nodes. When an ICMP message is required, the NGE packet is decrypted first and the original packet is restored to create a detailed ICMP message using the original packet's header information.

When the NGE domain crosses a Layer 3 service provider or crosses over routers that are not NGE-aware, it is not possible to create a detailed ICMP message using the original packet's information because the NGE packet protocol is always set to ESP. Furthermore, the NGE router that receives these ICMP messages will drop them because the messages are not NGE-encrypted.

The combination of dropping ICMP messages at the NGE border node and the missing unencrypted packet details in the ICMP information can cause problems with diagnosing network issues.

To help with diagnosing network issues, additional statistics are available on the interface to show whether ICMP messages are being returned from a foreign node. The following statistics are included in the group encryption NGE statistics for an interface:

- Group Enc Rx ICMP DestUnRch Pkts
- Group Enc Rx ICMP TimeExc Pkts
- Group Enc Rx ICMP Other Pkts

These statistics are used when clear text ICMP messages are received on an NGE router interface. The Invalid ESP statistics are not used in this situation even though the packet does not have a correct NGE ESP header. If there is no ingress exception ACL configured on the interface to allow the ICMP messages to be forwarded, the messages are counted and dropped.

If more information is required for these ICMP messages, such as source or destination address information, a second ICMP filter can be configured on the interface to allow logging of the ICMP messages. If the original packet information is also required, an egress exception ACL can be configured with the respective source or destination address information, or other criteria, to allow the original packet to enter the NGE domain in clear text and determine which flows are causing the ICMP failures.

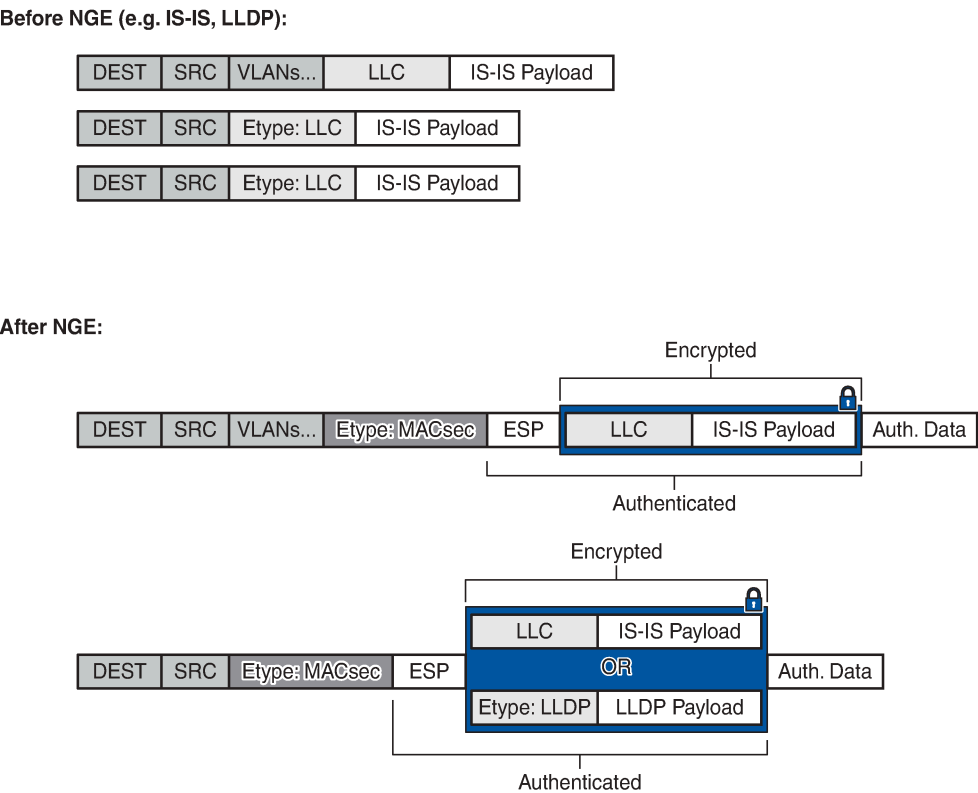
9.4.11 OAM considerations for router interface encryption

CSM timestamping is supported when NGE router interface encryption is enabled.

9.5 Layer 2 encryption

Layer 2 encryption allows IS-IS and LLDP Layer 2 protocols to be encrypted using NGE key groups. The following figure shows the packet format for Layer 2 encryption.

Figure 152: Encrypted Layer 2 packet



26178

Similar to router interface encryption, which reuses IPSec transport mode formatted packets, Layer 2 NGE encryption reuses the MACsec Ethertype and the ESP protocol, as defined in RFC 4303, to encrypt the payload of a Layer 2 packet.

Layer 2 NGE is configured on network Ethernet ports; access ports and hybrid ports do not support Layer 2 encryption. When Layer 2 encryption is configured, all IS-IS and LLDP packets are encrypted using NGE. Layer 2 encryption does not encrypt any Layer 3 packets using NGE router interface encryption or MPLS packets using MPLS service encryption.

Layer 2 NGE encryption is enabled by configuring an outbound key group and an inbound key group on the Ethernet port and is similar to assigning key groups for router interface encryption as described in [Assigning key groups to router interfaces](#). For example, to configure Layer 2 encryption on port 1/1/1:

config port 1/1/1 ethernet group-encryption encryption-keygroup 1 direction outbound

config port 1/1/1 ethernet group-encryption encryption-keygroup 1 direction inbound

The **group-encryption** command initializes the port for NGE Layer 2 encryption where both encrypted NGE Layer 2 packets and clear text packets can be received. The **group-encryption** command is used to add the port to an NGE domain and ensures that any received packets with the MACsec Ethertype are decrypted using NGE Layer 2 encryption.

The outbound key group references the key group to use for encrypting the supported Layer 2 packets egressing the router on the Ethernet port. The inbound key group is used to ensure that ingress traffic is using the correct key group. If ingress traffic that is not using the correct key group is detected, the packets are counted as errors.

If Layer 2 NGE is used for a LAG, each Ethernet port member of the LAG group must be configured with NGE. There is no NGE configuration inheritance from the primary port of a LAG group to its member ports. Each port operates based on its own NGE configuration.

For configuration guidelines for inbound and outbound key groups, see [Configuration notes](#).

9.6 NGE packet overhead and MTU considerations

NGE adds overhead packets to services. The following tables show the additional overhead for the worst-case scenario of MPLS services encryption and worst-case scenario of router interface and Layer 2 encryption. Additional overhead depends on which encryption and authentication algorithms are chosen.

Table 199: NGE overhead for MPLS

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.

Table 200: NGE overhead for router interface and Ethernet port NGE

Item	Number of bytes
ESP	24
ICV	32
Padding	17
Total	73

For Layer 3 packets for router interface encryption, and Layer 2 packets for Layer 2 encryption, the total is 73 bytes because the encryption label and control word copy are not required.

The overhead values in [Table 199: NGE overhead for MPLS](#) 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 described in the following table for MPLS-based and GRE-based services. For details on configuring MTU, see the "MTU configuration guidelines" section in the 7705 SAR Interface Configuration Guide.

The calculations in the table 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 201: Accounting for NGE overhead SDP and service MTU – calculation examples

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 <ul style="list-style-type: none"> 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)

Service type	MTU values with and without NGE enabled
	=> 1449 bytes (with NGE enabled) Service MTU (GRE) considerations with NGE enabled <ul style="list-style-type: none"> 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 the MTU value 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.

When an unencrypted Layer 3 packet ingresses the node and routing determines that the egress interface is a router interface NGE-enabled interface, the node calculates whether the packet size will be greater than the MTU of the egress interface after the router interface NGE overhead is added. If the packet cannot be forwarded out from the network interface, an ICMP message is sent back to the sender and the packet is dropped. Users must configure new MTU values to adjust for the overhead associated with NGE.

If an IP exception ACL that matches the ingressing packet exists on the egress interface, the MTU check applied to the ingress packet includes the router interface NGE overhead. This is because the ingress interface cannot determine which IP exceptions are configured on the egress interface, and therefore the worst-case MTU check that includes the router interface NGE overhead is performed.

9.6.1 GRE fragmentation for NGE packets

GRE fragmentation is supported on NGE-encrypted GRE SDP packets on the Ethernet interfaces of the adapter cards and platforms listed in [NGE overview](#).

To determine if an encrypted packet needs to be fragmented, the system compares the total packet size after NGE encryption to the network port MTU. If the encrypted packet size is larger than the MTU, the packet is fragmented. NGE decryption is performed after the packet is fully reassembled.

See [GRE fragmentation](#) for more information.

9.7 1588v2 encryption with NGE

If a router interface is enabled for encryption and Layer 3 1588v2 packets are sent, they are encrypted using NGE. This means that if port timestamping is enabled on a router interface with NGE, the port timestamp is applied to the Layer 3 1588v2 packet via software-based timestamping instead of hardware-based timestamping, and consequently timing accuracy may degrade. The exact level of timing or

synchronization degradation is dependent on many factors, and testing is recommended to measure any impact.

If there is a need to support Layer 3 1588v2 with better accuracy for frequency or better time using port timestamping, an NGE exception ACL is required to keep the Layer 3 1588v2 packets in clear text. The exception ACL must enable UDP packets with destination port 319 to be sent in clear text.

Layer 2 1588v2 packets are always sent and received in clear text. When NGE is enabled on a Layer 2 port, encryption is not applied and existing Layer 2 1588v2 functions are not impacted.

9.8 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 about QoS and QoS policies, see the 7705 SAR Quality of Service Guide. Also, because QoS for NGE is similar to QoS for IPsec, see [QoS for IPsec](#) for more information.

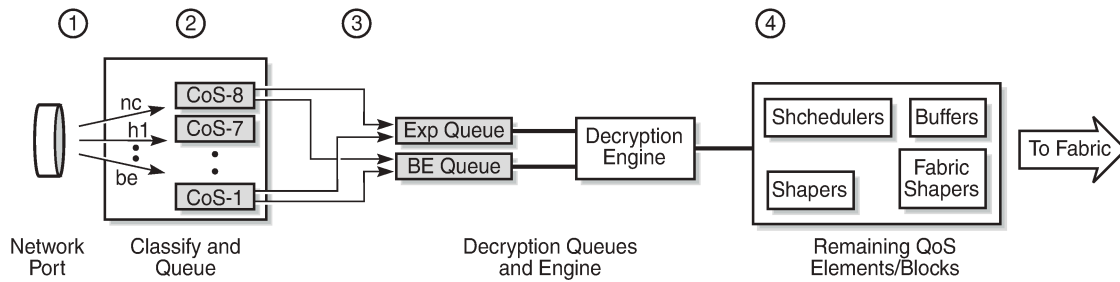
9.8.1 Network ingress

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 the following figure):

- 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 153: QoS for NGE traffic (network ingress)



25102

For more information about QoS for network ingress, see the “Network ingress” section in the 7705 SAR Quality of Service Guide.

9.8.2 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 about QoS for network egress, see the “Network egress” section in the 7705 SAR Quality of Service Guide.

9.9 Statistics

Statistics specific to NGE are counted for the following main areas:

- key group
- SPI
- MDA
- service

9.10 Remote network monitoring 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.

9.11 Configuration notes

This section describes NGE configuration guidelines and restrictions. For more information about configuring NGE using the NSP NFM-P, see the NSP NFM-P User Guide.

To enable NGE for an SDP or VPRN service:

1. Install the outbound direction key group on each node for the service.
2. Install the inbound direction key group on each node for the service.

To enable NGE for a router interface:

1. Enable **group-encryption** on the interface.
2. Configure the outbound key group.
3. Configure the inbound key group.

To change NGE from one key group to another key group for an SDP or VPRN service:

1. Remove the inbound direction key group from 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.

To change NGE from one key group to another key group for a router interface:

1. Remove the inbound key group.
2. Configure the new outbound key group.
3. Configure the new inbound key group.

To disable NGE for an SDP or VPRN service:

1. Remove the inbound direction key group from each node providing the service.
2. Remove the outbound direction key group from each node for the service.

To disable NGE for a router interface:

1. Remove the inbound key group.
2. Remove the outbound key group.
3. Disable **group-encryption** on the interface.

Restrictions:

- 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 NSP NFM-P provides configuration parameters that are not configurable using the CLI. See [Network Services Platform management](#) for more information.

9.12 Configuring NGE with CLI

NGE is fully managed by the NSP NFM-P. The NSP NFM-P ensures correct network synchronization of key groups, services, and NGE domains. Managing NGE without the NSP NFM-P is not recommended. See the NSP NFM-P User Guide for more information.

This section provides information about configuring NGE using the CLI.

Topics in this chapter include:

- [Basic NGE configuration overview](#)
- [Configuring NGE components](#)
- [NGE management tasks](#)

9.13 Basic NGE configuration overview

Perform the following steps to configure NGE for an MPLS service, router interface, or Ethernet port. 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.
 - a. Configure the encryption and authentication algorithms for the group.
 - b. Configure a security association (SA) that contains the encryption and authentication keys.
 - c. Configure the active outbound SA for the group.
3. Select the SDPs, VPRN services, router interfaces, or Ethernet ports that require encryption.
 - a. For each SDP, VPRN service, router interface, or Ethernet port, configure the outbound direction key group.
 - b. For each SDP, VPRN service, router interface, or Ethernet port, configure the inbound direction key group.

9.14 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 VPRN service](#)

- [Assigning a key group to a router interface](#)
- [Assigning a key group to an Ethernet port](#)

9.14.1 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#
```

9.14.2 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-encrypted form, which is indicated by the system-appended **crypto** keyword when an **info** command is run (see the CLI syntax, Example, and CLI output below). The 7705 SAR also includes the **crypto** keyword with an **admin>save** operation so that the 7705 SAR 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
    0x88433A6DB4FA4F8A490EF661CBE69F010BFAE9C2784BED7059E5ADAAB1A225C6
    encryption-key 0x63DCDD501B66F85441E4A55B597DA617
config>grp-encryp>encryp-keygrp# security-
association spi 6 authentication-key
    0x88433A6DB4FA4F8A490EF661CBE69F010BFAE9C2784BED7059E5ADAAB1A225C5
    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
4723a7d1d5f7b6ffa6fc encryption-
key 0x8d51db8f826239f672457442cecc73665f52cbe00aedfb4eda6166001247b4eb crypto
    security-association spi 6 authentication-key 0x7fb9fc5553630924ee29973f
7b0a48f801b0aeb1cb38b7666045274476a268e8d694ab6aa7ea050b7a43cdf8d80977625 encryption-
key 0x72bd9b87841dbebcb2d114031367ab5d9153a41b7c79c8f889ac56b950d8fffa crypto
    active-outbound-sa 6
    exit
-----
ALU-1>config>grp-encryp#

```

9.14.3 Assigning a key group to an SDP or VPRN service

A key group can be assigned to the following entities:

- SDPs
- VPRN services

NGE supports encryption of the following services when key groups are assigned to an SDP or VPRN service:

- VLL services (Epipe and Cpipe)
- VPRN services using Layer 3 spoke-SDP termination

- IES services using Layer 3 spoke-SDP termination
- VPLS services using spoke and mesh SDPs
- routed VPLS services into a VPRN or IES
- MP-BGP-based VPRNs
- NG-MVPN

For services that use SDPs, all tunnels may be either MPLS LSPs (RSVP-TE, LDP, or static LSP) or GRE tunnels. NGE is not supported on IP tunnels.

For VPRNs, the following encryptions are supported:

- unicast VPRN – MP-BGP-based VPRN-level encryption using spoke SDPs (**spoke-sdp**) or autobind SDPs (**auto-bind-tunnel**) with LDP, GRE, RSVP-TE, or segment routing (SR-ISIS, SR-OSPF, or SR-TE) tunnels
- multicast VPRN – NG-MVPN using mLDP with auto-discovery

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
```

```
...
-----
```

9.14.4 Assigning a key group to a router interface

Use the following CLI syntax to assign a key group to a router interface:

CLI syntax:

```
config>router# interface ip-int-name [create]
group-encryption
    encryption-keygroup keygroup-id direction {inbound | outbound}
```

The following example displays a key group assigned to a router interface:

Example:

```
config>router# interface demo
config>router>if# group-encryption
config>router>if>group-encryp# encryption-keygroup 6 direction inbound
config>router>if>group-encryp# encryption-keygroup 6 direction outbound
```

The following example displays key group configuration for a router interface.

```
ALU-1:Sar18>config>router# info
-----
...
    interface demo
        group-encryption
            encryption-keygroup 6 direction inbound
            encryption-keygroup 6 direction outbound
        exit
        no shutdown
        exit
    exit
...
-----
```

9.14.5 Assigning a key group to an Ethernet port

Use the following CLI syntax to assign a key group to an Ethernet port:

CLI syntax:

```
config# port port-id
ethernet
    group-encryption
        encryption-keygroup keygroup-id direction {inbound |
outbound}
```

The following example displays a key group assigned to an Ethernet port:

Example:

```
config# port 1/2/2
config>port# ethernet
config>port>ethernet# group-encryption
config>port>ethernet>group-encryp# encryption-keygroup 6 direction inbound
```

```
config>port>ethernet>group-encryp# encryption-keygroup 6 direction
outbound
```

The following example displays key group configuration for an Ethernet port.

```
ALU-1:Sar18>config>port# info
-----
...
    ethernet
        group-encryption
            encryption-keygroup 6 direction inbound
            encryption-keygroup 6 direction outbound
            exit
        no shutdown
        exit
    exit
...
-----
```

9.15 NGE management tasks

This section discusses the following NGE management tasks:

- [Modifying a key group](#)
- [Removing a key group](#)
- [Changing key groups](#)
- [Deleting a key group from a 7705 SAR](#)

9.15.1 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
0x0123456789012345678901234567890123456789012345678901234567890123
encryption-key
0x0123456789012345678901234567890123456789012345678901234567890123
config>grp-encryp>encryp-keygrp# security-
association spi 6 authentication-key
0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF
encryption-key
0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF
[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 0x0123456789012345678901234567890123456789012345678901234567890123 encryption-
key 0x0123456789012345678901234567890123456789012345678901234567890123
        security-association spi 6 authentication-
key 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF encryption-
key 0x0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF0123456789ABCDEF crypto
        active-outbound-sa 2
    exit
-----
ALU-1>config>grp-encryp#

```

9.15.2 Removing a key group

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.

9.15.2.1 Removing a key group from an SDP or VPRN service

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
```

9.15.2.2 Removing a key group from a router interface

Use the following CLI syntax to remove a key group from a router interface:

CLI syntax:

```
config>router# interface ip-int-name
    group-encryption
        no encryption-keygroup keygroup-id direction {inbound | outbound}
```

The following example displays a key group removed from a router interface:

Example:

```
config>router# interface demo
config>router>if# group-encryption
config>router>if>group-encryp# no encryption-keygroup 6 direction inbound
config>router>if>group-encryp# no encryption-keygroup 6 direction outbound
```

The following example shows that the key group configuration has been removed from a router interface.

```
ALU-1:Sar18>config>router# info
-----
...
    interface demo
        group-encryption
            exit
        no shutdown
        exit
    exit
...
-----
```

9.15.2.3 Removing a key group from an Ethernet port

Use the following CLI syntax to remove a key group from an Ethernet port:

CLI syntax:

```
config# port port-id
    ethernet
        group-encryption
            no encryption-keygroup keygroup-id direction {inbound |
outbound}
```

The following example displays a key group removed from an Ethernet port:

Example:

```
config# port 1/2/2
config>port# ethernet
config>port>ethernet# group-encryption
```

```
config>port>ethernet>group-encryp# no encryption-keygroup 6 direction
inbound
config>port>ethernet>group-encryp# no encryption-keygroup 6 direction
outbound
```

The following example shows that the key group configuration has been removed from an Ethernet port.

```
ALU-1: Sar18>config>port# info
-----
...
    ethernet
        group-encryption
            exit
        no shutdown
        exit
    exit
...
-----
```

9.15.3 Changing key groups

Use the following sequence of CLI commands to change key groups:

1. Remove the inbound direction key group.
2. Change the outbound direction key group.
3. Install the new inbound direction key group.

9.15.3.1 Changing the key group for an SDP or VPRN service

Changing key groups for an SDP or VPRN service must be done on all nodes 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

```

9.15.3.2 Changing the key group for a router interface

The following CLI syntax changes the key group on a router interface. In the example below, the inbound and outbound key groups are changed from key group 6 to key group 8.

CLI syntax:

```

config>router# interface ip-int-name
    group-encryption
        no encryption-keygroup keygroup-id direction {inbound|outbound}

```

Example:

```

config>router# interface demo
config>router>if# group-encryption
config>router>if>group-encryp# no encryption-keygroup 6 direction inbound
config>router>if>group-encryp# encryption-keygroup 8 direction outbound
config>router>if>group-encryp# encryption-keygroup 8 direction inbound

```

The following example shows that the key group configuration has been changed for the router interface.

```

ALU-1:Sar18>config>router# info
-----
...
    interface demo
        group-encryption
            encryption-keygroup 8 direction inbound
            encryption-keygroup 8 direction outbound
        exit
        no shutdown
        exit
    exit
...
-----

```

9.15.3.3 Changing the key group for an Ethernet port

The following CLI syntax changes the key group on an Ethernet port. In the example below, the inbound and outbound key groups are changed from key group 6 to key group 8.

CLI syntax:

```

config# port port-id
    ethernet
        group-encryption

```

```
no encryption-keygroup keygroup-id direction {inbound|
outbound}
```

Example:

```
config# port 1/2/2
config>port# ethernet
config>port>ethernet# group-encryption
config>port>ethernet>group-encryp# no encryption-keygroup 6 direction
inbound
config>port>ethernet>group-encryp# encryption-keygroup 8 direction
outbound
config>port>ethernet>group-encryp# encryption-keygroup 8 direction inbound
```

The following example shows that the key group configuration has been changed for the Ethernet port.

```
ALU-1:Sar18>config>port# info
-----
...
    ethernet
        group-encryption
            encryption-keygroup 8 direction inbound
            encryption-keygroup 8 direction outbound
        exit
        no shutdown
    exit
exit
...
-----
```

9.15.4 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, VPRN services, router interfaces, and Layer 2 NGE-encrypted Ethernet ports 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>grp-encryp# no encryption-keygroup 8
```

9.16 NGE command reference

9.16.1 Command hierarchies

- Configuration commands
 - NGE commands
 - Services commands
 - Router interface encryption commands
 - Ethernet port encryption commands
- Show commands
- Clear commands

9.16.1.1 Configuration commands

9.16.1.1.1 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

```

9.16.1.1.2 Services commands

```

config
- service
- sdp
- encryption-keygroup keygroup-id direction {inbound | outbound}
- no encryption-keygroup direction {inbound | outbound}
- vpn
- encryption-keygroup keygroup-id direction {inbound | outbound}
- no encryption-keygroup direction {inbound | outbound}

```

See [Global service command reference](#) for information about encryption key groups for an SDP and [VPRN services command reference](#) for information about encryption key groups for a VPRN service.

9.16.1.1.3 Router interface encryption commands

```
config
- router
- [no] interface ip-int-name
- [no] group-encryption
- encryption-keygroup keygroup-id direction {inbound | outbound}
- no encryption-keygroup direction {inbound | outbound}
- ip-exception filter-id direction {inbound | outbound}
- no ip-exception direction {inbound | outbound}
```

See the "IP router command reference" section in the 7705 SAR Router Configuration Guide for information about router interface encryption commands.

9.16.1.1.4 Ethernet port encryption commands

```
config
- [no] port port-id
- ethernet
- [no] group-encryption
- encryption-keygroup keygroup-id direction {inbound | outbound}
- no encryption-keygroup direction {inbound | outbound}
```

See the "Configuration command reference" section in the 7705 SAR Interface Configuration Guide for information about Ethernet port encryption commands.

9.16.1.2 Show commands

```
show
- group-encryption
- encryption-keygroup keygroup-id
- encryption-keygroup keygroup-id spi spi
- summary
```

9.16.1.3 Clear commands

```
clear
- group-encryption
- encryption-keygroup keygroup-id
- encryption-keygroup keygroup-id spi spi
```

9.16.2 Command descriptions

- [Configuration commands](#)
- [Show commands](#)
- [Clear commands](#)

9.16.2.1 Configuration commands

- [Generic commands](#)
- [Group encryption commands](#)

9.16.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>grp-encryp>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

9.16.2.1.2 Group encryption commands

group-encryption

Syntax

group-encryption

Context

config

Description

This command enables the context to configure group encryption parameters.

encryption-keygroup

Syntax

encryption-keygroup *keygroup-id* [**create**]

no encryption-keygroup *keygroup-id*

Context

config>grp-encryp

Description

This command is used to create a key group. When 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 or name 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

active-outbound-sa

Syntax

active-outbound-sa *spi*

no active-outbound-sa

Context

config>grp-encryp>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 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

esp-auth-algorithm

Syntax

esp-auth-algorithm {**sha256** | **sha512**}

no esp-auth-algorithm

Context

config>grp-encryp>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-encryp>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 resets 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.

keygroup-name**Syntax**

keygroup-name *keygroup-name*

no keygroup-name

Context

config>grp-encryp>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-encryp>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. When configured, they are never displayed in their original, clear text form. Keys are displayed in a 7705 SAR-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-encrypted form

group-encryption-label

Syntax

group-encryption-label *encryption-label*

no group-encryption-label

Context

config>grp-encryp

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

9.16.2.2 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 202: Encryption key group field descriptions](#) describes the fields.

Output example

```
*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 202: Encryption key group field descriptions

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

Label	Description
Decrypted Pkts	The number of decrypted packets forwarded by the key group
Decrypted Bytes	The number of decrypted bytes forwarded by the key group
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 enqueueing 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 enqueueing 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

Label	Description
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
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 203: Group encryption summary field descriptions](#) describes the fields.

Output example

```
A:ALU-1:Sa18>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 203: Group encryption summary field descriptions

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

9.16.2.3 Clear commands

group-encryption

Syntax
group-encryption

Context
clear

Description
This command accesses the context to clear group encryption parameters.

encryption-keygroup

Syntax

```
encryption-keygroup keygroup-id  
encryption-keygroup keygroup-id spi spi
```

Context

```
clear>grp-encryp
```

Description

This command clears NGE information for a key group.

Parameters

- keygroup-id*
specifies the key group identifier
Values 1 to 127 or *keygroup-name* (up to 64 characters)
- spi*
specifies the SPI ID
Values 1 to 127

10 Ethernet virtual private networks

This chapter provides an overview and configuration information about Ethernet virtual private networks (EVPNs).

Topics in this chapter include:

- [Overview and EVPN applications](#)
- [EVPN for MPLS tunnels](#)
- [General EVPN topics](#)
- [Configuring an EVPN service with CLI](#)
- [EVPN command reference](#)

10.1 Overview and EVPN applications

EVPN is an IETF technology as defined in RFC 7432, *BGP MPLS-Based Ethernet VPN*, that uses a new BGP address family and allows VPLS services to be operated as IP-VPNs, where the MAC addresses and the information to set up the flooding trees are distributed by BGP.

EVPN is designed to fill the gaps in other Layer 2 VPN technologies such as VPLS. The main objective of EVPN is to build Ethernet LAN (E-LAN) services similar to RFC 4364 IP-VPNs while supporting MAC address learning in the control plane (distributed by MP-BGP), efficient multi-destination traffic delivery, and active/active multihoming.



Note: EVPN is not supported on non-Ethernet adapter cards. Traffic that ingresses a non-Ethernet adapter card may not be allowed to egress an EVPN endpoint, and EVPN egress traffic is not allowed on a non-Ethernet adapter card. For information about adapter card generations, see the “Evolution of Ethernet adapter cards, modules, and platforms” section in the 7705 SAR Interface Configuration Guide.

EVPN can be used as the control plane for different data plane encapsulations. The 7705 SAR implementation supports the following data plane encapsulation:


- **EVPN for MPLS tunnels (EVPN-MPLS)**

In EVPN-MPLS, PEs are connected by any type of MPLS tunnel. Typically, EVPN-MPLS is used as an evolutionary step for VPLS services in the WAN, with data center interconnect being one of the main applications.

10.1.1 EVPN for MPLS tunnels in E-LAN services

The following figure shows the use of EVPN for MPLS tunnels on the 7705 SAR. In this example, EVPN is used as the control plane for E-LAN services in the WAN.

Figure 154: EVPN for MPLS in VPLS services



28731

EVPN-MPLS is standardized in RFC 7432 as a Layer 2 VPN technology that can fill the gaps in VPLS for E-LAN services. A significant number of service providers offering E-LAN services today require EVPN for its multihoming capabilities as well as for the optimization that EVPN provides. EVPN supports all-active multihoming and single-active multihoming.

Although VPLS already supports single-active multihoming, EVPN single-active multihoming is considered to be a superior technology due to its mass-withdrawal capabilities to speed up convergence in scaled environments.

EVPN technology provides significant benefits, including:

- superior multihoming capabilities
- IP-VPN-like operation and control for E-LAN services
- reduction and (in some cases) suppression of the BMU (broadcast, multicast, and unknown unicast) traffic in the network
- simple provisioning and management
- new set of tools to control the distribution of MAC addresses and ARP entries in the network

10.1.2 EVPN for MPLS tunnels in E-Line services

The MPLS network used by EVPN for E-LAN services can also be shared by Ethernet line (E-Line) services using EVPN in the control plane. EVPN for E-Line services (EVPN-VPWS, virtual private wire service) is a simplification of the RFC 7432 procedure, and is supported on the 7705 SAR in compliance with IETF *draft-ietf-bess-evpn-vpws*.

10.2 EVPN for MPLS tunnels

This section provides information about EVPN for MPLS tunnels:

- [BGP-EVPN control plane for MPLS tunnels](#)
- [EVPN-VPLS for MPLS tunnels](#)
- [EVPN-VPWS for MPLS tunnels](#)
- [EVPN for MPLS tunnels in r-VPLS services](#)
- [MPLS entropy label](#)
- [Preference-based and non-revertive designated forwarder election](#)

10.2.1 BGP-EVPN control plane for MPLS tunnels

The following table lists the EVPN route types supported on the 7705 SAR and their use in EVPN-MPLS.

Table 204: EVPN route types and usage

EVPN route type	Use
Type 1 – Ethernet auto-discovery route	Mass-withdrawal, Ethernet segment identifier (ESI) labels, aliasing
Type 2 – MAC/IP advertisement route	MAC/IP advertisement, IP advertisement for ARP resolution
Type 3 – inclusive multicast Ethernet tag route	Flooding tree setup (BMU flooding)
Type 4 – Ethernet segment (ES) route	ES discovery and DF election
Type 5 – IP prefix advertisement route	IP routing

If EVPN multihoming is not required, two route types are needed to set up a basic EVPN instance (EVI):

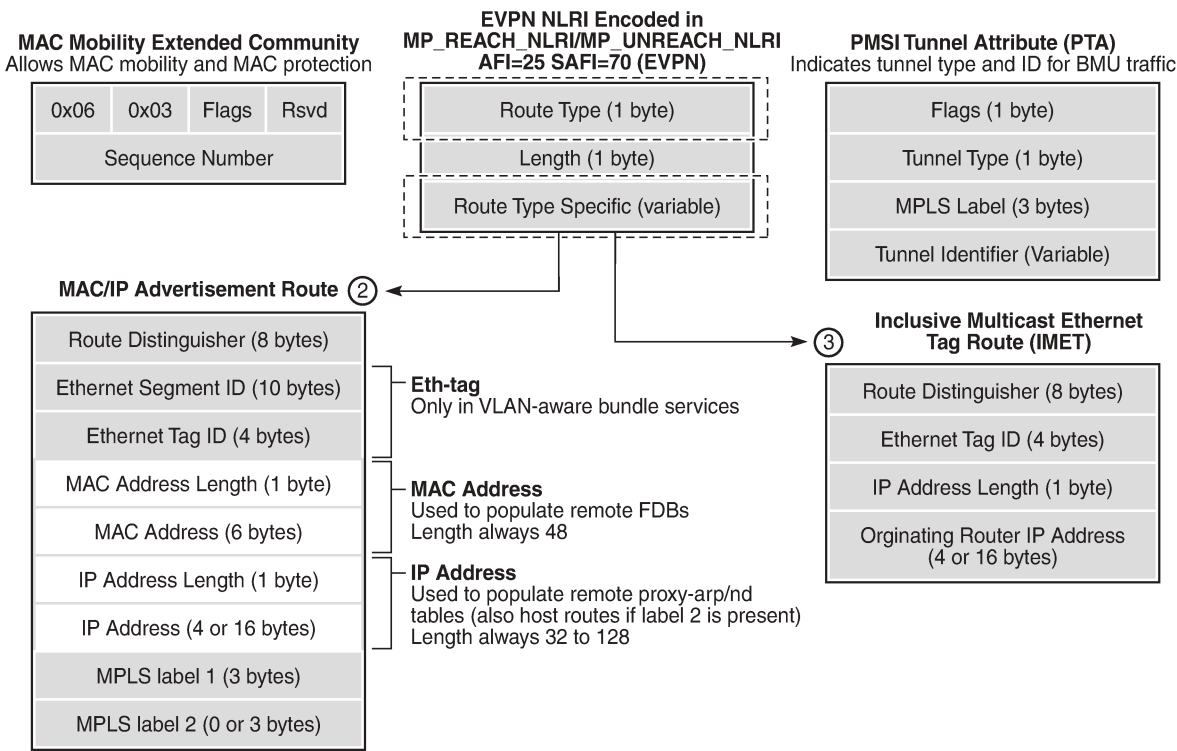
- [EVPN route type 2 – MAC/IP advertisement route](#)
- [EVPN route type 3 – inclusive multicast Ethernet tag route](#)

If multihoming is required, two additional route types are needed:

- [EVPN route type 1 – Ethernet auto-discovery route](#)
- [EVPN route type 4 – Ethernet segment route](#)

The route fields and extended communities for route types 2 and 3 are shown in the following figure.

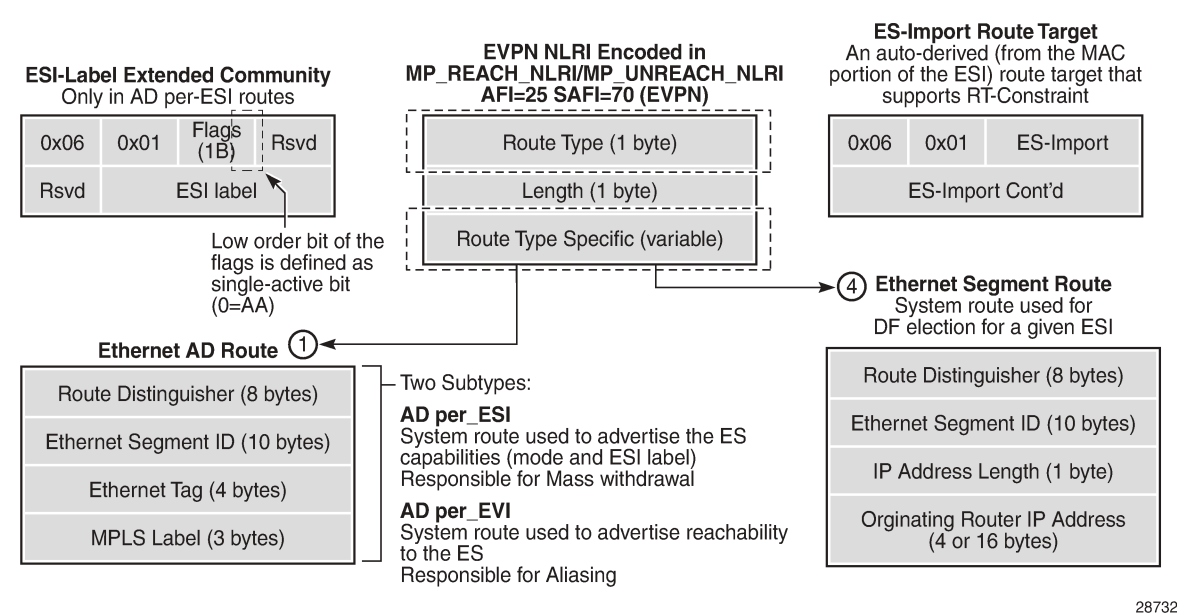
Figure 155: EVPN-MPLS route type 2 and type 3 (required routes and communities)



28734

When EVPN multihoming is enabled in the system, two more routes (route types 1 and 4) are required. The following figure shows the fields in route types 1 and 4 and their associated extended communities.

Figure 156: EVPN route type 1 and type 4



10.2.1.1 EVPN route type 2 – MAC/IP advertisement route

The 7705 SAR generates route type 2 for advertising MAC addresses. If **mac-advertisement** is enabled, the router generates MAC advertisement routes for the following:

- learned MAC addresses on SAPs or SDP bindings
- conditional static MAC addresses



Note: The **unknown-mac-route** command is not supported for EVPN-MPLS services.

Route type 2 uses the fields and values shown in [Figure 155: EVPN-MPLS route type 2 and type 3 \(required routes and communities\)](#) and described in the following table. For type 2 BGP route key processing, the following fields are considered to be part of the prefix in the NLRI: Ethernet tag ID, MAC address length, MAC address, IP address length, and IP address.

Table 205: Route type 2 fields and values

Field	Value
Route distinguisher	Taken from the RD of the VPLS service within the BGP context. The RD can be configured or derived from the bgp-evpn>evi value.
Ethernet segment identifier (ESI)	Zero for MAC addresses learned from single-homed CEs and non-zero for MAC addresses learned from multihomed CEs

Field	Value
Ethernet tag ID	0
MAC address length	Always 48
MAC address	MAC address learned or statically configured
IP address and IP address length	The IP address associated with the MAC address being advertised, with a length of 32 (or 128 for IPv6) In general, any MAC route without an IP address has IPL= 0 (IP length) and the IP address is omitted When received, any IPL value not equal to 0, 32, or 128 discards the route
MPLS label 1	Carries the MPLS label allocated by the system to the VPLS service. The label value is encoded in the high-order 20 bits of the field and is the same label used in route type 3 for the same service unless bgp-evpn>mpls>ingress-replication-bum-label is configured in the service.
MPLS label 2	0
MAC mobility extended community	Used for signaling the sequence number in case of MAC moves and the sticky bit in case of advertising conditional static MAC addresses. If a MAC route is received with a MAC mobility extended community, the sequence number and the sticky bit are considered during route selection.

10.2.1.2 EVPN route type 3 – inclusive multicast Ethernet tag route

Route type 3 is used for setting up the flooding tree (BMU flooding) for a specified VPLS service. The received inclusive multicast routes add entries to the VPLS flood list in the 7705 SAR. Ingress replication is supported as the tunnel type in route type 3 when BGP-EVPN MPLS is enabled.

Route type 3 uses the fields and values shown in [Figure 155: EVPN-MPLS route type 2 and type 3 \(required routes and communities\)](#) and described in the following table. For type 3 BGP route key processing, the following fields are considered to be part of the prefix in the NLRI: Ethernet tag ID, IP address length, and originating router IP address.

Table 206: Route type 3 fields and values

Field	Value
Route distinguisher	Taken from the RD of the VPLS service within the BGP context. The RD can be configured or derived from the bgp-evpn>evi value.
Ethernet tag ID	0
IP address length	Always 32

Field	Value
Originating router IP address	Carries the system address (IPv4 only)
PMSI attribute	<p>The PMSI attribute can have different formats depending on the tunnel type enabled in the service</p> <p>Tunnel type = ingress replication (6)</p> <p>The route is referred to as an Inclusive Multicast Ethernet Tag IR (IMET-IR) route and the PMSI Tunnel Attribute (PTA) fields are populated as follows:</p> <p>Flags – leaf not required</p> <p>MPLS label – carries the MPLS label allocated for the service in the high-order 20 bits of the label field. Unless bgp-evpn>mpls>ingress-replication-bum-label is configured in the service, the MPLS label used is the same as that used in the MAC/IP routes for the service.</p> <p>Tunnel endpoint – equal to the originating IP address</p>

10.2.1.3 EVPN route type 1 – Ethernet auto-discovery route

The 7705 SAR generates route type 1 for advertising for multihoming functions. The system can generate the following two subtypes of Ethernet auto-discovery (AD) routes:

- Ethernet AD per-ESI route (Ethernet segment ID)
- Ethernet AD per-EVI route (EVPN instance)

The Ethernet AD per-ESI route uses the fields and values shown in [Figure 156: EVPN route type 1 and type 4](#) and described in the following table. For type 1 BGP route key processing, the following fields are considered to be part of the prefix in the NLR: Ethernet segment identifier and Ethernet tag ID.

Table 207: Route type 1 fields and values (Ethernet AD per-ESI route)

Field	Value
Route distinguisher	Taken from the system-level RD or service-level RD
Ethernet segment identifier (ESI)	Contains a 10-byte identifier as configured in the system for a specified Ethernet segment
Ethernet tag ID	MAX-ET (0xFFFFFFFF). This value is reserved and used only for AD routes per ESI.
MPLS label	0
ESI label extended community	Includes the single-active bit (0 for all-active and 1 for single-active) and ESI label for all-active multihoming split-horizon
Route target extended community	Taken from the service-level RT or an RT set for the services defined on the Ethernet segment

The system can send either a separate Ethernet AD per-ESI route per service or several Ethernet AD per-ESI routes aggregating the route targets for multiple services. While both alternatives can interoperate, RFC 7432 states that the EVPN AD per-ES route must be sent with a set of route targets corresponding to all the EVIs defined on the Ethernet segment. Either alternative can be enabled using the **ad-per-es-route-target** command options.

The default option, **evi-rt**, configures the system to send a separate AD per-ES route per service.

The **evi-rt-set route-distinguisher ip-address** option, when enabled, allows the aggregation of routes; that is, a single AD per-ES route with the associated RD (*ip-address:1*) and a set of EVI route targets are advertised (to a maximum of 128 route targets). If the number of EVIs defined in the Ethernet segment is significantly large for the packet size, the system will send more than one route. For example:

- AD per-ES route for EVI-route-set 1 is sent with RD *ip-address:1*
- AD per-ES route for EVI-route-set 2 is sent with RD *ip-address:2*

The Ethernet AD per-EVI route uses the fields and values shown in [Figure 156: EVPN route type 1 and type 4](#) and described in the following table.


 **Note:** The AD per-EVI route does not send the ESI label extended community field as is done for Ethernet AD per-ESI (see [Table 207: Route type 1 fields and values \(Ethernet AD per-ESI route\)](#)).

Table 208: Route type 1 fields and values (Ethernet AD per-EVI route)

Field	Value
Route distinguisher	Taken from the service-level RD
Ethernet segment identifier (ESI)	Contains a 10-byte identifier as configured in the system for a specified Ethernet segment
Ethernet tag ID	0
MPLS label	Encodes the unicast label allocated for the service (high-order 20 bits)
Route target extended community	Taken from the service-level RT

10.2.1.4 EVPN route type 4 – Ethernet segment route

The 7705 SAR generates route type 4 for multihoming Ethernet segment (ES) discovery and designated forwarder (DF) election.

The Ethernet segment route uses the fields and values shown in [Figure 156: EVPN route type 1 and type 4](#) and described in the following table. For type 4 BGP route key processing, the following fields are considered to be part of the prefix in the NLRI: Ethernet segment ID, IP address length, and originating router IP address.

Table 209: Route type 4 fields and values

Field	Value
Route distinguisher	Taken from the service-level RD
Ethernet segment identifier (ESI)	Contains a 10-byte identifier as configured in the system for a specified Ethernet segment
ES-import route target community	Automatically derived value from the MAC address portion of the ESI. This extended community is treated as a route target and is supported by RT-constraint (route target BGP family).

10.2.1.5 EVPN route type 5 – IP prefix route


EVPN route type 5 (IP prefix route) is supported for MPLS tunnels. The route fields for route type 5 are shown in the following figure and described in the table. For type 5 BGP route key processing, the following fields are considered to be part of the prefix in the NLRI: Ethernet tag ID, IP prefix length, and IP prefix.

All the routes in EVPN-MPLS are sent with the RFC 5512 tunnel encapsulation extended community, with the tunnel type value set to MPLS.

The router generates route type 5 for advertising IP prefixes in EVPN. The router generates IP prefix advertisement routes for:

- IP prefixes existing in a VPRN linked to the integrated routing and bridging (IRB) backhaul r-VPLS service
- The IRB interface refers to an r-VPLS service bound to a VPRN IP interface.

Figure 157: EVPN route type 5



28912

Table 210: Route type 5 fields and values

Field	Value
Route distinguisher	Taken from the RD configured in the IRB backhaul r-VPLS service within the BGP context
Ethernet segment identifier (ESI)	0:0:0:0:0:0:0:0
Ethernet tag ID	0
IP address length	Any value in the range 0 to 128
IP address	Any valid IPv4 or IPv6 address
GW IP address	Can carry two different values: - If different from 0, route type 5 carries the primary IP interface address of the VPRN behind which the IP prefix is known. This is the case for the regular IRB backhaul r-VPLS model. - If 0.0.0.0, the route type 5 is sent with a MAC next-hop extended community that carries the VPRN interface MAC address. This is the case for the EVPN tunnel r-VPLS model.
MPLS label	Carries the MPLS label allocated for the service Only one MPLS label can be configured per VPLS service

10.2.1.6 RFC 5512 – BGP tunnel encapsulation extended community

The following routes are sent with the RFC 5512 BGP tunnel encapsulation extended community: route type 2 (MAC/IP), route type 3 (inclusive multicast Ethernet tag), and route type 1 (Ethernet AD per-EVI). Route type 4 (Ethernet segment) and route type 1 (AD per-ESI) routes are not sent with this extended community.

The router processes the following BGP tunnel-encapsulation tunnel values registered by IANA for RFC 5512:

- MPLS encapsulation: 10

Any other tunnel value gives the route "treat-as-withdraw" status.

If the encapsulation value is MPLS, the BGP validates the high-order 20 bits of the label field, ignoring the low-order 4 bits.

If the encapsulation extended community is not present in a received route, BGP treats the route as an MPLS-based configuration of the **config>router>bgp>group>neighbor>def-recv-evpn-encap mpls** command. The command is also available at the **bgp** and **group** levels.

10.2.2 EVPN-VPLS for MPLS tunnels

This section provides information about the following topics:

- [Overview](#)
- [EVPN and VPLS integration](#)
- [Auto-derived route distinguisher in services](#)
- [EVPN multihoming in VPLS services](#)

10.2.2.1 Overview

EVPN can be used in MPLS networks where PEs are interconnected through any type of tunnel, including RSVP-TE, segment routing TE, LDP, BGP, segment routing IS-IS, and segment routing OSPF. The selection of the tunnel to be used in a VPLS service with BGP-EVPN MPLS enabled is based on the **auto-bind-tunnel** command, which is similar to the way that VPRN services operate. The BGP-EVPN routes next-hops can be IPv4 or IPv6 addresses and can be resolved to a tunnel in the IPv4 tunnel-table or IPv6 tunnel-table.

EVPN-MPLS is modeled using a VPLS service where EVPN-MPLS bindings can coexist with SAPs and SDP bindings. The following output shows an example of a VPLS service with EVPN-MPLS.

```
*A:PE-1>config>service>vpls# info
-----
description "evpn-mpls-service"
bgp
bgp-evpn
  evi 10
  mpls
    no shutdown
    auto-bind-tunnel resolution any
sap 1/1/1:1 create
exit
spoke-sdp 1:1 create
```

The **bgp-evpn** context must be enabled when MPLS is not **shutdown**. In addition to the **mpls>no shutdown** command, the minimum set of commands needed to set up the EVPN-MPLS instance are the **bgp-evpn>evi** and the **bgp-evpn>mpls>auto-bind-tunnel resolution** commands. Users can also configure other command options. The minimum set and optional commands are listed below:

- **evi value** – the EVPN identifier value (EVI *value*) is unique in the system and can have a value between 1 and 65535. The EVI *value* is used for the service-carving algorithm (which is used for multihoming, if configured) and for auto-deriving route target and route distinguishers in the service.

If the **evi value** is not specified, its value is 0 and no route distinguisher or route targets are auto-derived from it. If the **evi value** is specified and no other route distinguisher or route targets are configured in the service, the following derivations apply:

- the route distinguisher is derived from *system-ip:evi value*
- the route target is derived from *autonomous-system:evi value*



Note: When the **vsi-import** and **vsi-export** policies are configured, the route target must be configured in the policies and the policy values take preference over the values of auto-

derived route targets. The operational route target for a service is displayed using the **show service id service-id bgp** command.

When the **evi** command is configured, a **config>service>vpls>bgp** command is required—even if the context is empty—to allow the user to see the correct information when using the **show>service>id service-id>bgp** and **show>service>system>bgp-route-distinguisher** commands.

Although it is not mandatory, if multihoming is not configured, the configuration of the **evi** command is enforced for EVPN services with SAPs or SDP bindings defined in the **ethernet-segment** command. See [EVPN multihoming in VPLS services](#) for more information about the **ethernet-segment** command.

The following options are specific to EVPN-MPLS and defined in the **bgp-evpn>mpls** context:

- **control-word** – this command is required, as per RFC 7432, to avoid frame disordering. The user can enable and disable it so that interoperability with other vendors can be guaranteed.
- **auto-bind-tunnel** – this command is required and allows the user to decide what type of MPLS transport tunnels are used for a particular instance. The command is used in the same way as it is used in VPRN services.

For BGP-EVPN MPLS, **bgp** is explicitly added to the **resolution-filter** in EVPN (**bgp** is implicit in VPRNs).

- **force-vlan-vc-forwarding** – this command allows the system to preserve the VLAN ID and P-bits of the service-delimiting qtag in a new tag added in the customer frame before sending the customer frame to the EVPN core.
- **split-horizon-group** – this command allows the association of a user-created split horizon group with all the EVPN-MPLS destinations. See [EVPN and VPLS integration](#) for more information.
- **ecmp** – this command, when set to a value greater than 1, activates aliasing to the remote PEs that are defined in the same all-active multihoming Ethernet segment. See [EVPN multihoming in VPLS services](#) for more information.
- **ingress-replication-bum-label** – this command is only enabled when the user wants the PE to advertise a label for BMU traffic (Inclusive Multicast Ethernet Tag routes, route type 3) that is different from the label advertised for unicast traffic (with the MAC/IP routes). This is useful to avoid potential transient packet duplication in all-active multihoming.

In addition to the options above, the following **bgp-evpn** commands are also available for EVPN-MPLS services:

- **[no] mac-advertisement**
- **mac-duplication** and settings
- **route-next-hop**

When EVPN-MPLS is established among some PEs in the network, EVPN unicast and multicast bindings are created on each PE to the remote EVPN destinations. A specified ingress PE creates:

- a unicast EVPN-MPLS destination binding to a remote egress PE as soon as a MAC/IP route is received from that egress PE
- a multicast EVPN-MPLS destination binding to a remote egress PE, only if the egress PE advertises an Inclusive Multicast Ethernet Tag route (route type 3) with a BMU label. This is only possible if the egress PE is configured with **ingress-replication-bum-label**.

The unicast and multicast bindings, as well as the MAC addresses learned on them, can be checked using the **show** commands in the following example. In the example, the remote PE 192.0.2.69 is configured with **no ingress-replication-bum-label** and PE 192.0.2.70 is configured with **ingress-replication-bum-**

label. Therefore, "Dut" has a single EVPN-MPLS destination binding to PE 192.0.2.69 and two bindings (unicast and multicast) to PE 192.0.2.70.

```
*A:Dut# show service id 1 evpn-mpls
```

```
=====
BGP EVPN-MPLS Dest
=====
```

TEP Address	Egr Label Transport	Num. MACs	Mcast	Last Change
192.0.2.69	262118 ldp	1	Yes	06/11/2015 19:59:03
192.0.2.70	262139 ldp	0	Yes	06/11/2015 19:59:03
192.0.2.70	262140 ldp	1	No	06/11/2015 19:59:03
192.0.2.72	262140 ldp	0	Yes	06/11/2015 19:59:03
192.0.2.72	262141 ldp	1	No	06/11/2015 19:59:03
192.0.2.73	262139 ldp	0	Yes	06/11/2015 19:59:03
192.0.2.254	262142 bgp	0	Yes	06/11/2015 19:59:03

Number of entries : 7				

=====				

```
*A:Dut# show service id 1 fdb detail
```

```
=====
Forwarding Database, Service 1
=====
```

ServId	MAC	Source-Identifier	Type Age	Last Change
1	00:ca:fe:ca:fe:69	eMpls: 192.0.2.69:262118	EvpnS	06/11/15 21:53:48
1	00:ca:fe:ca:fe:70	eMpls: 192.0.2.70:262140	EvpnS	06/11/15 19:59:57
1	00:ca:fe:ca:fe:72	eMpls: 192.0.2.72:262141	EvpnS	06/11/15 19:59:57

No. of MAC Entries: 3				

Legend: L=Learned O=Oam P=Protected-MAC C=Conditional S=Static				
=====				

10.2.2.2 EVPN and VPLS integration

The 7705 SAR EVPN implementation supports *draft-ietf-bess-evpn-vpls-seamless-integ* so that EVPN-MPLS and VPLS can be integrated into the same network and within the same service. This feature is useful for the integration between both technologies and for the migration of VPLS services to EVPN-MPLS.

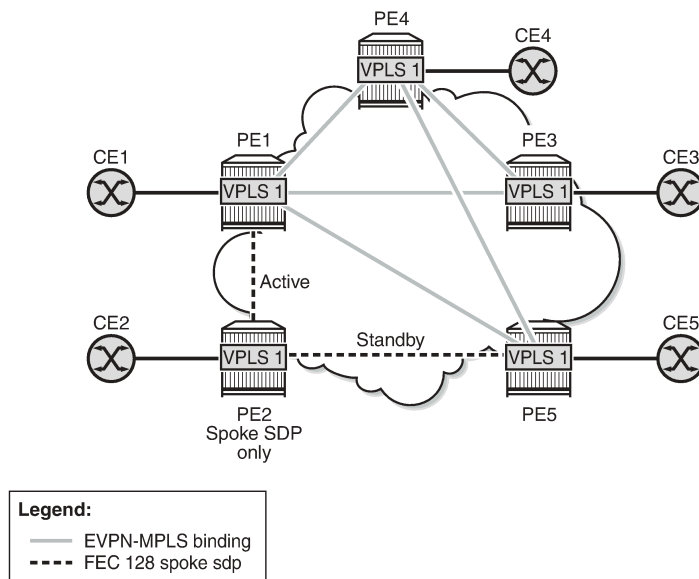
The following behavior enables the integration of EVPN and SDP bindings in the same VPLS network. An illustration and configuration example follow the list:

- Systems with EVPN endpoints and SDP bindings to the same far end bring down the SDP bindings.
 - The router allows the establishment of an EVPN endpoint and an SDP binding to the same far end but the SDP binding is kept operationally down. Only the EVPN endpoint is operationally up. This is true for spoke SDPs and mesh SDPs.
 - If there is an existing EVPN endpoint to a specified far end and the establishment of a spoke SDP is attempted, the spoke SDP is set up but kept operationally down with an operational flag indicating that there is an EVPN route to the same far end.
 - If there is an existing spoke SDP and a valid (used) EVPN route arrives, the EVPN endpoint is set up and the spoke SDP is brought operationally down with an operational flag indicating that there is an EVPN route to the same far end.
 - If there is an SDP binding and EVPN endpoint to different far-end IP addresses on the same remote PE, both links will be operationally up. This can happen if the SDP binding is terminated in an IPv6 address or IPv4 address that is different from the system address where the EVPN endpoint is terminated.
- The user can add spoke SDPs and all the EVPN-MPLS endpoints in the same split horizon group.
 - The **split-horizon-group** *group-name* command under the **vppls>bgp-evpn>mpls** context allows the EVPN-MPLS endpoints to be added to a split horizon group.
 - The **bgp-evpn>mpls>split-horizon-group** must reference a user-configured split horizon group. User-configured split horizon groups can be configured within the **service>vppls** context. The same *group-name* can be associated with SAPs, spoke SDPs, and EVPN-MPLS endpoints.
 - If the **split-horizon-group** command is not used, the default split horizon group—which contains all the EVPN endpoints—is still used, but it is not possible to refer to it on SAPs and spoke SDPs.
- The system disables the advertisement of MAC addresses learned on SAPs and spoke SDPs that are part of an EVPN split horizon group.
 - When the SAPs and spoke SDPs are configured within the same split horizon group as the EVPN endpoints, MAC addresses are still learned on them but are not advertised in EVPN.
 - The SAPs and spoke SDPs added to an EVPN split horizon group should not be part of any EVPN multihomed ES. In this case, the PE still advertises the AD per-EVI route for the SAP or spoke SDP, attracting EVPN traffic that could not be forwarded to that SAP or SDP binding.

The following figure shows an example of EVPN-VPLS integration. In this example, if EVPN and SAPs and spoke SDPs are part of the same split horizon group, the traffic arriving on the SAPs and spoke SDPs is not forwarded to EVPN.

The spoke SDPs on PE2 are not part of an split horizon group, so they can forward traffic to EVPN. Spoke SDPs on PE5 are part of same split horizon group, so they cannot forward traffic to EVPN.

Figure 158: EVPN-VPLS integration



28733

A CLI configuration example for PE1, PE5, and PE2 is provided below.

```
*A:PE1>config>service# info
```

```
-----
vpls 1 customer 1 create
split-horizon-group "SHG-1" create
bgp
  route-target target:65000:1
bgp-evpn
  evi 1
  mpls
    no shutdown
    split-horizon-group SHG-1
spoke-sdp 12:1 create
exit
sap 1/1/1:1 create
exit
```

```
PE5#config>service# info
```

```
-----
vpls 1 customer 1 create
split-horizon-group "SHG-1" create
exit
stp
  shutdown
exit
endpoint "vpls20" create
exit
sap 1/1/7:2 create
  no shutdown
exit
spoke-sdp 64:2 split-horizon-group "SHG-1" endpoint "vpls20" create
  no shutdown
exit
```

```
spoke-sdp 65:2 split-horizon-group "SHG-1" endpoint "vpls20" create
no shutdown
exit
no shutdown
```

```
*A:PE2>config>service# info
-----
vpls 1 customer 1 create
end-point CORE create
no suppress-standby-signaling
spoke-sdp 21:1 end-point CORE
precedence primary
spoke-sdp 25:1 end-point CORE
```

PE1, PE3, PE4, and PE5 have BGP-EVPN enabled in VPLS-1. PE2 has active/standby spoke SDPs to PE1 and PE5. In this configuration:

- EVPN endpoints are instantiated within the same split horizon group, for example, SHG-1
- manual spoke SDPs from PE1 and PE5 to PE2 are not part of SHG-1

EVPN MAC advertisements are as follows:

- MAC addresses learned on FEC128 spoke SDPs are advertised normally in EVPN.
- MAC addresses learned on spoke SDPs that are part of SHG-1 are not advertised in EVPN because SHG-1 is the split horizon group used for **bgp-evpn>mpls**. This prevents any data plane MAC addresses learned on the split horizon group from being advertised in EVPN.

BMU operation on PE1 is as follows:

- When CE1 sends BMU traffic, PE1 floods it to all the active bindings.
- When CE2 sends BMU traffic, PE2 sends it to PE1 (active spoke SDP) and PE1 floods it to all the bindings and SAPs.
- When CE5 sends BMU traffic, PE5 floods it to the EVPN PEs. PE1 will flood the traffic to the active spoke SDP and SAPs but never to the EVPN PEs because they are part of the same split horizon group.

10.2.2.3 Auto-derived route distinguisher in services

In a VPLS service, a single route distinguisher (RD) is used per service.

The following rules apply:

- The VPLS RD is selected based on the following precedence:
 - a manual RD always take precedence when configured
 - if there is no manual configuration, the RD is derived from the **bgp-evpn>evi** configuration
 - if there is no manual or **evi** configuration, there will not be an RD and the service will fail
- The selected RD is displayed in the "Oper Route Dist" field of the **show>service>id>bgp** command.
- The service supports dynamic RD changes, such as a new manual RD configuration.



Note: When the RD changes, the active routes for that VPLS are withdrawn and readvertised with the new RD.

- If the manual mechanism to derive the RD for a specified service is removed from the configuration, the system will select a new RD based on the above rules. In this case, the routes are withdrawn, the new RD is selected from the **evi** configuration, and the routes are readvertised with the new RD.



Note: This reconfiguration will fail if the new RD already exists in a different VPLS or Epipe service.

10.2.2.4 EVPN multihoming in VPLS services

EVPN multihoming implementation is based on the concept of the Ethernet segment and configured through the **ethernet-segment** command. An Ethernet segment is a logical structure that can be defined in one or more PEs and identifies the CE (or access network) that is multihomed to the EVPN PEs. An Ethernet segment is associated with port, LAG, or SDP objects and is shared by all the services defined on those objects. For virtual Ethernet segments, individual VID or VC-ID ranges can be associated with the port, LAG, or SDP objects defined in the Ethernet segment.

Each Ethernet segment has a unique identifier called the Ethernet segment identifier (ESI) that is 10 bytes long and is manually configured in the router.



Note: The *esi* value is advertised in the control plane to all the PEs in an EVPN network. Therefore, it is very important to ensure that the 10-byte *esi* value is unique throughout the entire network. Single-homed CEs are assumed to be connected to an Ethernet segment with *esi* = 0 (that is, single-homed Ethernet segments do not need to be explicitly configured).

This section describes the behavior of the EVPN multihoming implementation in an EVPN-MPLS service and includes the following topics:

- [EVPN all-active multihoming](#)
- [EVPN single-active multihoming](#)

10.2.2.4.1 EVPN all-active multihoming

This section contains information about the following topics:

- all-active multihoming service mode
- ES discovery and DF election procedures (all-active multihoming)
- aliasing
- network failures and convergence for all-active multihoming

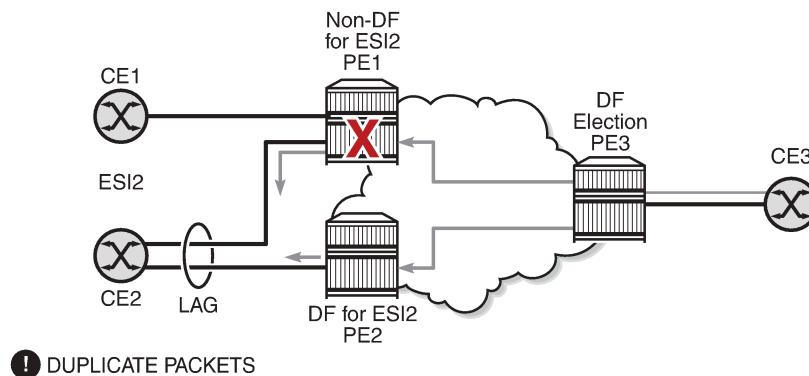
As described in RFC 7432, all-active multihoming is only supported on access LAG SAPs. The CE must be configured with a LAG to avoid duplicated packets to the network. The use of LACP is optional.

Three different procedures are implemented in the 7705 SAR to provide all-active multihoming for a specified Ethernet segment:

- designated forwarder (DF) election (see [Figure 159: DF election](#))
- split horizon (see [Figure 160: Split horizon](#))
- aliasing (see [Figure 161: Aliasing](#))

The following figure shows the need for DF election in all-active multihoming.

Figure 159: DF election



28729


The DF election in an EVPN all-active multihoming scenario avoids duplicate packets on the multihomed CE. The DF election procedure is responsible for electing one DF PE per ESI per service; the other PEs are non-DF for the ESI and service. Only the DF forwards BMU traffic from the EVPN network toward the ES SAPs (the multihomed CE). The non-DF PEs do not forward BMU traffic to the local Ethernet segment SAPs (see [ES discovery and DF election procedures \(all-active multihoming\)](#) for more information).



Note: BMU traffic from the CE to the network and known unicast traffic in any direction is allowed on both the DF and non-DF PEs.

The following figure shows the EVPN split horizon concept for all-active multihoming.

Figure 160: Split horizon




28737

The EVPN split horizon procedure ensures that the BMU traffic originated by the multihomed PE and sent from the non-DF to the DF is not replicated back to the CE (echoed packets on the CE). To avoid these echoed packets, the non-DF (PE1) sends all the BMU packets to the DF (PE2) with an indication of the source Ethernet segment. That indication is the ESI label (ESI2 in the example), previously signaled by PE2 in the AD per-ESI route for the Ethernet segment. When PE2 receives an EVPN packet (after the EVPN label lookup), PE2 finds the ESI label that identifies its local Ethernet segment (ESI2). The BMU packet is replicated to other local CEs but not to the ESI2 SAP.

The following figure shows the EVPN aliasing concept for all-active multihoming.

Figure 161: Aliasing



28725

Because CE2 is multihomed to PE1 and PE2 using an all-active Ethernet segment, aliasing is the procedure by which PE3 can load-balance the known unicast traffic between PE1 and PE2, even if the destination MAC address was only advertised by PE1, as shown in the example. When PE3 installs MAC1 in the FDB, it associates MAC1 not only with the advertising PE (PE1) but also with all the PEs advertising the same **esi esi** (ESI2) for the service. In this example, PE1 and PE2 advertise an AD per-EVI route for ESI2. Therefore, PE3 installs the two next hops associated with MAC1.

Aliasing is enabled by configuring ECMP to be greater than 1 in the **bgp-evpn>mpls** context (see [Aliasing](#) for more information).

10.2.2.4.1.1 All-active multihoming service model

The following shows an example where the PE1 configuration provides all-active multihoming to the CE2 shown in [Figure 161: Aliasing](#).

```
*A:PE1>config>lag(1)# info
-----
mode access
encap-type dot1q
port 1/1/2
lacp active administrative-key 1 system-id 00:00:00:00:00:22
no shutdown

*A:PE1>config>service>system>bgp-evpn# info
-----
route-distinguisher 1.1.1.1:0
ethernet-segment "ESI2" create
esi 01:12:12:12:12:12:12:12:12
multi-homing all-active
service-carving
lag 1
no shutdown

*A:PE1>config>redundancy>evpn-multi-homing# info
-----
boot-timer 120
es-activation-timer 10

*A:PE1>config>service>vpls# info
-----
description "evpn-mpls-service with all-active multihoming"
```

```

bgp
bgp-evpn
  evi 10
  mpls
    no shutdown
    auto-bind-tunnel resolution any
  sap lag-1:1 create
exit

```

In the same way, PE2 is configured as follows:

```

*A:PE2>config>lag(1)# info
-----
mode access
encap-type dot1q
port 1/1/1
lacp active administrative-key 1 system-id 00:00:00:00:00:22
no shutdown

*A:PE2>config>service>system>bgp-evpn# info
-----
route-distinguisher 1.1.1.1:0
ethernet-segment "ESI12" create
  esi 01:12:12:12:12:12:12:12:12
  multi-homing all-active
  service-carving
  lag 1
  no shutdown

*A:PE2>config>redundancy>evpn-multi-homing# info
-----
boot-timer 120
es-activation-timer 10

*A:PE2>config>service>vpls# info
-----
description "evpn-mpls-service with all-active multihoming"
bgp
  route-distinguisher 65001:60
  route-target target:65000:60
bgp-evpn
  evi 10
  mpls
    no shutdown
    auto-bind-tunnel resolution any
  sap lag-1:1 create
exit

```

The preceding configuration enables the all-active multihoming procedures. The following must be considered:

- The Ethernet segment must be configured with a name and a 10-byte **esi esi** value:
 - **config>service>system>bgp-evpn>ethernet-segment *name* create**
 - **config>service>system>bgp-evpn>ethernet-segment>esi esi**
- When configuring the **esi** value, the system enforces the rule that the six high-order octets after the type field are not 0 (so that the auto-derived route target for the ES route is different from 0). Otherwise, the entire **esi** value must be unique in the system.

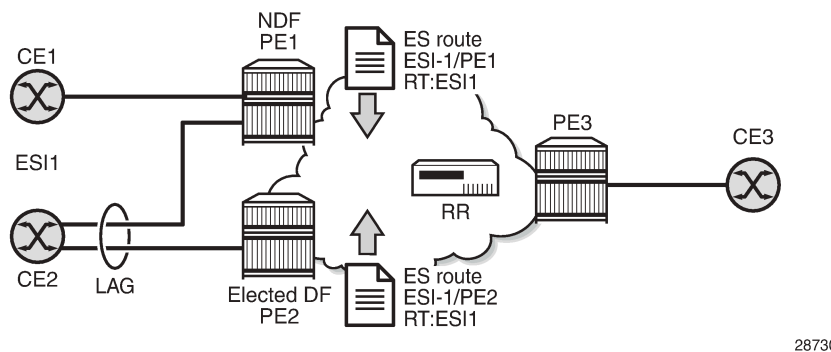
- Only a LAG can be associated with the Ethernet segment. The LAG is used exclusively for EVPN multihoming. Other LAG ports in the system can still be used for MC-LAG and other services.
- When the LAG is configured on PE1 and PE2, the same *admin-key*, *system-priority*, and *system-id* must be configured on both PEs so that CE2 responds as though it is connected to the same system.
- Only one SAP per service can be part of the same Ethernet segment.

10.2.2.4.1.2 ES discovery and DF election procedures (all-active multihoming)

The Ethernet segment (ES) discovery and DF election is implemented in three steps (illustrated in the following figure).

- [Step 1 – ES advertisement and discovery](#)
- [Step 2 – DF election](#)
- [Step 3 – DF and non-DF service behavior](#)

Figure 162: ES discovery and DF election



28730

Step 1 – ES advertisement and discovery

Ethernet segment ESI-1 is configured as described in the previous section ([All-active multihoming service model](#)), with all the required parameters. When **ethernet-segment>no shutdown** is executed, PE1 and PE2 advertise an ES route for ESI-1. Both PEs include the route target auto-derived from the MAC address portion of the configured ESI. If the route target address family is configured in the network, it allows the RR to keep the dissemination of the ES routes under control.

In addition to the ES route, PE1 and PE2 advertise AD per-ESI routes and AD per-EVI routes:

- AD per-ESI routes announce the Ethernet segment capabilities, including the mode (single-active or all-active) and the ESI label for split horizon.
- AD per-EVI routes are advertised so that PE3 knows which services (EVIs) are associated with the ESI. These routes are used by PE3 for its aliasing and backup procedures.

Step 2 – DF election

When the exchange of ES routes between PE1 and PE2 is complete, both PEs run the DF election for all the services in the Ethernet segment.

PE1 and PE2 elect a DF for an ESI-service pair (that is, per ESI, per service). The default DF election mechanism in the 7705 SAR is **service-carving mode auto**, as per RFC 7432. It is possible to manually control the method of selecting a DF by using a preference-based algorithm as described in *draft-rabadan-bess-evpn-pref-df* through the **service-carving manual** CLI context; see [Preference-based and non-revertive designated forwarder election](#) for information.

The following items apply when the default service carving mode is used on a specified PE. The DF election does not occur until the Ethernet segment is configured as **no shutdown**:

- An ordered list of PE IP addresses where ESI-1 resides is built. The IP addresses are obtained from the "Origin IP" fields of all the ES routes received for ESI-1 and also include the local system address. The lowest IP address is considered ordinal "0" in the list.
- The local IP address can only be considered a candidate after a successful **ethernet-segment>no shutdown** command for a specified service.



Note: The remote PE IP addresses must be present in the local PE RTM so that they can participate in the DF election.

- A PE only considers a specified remote IP address as a candidate for the DF election algorithm for a specified service if, in addition to the ES route, the corresponding AD per-ESI and per-EVI routes for that PE have been received and properly activated.
- All remote PEs receiving the AD per-ES routes (for example, PE3) interpret ESI-1 as all-active if all the PEs send their AD per-ES routes with the single-active bit = 0. Otherwise, if at least one PE sends an AD per-ESI route with the single-active flag set or if the local ESI configuration is single-active, the ESI behavior is single-active.
- An **es-activation-timer** can be configured at the **redundancy>bgp-evpn-multi-homing>es-activation-timer** level or at the **service>system>bgp-evpn>ethernet-segment>es-activation-timer** level. This timer, which is 3 seconds by default, delays the transition from non-DF to DF for a specified service after the DF election has run:
 - This use of the **es-activation-timer** with a value different from 0 minimizes the risk of loops and packet duplication due to multiple DFs in transient states.
 - The same **es-activation-timer** value should be configured for all the PEs that are part of the same ESI. The user can configure a long timer to minimize the risk of loops or duplication or a short timer (even **es-activation-timer = 0**) to speed up the convergence for non-DF to DF transitions. When the user configures a specific value, the value configured at the Ethernet segment level supersedes the configured global value.
- The DF election is triggered by the following events:
 - The **config>service>system>bgp-evpn>eth-seg>no shutdown** command triggers the DF election for all the services in the ESI.
 - Reception of a new update or withdrawal of an ES route (containing an ESI configured locally) triggers the DF election for all the services in the ESI.
 - Reception of a new update or withdrawal of an AD per-ES route (containing an ESI configured locally) triggers the DF election for all the services associated with the list of route targets received along with the route.
 - Reception of a new update of an AD per-ES route with a change in the ESI-label extended community (single-active bit or MPLS label) triggers the DF election for all the services associated with the list of route targets received along with the route.

- Reception of a new update or withdrawal of an AD per-EVI route (containing an ESI configured locally) triggers the DF election for that service.
- When the PE boots up, the boot timer allows the necessary time for the control plane protocols to come up before bringing up the Ethernet segment and running the DF algorithm. The boot timer is configured at the system level (**config>redundancy>bgp-evpn-multi-homing>boot-timer**) and should have a value long enough to allow the IOMs and BGP sessions to come up before exchanging ES routes and running the DF election for each EVI.
 - The system does not advertise ES routes until the boot timer expires. This guarantees that the peer ES PEs do not run the DF election until the PE is ready to become the DF, if necessary.
 - The following **show** command displays the configured boot timer value as well as the remaining timer value if the system is still in boot stage.

```
A:PE1# show redundancy bgp-evpn-multi-homing
=====
Redundancy BGP EVPN Multi-homing Information
=====
Boot-Timer           : 10 secs
Boot-Timer Remaining : 0 secs
ES Activation Timer   : 3 secs
=====
```

- When the **service-carving mode auto** command is configured (**auto** is the default mode), the DF election algorithm runs the following function to identify the DF for a specified service and ESI:

V mod N = i (ordinal)

where **V** is the EVI and **N** is the number of peers, for example:

- as shown in [Figure 162: ES discovery and DF election](#), PE1 and PE2 are configured with ESI-1. If $V = 10$ and $N = 2$, then $V \bmod N = 0$, and PE1 would be elected DF—its IP address is lower than PE2's IP address and it is the first PE in the candidate list.



Note: The algorithm uses the configured **evi** value in the service rather than the *service-id*. The **evi** value for a service must match for all the PEs that are part of the ESI. This guarantees that the election algorithm is consistent across all the PEs of the ESI. The **evi** value must always be configured in a service with SAPs or SDP bindings that are created in an ES.

- The user can manually configure the **evi** identifiers for which the PE is primary, using the commands: **service-carving>mode manual** and **service-carving>manual evi start [to to]**
 - The system will be the primary PE, forwarding or multicasting traffic for the **evi** identifiers included in the configuration. The PE will be secondary (non-DF) for the non-specified EVIs.
 - If a range is configured but the **service-carving** mode is not **manual**, the range has no effect.
 - Only two PEs are supported when **service-carving mode manual** is configured. If a third PE is configured with **service-carving mode manual** for an ESI, the two non-primary PEs remain non-DF regardless of the primary status.
 - For example, as shown in [Figure 162: ES discovery and DF election](#), if PE1 is configured with **service-carving manual evi 1 to 100** and PE2 with **service-carving manual evi 101 to 200**, then PE1 will be the primary PE and PE2 will be the secondary PE.

- When **service-carving** is disabled, the lowest originator IP address will win the election for a specified service and ESI. The following command disables service carving:

config>service>system>bgp-evpn>ethernet-segment>service-carving mode off

The following **show** command displays the Ethernet segment configuration and DF status for all EVIs configured in the Ethernet segment.

```
*A:PE1# show service system bgp-evpn ethernet-segment name "ESI-1" all
=====
Service Ethernet Segment
=====
Name                : ESI-1
Admin State         : Up                Oper State         : Up
ESI                 : 01:00:00:00:00:71:00:00:00:01
Multi-homing        : allActive         Oper Multi-homing   : allActive
Source BMac LSB     : 71-71
ES BMac Tbl Size    : 8                  ES BMac Entries     : 1
Lag Id              : 1
ES Activation Timer  : 0 secs
Exp/Imp Route-Target : target:00:00:00:00:71:00

Svc Carving         : auto
ES SHG Label        : 262142
=====

EVI Information
=====
EVI      SvcId      Actv Timer Rem    DF
-----
1         1         0                no
-----
Number of entries: 1
=====

DF Candidate list
-----
EVI      DF Address
-----
1        192.0.2.69
1        192.0.2.72
-----
Number of entries: 2
=====

ISID Information
=====
ISID      SvcId      Actv Timer Rem    DF
-----
20001     20001     0                no
-----
Number of entries: 1
=====

DF Candidate list
-----
ISID      DF Address
-----
20001     192.0.2.69
20001     192.0.2.72
-----
Number of entries: 2
```

```

=====
BMAC Information
=====

```

```

-----
SvcId                      BMacAddress
-----
20000                      00:00:00:00:71:71
-----

```

```

-----
Number of entries: 1
=====

```

Step 3 – DF and non-DF service behavior

Based on the result of the DF election or the manual service carving, the control plane on the non-DF (PE1) instructs the data path to remove the LAG SAP associated with the ESI from the default flooding list for broadcast and multicast (BM) traffic. Unknown unicast traffic may still be sent if the EVI label is a unicast label and the source MAC address is not associated with the ESI.

On PE1 and PE2, both LAG SAPs learn the same MAC address (coming from the CE). For example, in the following **show** commands, 00:ca:ca:ba:ce:03 is learned on both the PE1 and PE2 access LAG (on ESI-1). However, PE1 learns the MAC address as "Learned" whereas PE2 learns it as "Evpn". This is due to CE2 switching the traffic for that source MAC address to PE1. PE2 learns the MAC address through EVPN but it associates the MAC address with the ESI SAP because the MAC address belongs to the ESI.

```
*A:PE1# show service id 1 fdb detail
```

```
=====
Forwarding Database, Service 1
=====
```

ServId	MAC	Source-Identifier	Type Age	Last Change
1	00:ca:ca:ba:ce:03	sap:lag-1:1	L/0	06/11/15 00:14:47
1	00:ca:fe:ca:fe:70	eMpls: 192.0.2.70:262140	EvpnS	06/11/15 00:09:06
1	00:ca:fe:ca:fe:72	eMpls: 192.0.2.72:262141	EvpnS	06/11/15 00:09:39

```
-----
No. of MAC Entries: 3
-----
```

```
Legend: L=Learned O=0am P=Protected-MAC C=Conditional S=Static
=====
```

```
*A:PE2# show service id 1 fdb detail
```

```
=====
Forwarding Database, Service 1
=====
```

ServId	MAC	Source-Identifier	Type Age	Last Change
1	00:ca:ca:ba:ce:03	sap:lag-1:1	Evpn	06/11/15 00:14:47
1	00:ca:fe:ca:fe:69	eMpls: 192.0.2.69:262141	EvpnS	06/11/15 00:09:40
1	00:ca:fe:ca:fe:70	eMpls: 192.0.2.70:262140	EvpnS	06/11/15 00:09:40

```
-----
No. of MAC Entries: 3
-----
```

```
Legend: L=Learned O=0am P=Protected-MAC C=Conditional S=Static
```

When PE1 (non-DF) and PE2 (DF) exchange BMU packets for **evi 1**, all the packets are sent with the ESI label included at the bottom of the stack (in both directions). The ESI label advertised by PE1 and PE2 for ESI-1 can be displayed using the following commands:

```
*A:PE1# show service system bgp-evpn ethernet-segment name "ESI-1"
=====
Service Ethernet Segment
=====
Name                : ESI-1
Admin State         : Up
Oper State          : Up
ESI                 : 01:00:00:00:00:71:00:00:00:01
Multi-homing        : allActive
Oper Multi-homing   : allActive
Source BMac LSB     : 71-71
ES BMac Tbl Size    : 8
ES BMac Entries     : 1
Lag Id              : 1
ES Activation Timer  : 0 secs
Exp/Imp Route-Target : target:00:00:00:00:71:00

Svc Carving         : auto
ES SHG Label        : 262142
=====

*A:PE2# show service system bgp-evpn ethernet-segment name "ESI-1"
=====
Service Ethernet Segment
=====
Name                : ESI-1
Admin State         : Up
Oper State          : Up
ESI                 : 01:00:00:00:00:71:00:00:00:01
Multi-homing        : allActive
Oper Multi-homing   : allActive
Source BMac LSB     : 71-71
ES BMac Tbl Size    : 8
ES BMac Entries     : 0
Lag Id              : 1
ES Activation Timer  : 20 secs
Exp/Imp Route-Target : target:00:00:00:00:71:00

Svc Carving         : auto
ES SHG Label        : 262142
=====
```

10.2.2.4.1.3 Aliasing

Following the example in [Figure 162: ES discovery and DF election](#), if the service configuration on PE3 has ECMP > 1, then PE3 adds PE1 and PE2 to the list of next hops for ESI-1. As soon as PE3 receives a MAC address for ESI-1, it starts load-balancing the flows, per service, between PE1 and PE2 to the remote ESI CE. The following command shows the FDB in PE3.



Note: MAC address 00:ca:ca:ba:ce:03 is associated with the Ethernet segment eES:01:00:00:00:00:71:00:00:00:01 (**esi esi** configured on PE1 and PE2 for ESI-1).

```
*A:PE3# show service id 1 fdb detail
=====
Forwarding Database, Service 1
=====
ServId  MAC                Source-Identifier      Type      Last Change
```

				Age
1	00:ca:ca:ba:ce:03	eES:	Evpn	06/11/15 00:14:47
		01:00:00:00:00:71:00:00:00:01		
1	00:ca:fe:ca:fe:69	eMpls:	EvpnS	06/11/15 00:09:18
		192.0.2.69:262141		
1	00:ca:fe:ca:fe:70	eMpls:	EvpnS	06/11/15 00:09:18
		192.0.2.70:262140		
1	00:ca:fe:ca:fe:72	eMpls:	EvpnS	06/11/15 00:09:39
		192.0.2.72:262141		

No. of MAC Entries: 4				

Legend: L=Learned O=Oam P=Protected-MAC C=Conditional S=Static				
=====				

The following command shows all the EVPN-MPLS destination bindings on PE3, including the ES destination bindings.

The Ethernet segment eES:01:00:00:00:00:71:00:00:00:01 is resolved to PE1 and PE2 addresses:

```
*A:PE3# show service id 1 evpn-mpls
```

=====				
BGP EVPN-MPLS Dest				
=====				
TEP Address	Egr Label Transport	Num. MACs	Mcast	Last Change

192.0.2.69	262140	0	Yes	06/10/2015 14:33:30
	ldp			
192.0.2.69	262141	1	No	06/10/2015 14:33:30
	ldp			
192.0.2.70	262139	0	Yes	06/10/2015 14:33:30
	ldp			
192.0.2.70	262140	1	No	06/10/2015 14:33:30
	ldp			
192.0.2.72	262140	0	Yes	06/10/2015 14:33:30
	ldp			
192.0.2.72	262141	1	No	06/10/2015 14:33:30
	ldp			
192.0.2.73	262139	0	Yes	06/10/2015 14:33:30
	ldp			
192.0.2.254	262142	0	Yes	06/10/2015 14:33:30
	bgp			

Number of entries : 8				

=====				
BGP EVPN-MPLS Ethernet Segment Dest				
=====				
Eth SegId	TEP Address	Egr Label Transport	Last Change	

01:00:00:00:00:71:00:00:00:01	192.0.2.69	262141	06/10/2015 14:33:30	
		ldp		
01:00:00:00:00:71:00:00:00:01	192.0.2.72	262141	06/10/2015 14:33:30	
		ldp		
01:74:13:00:74:13:00:00:74:13	192.0.2.73	262140	06/10/2015 14:33:30	
		ldp		

Number of entries : 3				


PE3 performs aliasing for all the MAC addresses associated with that ESI. This is possible because PE1 is configured with ECMP > 1:

```
*A:PE3>config>service>vpls# info
-----
      bgp
      exit
      bgp-evpn
      evi 1
      exit
      mpls
      ecmp 4
      auto-bind-tunnel
      resolution any
      exit
      no shutdown
    exit
  exit
  proxy-arp
  shutdown
  exit
  stp
  shutdown
  exit
  sap 1/1/1:2 create
  exit
  no shutdown
```

10.2.2.4.1.4 Network failures and convergence for all-active multihoming

The following figure shows the behavior on the remote PE (PE3) when there is an Ethernet segment failure.

Figure 163: All-active multihoming ES failure



28726

The unicast traffic behavior on PE3 is as follows:

1. PE3 can only forward MAC DA = CE2 to both PE1 and PE2 when the MAC advertisement route from PE1 or PE2 and the set of Ethernet AD per-ES routes and Ethernet AD per-EVI routes from PE1 and PE2 are active at PE3.
2. If there is a failure between CE2 and PE2, PE2 withdraws its set of Ethernet AD and ES routes, then PE3 forwards traffic destined for CE2 to PE1 only. PE3 does not need to wait for the withdrawal of the individual MAC address.
3. The same behavior would be followed if the failure had been at PE1.
4. If after (2), PE2 withdraws its MAC advertisement route, PE3 treats traffic to MAC DA = CE2 as unknown unicast, unless the MAC address had been previously advertised by PE1.


For BMU traffic, the following events trigger a DF election on a PE and only the DF forwards BMU traffic after the **es-activation-timer** expiration (if there was a transition from non-DF to DF):

- reception of ES route update (local ES **shutdown/no shutdown** or remote route)
- new AD-ES route update or withdrawal
- new AD-EVI route update or withdrawal
- local ES port or SAP or service shutdown
- service-carving range change (affecting the **evi**)
- multihoming mode change (single-active to all-active or all-active to single-active)

Logical failures on Ethernet segments and black holes

The following figure shows a black hole caused by a SAP or service shutdown.

Figure 164: Black hole caused by SAP or service shutdown



28728

If an individual VPLS service is **shutdown** in PE1 (the example is also valid for PE2), the corresponding LAG SAP goes operationally down. This event triggers the withdrawal of the AD per-EVI route for that particular SAP. PE3 removes PE1 from its list of aliased next hops and PE2 takes over as DF (if it was not the DF already). However, this does not prevent the network from blackholing the traffic that CE2 switches to the link to PE1. Traffic sent from CE2 to PE2 or traffic from the rest of the CEs to CE2 is unaffected, so this situation is not easily detected on the CE.

The same result occurs if the ES SAP is administratively **shutdown** instead of the service.




Note: When `bgp-evpn>mpls shutdown` is executed, the SAP associated with the ES is brought operationally down (**StandbyForMHPProtocol**) and so is the entire service if there are no other SAPs or SDP bindings in the service. However, if there are other SAPs or SDP bindings, the service remains operationally up.

Transient issues caused by MAC route delays

Some situations may cause transient issues to occur, such as a transient packet duplication and a transient black hole. These are shown in the following figure and described below.

Figure 165: Transient issues caused by slow MAC learning



28735

- **Transient packet duplication caused by delay in PE3 to learn MAC1:**

This scenario is illustrated by the diagram at the top of the figure.

In an all-active multihoming scenario, if a specified MAC address (for example, MAC1), is not learned yet in a remote PE (for example, PE3) but it is known in the two PEs of the ES (for example, PE1 and PE2), the latter PEs may send duplicated packets to the CE.

This issue is solved by the use of the **ingress-replication-bum-label** command in PE1 and PE2. If the command is configured, PE1 and PE2 will know that the received packet is an unknown unicast packet; therefore, the non-DF (PE1) will not send the packets to the CE and there will not be duplication.

Even if the **ingress-replication-bum-label** command is not used, this is only a transient situation that is solved as soon as MAC1 is learned in PE3.

- **Transient black hole caused by delay in PE1 to learn MAC1:**

This scenario is illustrated by the diagram at the bottom of the figure.

In an all-active multihoming scenario, MAC1 is known in PE3 and aliasing is applied to MAC1. However, MAC1 is not known yet in PE1, the non-DF for the ES. If PE3 hashing picks up PE1 as the destination for the aliased MAC1, the packets will be blackholed. This case is solved on the non-DF by not blocking

unknown unicast traffic that arrives with a unicast label, which is possible if PE1 and PE2 are configured using **ingress-replication-bum-label**.

As soon as PE1 learns MAC1, the black hole is resolved even if **ingress-replication-bum-label** is not used.

10.2.2.4.2 EVPN single-active multihoming

This section provides information about the following topics:

- single-active multihoming service model
- ES and DF election procedures (single-active multihoming)
- backup PE function
- network failures and convergence for single-active multihoming

The 7705 SAR supports single-active multihoming on access SAPs, LAG SAPs, and spoke SDPs for a specified VPLS service. For LAG SAPs, the CE is configured with a different LAG to each PE in the Ethernet segment (as opposed to a single LAG in all-active multihoming).

The following procedures support EVPN single-active multihoming for a specified Ethernet segment:

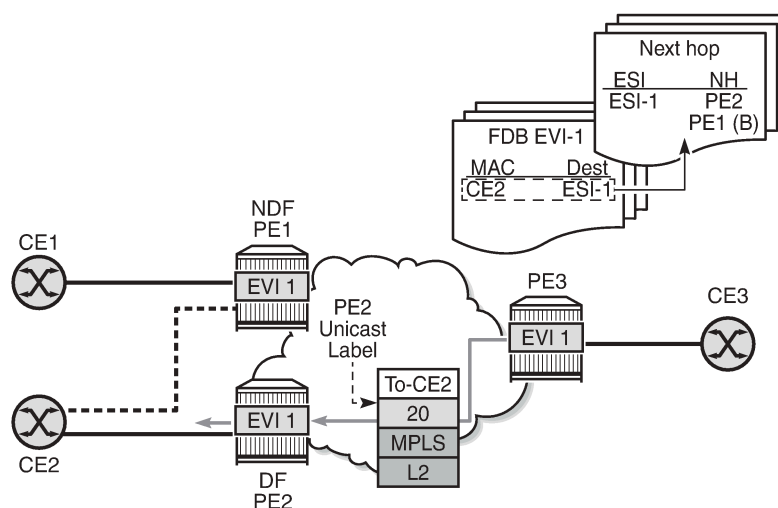
- **DF election**

As in all-active multihoming, DF election in single-active multihoming determines the forwarding for BMU traffic from the EVPN network to the Ethernet segment CE. Also, in single-active multihoming, DF election determines the forwarding of any traffic (unicast and BMU) in any direction (to or from the CE).

- **backup PE**

In single-active multihoming, the remote PEs do not perform aliasing to the PEs in the Ethernet segment. The remote PEs identify the DF based on the MAC routes, send the unicast flows for the Ethernet segment to the PE in the DF, and program a backup PE as an alternative next hop for the remote ESI in case of failure, as shown in the following figure for PE3.

Figure 166: Backup PE



28727

10.2.2.4.2.1 Single-active multihoming service model

The following is an example of PE1 configuration that provides single-active multihoming to CE2, as shown in [Figure 166: Backup PE](#).

```
*A:PE1>config>service>system>bgp-evpn# info
-----
route-distinguisher 1.1.1.1:0
ethernet-segment "ESI2" create
esi 01:12:12:12:12:12:12:12:12
multi-homing single-active
service-carving
sdp 1
no shutdown

*A:PE1>config>redundancy>bgp-evpn-multi-homing# info
-----
boot-timer 120
es-activation-timer 10

*A:PE1>config>service>vpls# info
-----
description "evpn-mpls-service with single-active multihoming"
bgp
bgp-evpn
evi 10
mpls
no shutdown
auto-bind-tunnel resolution any
spoke-sdp 1:1 create
exit
```

The PE2 configuration for this scenario is as follows:

```
*A:PE2>config>service>system>bgp-evpn# info
-----
route-distinguisher 1.1.1.1:0
ethernet-segment "ESI2" create
esi 01:12:12:12:12:12:12:12:12
multi-homing single-active
service-carving
sdp 2
no shutdown

*A:PE2>config>redundancy>bgp-evpn-multi-homing# info
-----
boot-timer 120
es-activation-timer 10

*A:PE2>config>service>vpls# info
-----
description "evpn-mpls-service with single-active multihoming"
bgp
bgp-evpn
evi 10
mpls
no shutdown
auto-bind-tunnel resolution any
spoke-sdp 2:1 create
exit
```

In single-active multihoming, the non-DF PEs for a specified ESI block unicast and BMU traffic in both directions (upstream and downstream) on the object associated with the ESI. Single-active multihoming is similar to all-active multihoming with the following differences:

- The Ethernet segment is configured for single-active: **service>system>bgp-evpn>ethernet-segment>multi-homing single-active**.
- The advertisement of the ESI label in an AD per-ESI route is optional for **single-active** Ethernet segments. By default, the ESI label is used for single-active Ethernet segments. The user can control the **no advertisement** of the ESI label by using the following command: **service>system>bgp-evpn>ethernet-segment>multi-homing single-active no-esi-label**.
- For single-active multihoming, the Ethernet segment can be associated with a port, SDP, or LAG ID, as shown in [Figure 166: Backup PE](#), where:
 - a port is used for single-active SAP redundancy without the need for a LAG
 - an SDP is used for single-active spoke SDP redundancy
 - a LAG ID is used for single-active LAG redundancy



Note: In the last case (LAG ID for single-active LAG redundancy), the *admin-key*, *system-id*, and *system-priority* values must be different on the PEs that are part of the Ethernet segment.

- For single-active multihoming, when the PE is non-DF for the service, the SAPs and spoke SDPs on the Ethernet segment are operationally down and show **StandbyForMHPProtocol** as the reason.
- From a service perspective, single-active multihoming can provide redundancy to CEs (multihomed devices (MHD)) or networks (multihomed networks (MHN)) with the following setup:
 - LAG with or without LACP – in this case, the multihomed ports on the CE are part of different LAGs (one LAG per multihomed PE is used in the CE)
 - regular Ethernet 802.1q/ad ports – in this case, the multihomed ports on the CE or network are not part of any LAG
 - active/standby PWs – in this case, the multihomed CE or network is connected to the PEs through an MPLS network and an active/standby spoke SDP per service. The non-DF PE for each service uses the LDP PW status bits to signal that the spoke SDP is operationally down on the PE side.

10.2.2.4.2.2 ES and DF election procedures (single-active multihoming)

In all-active multihoming, the non-DF keeps the SAP operationally up, although it removes the SAP from the default flooding list. See the [ES discovery and DF election procedures \(all-active multihoming\)](#) for more information. In the single-active multihoming implementation, the non-DF brings the SAP or SDP binding operationally down.

The following **show** commands display the status of the single-active Ethernet segment (ESI-7413) in the non-DF. The associated spoke SDP is operationally down and it signals PW status standby to the multihomed CE.

```
*A:PE1# show service system bgp-evpn ethernet-segment name "ESI-7413"
```

```
=====
Service Ethernet Segment
```

```

=====
Name                : ESI-7413
Admin State         : Up                      Oper State       : Up
ESI                 : 01:74:13:00:74:13:00:00:74:13
Multi-homing        : singleActive           Oper Multi-homing  : singleActive
Source BMAC LSB     : <none>
Sdp Id              : 4
ES Activation Timer  : 0 secs
Exp/Imp Route-Target : target:74:13:00:74:13:00

Svc Carving         : auto
ES SHG Label        : 262141
=====

*A:PE1# show service system bgp-evpn ethernet-segment name "ESI-7413" evi 1
=====
EVI DF and Candidate List
=====
EVI      SvcId      Actv Timer Rem      DF  DF Last Change
-----
1         1          0                  no  06/11/2015 20:05:32
=====

DF Candidates                               Time Added
-----
192.0.2.70                                06/11/2015 20:05:20
192.0.2.73                                06/11/2015 20:05:32
-----
Number of entries: 2
=====

*A:PE1# show service id 1 base
=====
Service Basic Information
=====
Service Id      : 1                      Vpn Id          : 0
Service Type    : VPLS
Name            : (Not Specified)
Description     : (Not Specified)

...<snip>...

-----
Service Access & Destination Points
-----
Identifier                               Type      AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/1:1                             q-tag     9000    9000    Up   Up
sdp:4:13 S(192.0.2.74)                   Spok      0       8978    Up   Down
=====
* indicates that the corresponding row element may have been truncated.

*A:PE1# show service id 1 all | match Pw
Local Pw Bits      : pwFwdingStandby
Peer Pw Bits       : None

*A:PE1# show service id 1 all | match Flag
Flags              : StandbyForMHPProtocol
Flags              : None

```

10.2.2.4.2.3 Backup PE function

A remote PE (PE3 in [Figure 166: Backup PE](#)) imports the AD routes per ESI, where the single-active flag is set. PE3 interprets the Ethernet segment as single-active if at least one PE sends an AD per-ESI route with the single-active flag set. MAC addresses for a specified service and ESI are learned from a single PE, that is, the DF for that ESI-EVI pair (per ESI, per EVI).

The remote PE installs both a single EVPN-MPLS destination (TEP, label) for a received MAC address and a backup next hop to the PE for which the AD per-ESI and per-EVI routes are received. For example, in the following **show** command output, 00:ca:ca:ba:ca:06 is associated with the remote Ethernet segment **eES 01:74:13:00:74:13:00:00:74:13**. That eES resolves to PE 192.0.2.73, which is the DF on the Ethernet segment.

```
*A:PE3# show service id 1 fdb detail
=====
Forwarding Database, Service 1
=====
```

ServId	MAC	Source-Identifier	Type Age	Last Change
1	00:ca:ca:ba:ca:02	sap:1/1/1:2	L/0	06/12/15 00:33:39
1	00:ca:ca:ba:ca:06	eES: 01:74:13:00:74:13:00:00:74:13	Evpn	06/12/15 00:33:39
1	00:ca:fe:ca:fe:69	eMpls: 192.0.2.69:262118	EvpnS	06/11/15 21:53:47
1	00:ca:fe:ca:fe:70	eMpls: 192.0.2.70:262140	EvpnS	06/11/15 19:59:57
1	00:ca:fe:ca:fe:72	eMpls: 192.0.2.72:262141	EvpnS	06/11/15 19:59:57

```
-----
No. of MAC Entries: 5
-----
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
=====

*A:PE3# show service id 1 evpn-mpls
=====
BGP EVPN-MPLS Dest
=====
```

TEP Address	Egr Label Transport	Num. MACs	Mcast	Last Change
192.0.2.69	262118 ldp	1	Yes	06/11/2015 19:59:03
192.0.2.70	262139 ldp	0	Yes	06/11/2015 19:59:03
192.0.2.70	262140 ldp	1	No	06/11/2015 19:59:03
192.0.2.72	262140 ldp	0	Yes	06/11/2015 19:59:03
192.0.2.72	262141 ldp	1	No	06/11/2015 19:59:03
192.0.2.73	262139 ldp	0	Yes	06/11/2015 19:59:03
192.0.2.254	262142 bgp	0	Yes	06/11/2015 19:59:03

```
-----
Number of entries : 7
=====
=====
```


BGP EVPN-MPLS Ethernet Segment Dest			
Eth SegId	TEP Address	Egr Label Transport	Last Change
01:74:13:00:74:13:00:00:74:13	192.0.2.73	262140 ldp	06/11/2015 19:59:03
Number of entries : 1			

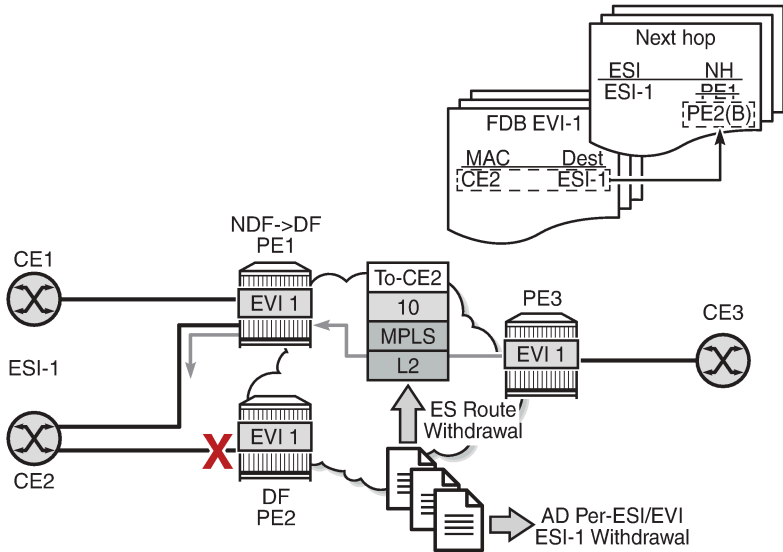
If PE3 sees only two single-active PEs in the same ESI, the second PE is the backup PE. Upon receiving an AD per-ES or per-EVI route withdrawal for the ESI from the primary PE, PE3 starts sending the unicast traffic to the backup PE immediately.

If PE3 receives AD routes for the same ESI and EVI from more than two PEs, the PE does not install any backup route in the data path. Upon receiving an AD per-ES or per-EVI route withdrawal for the ESI, the PE flushes the MAC addresses associated with the ESI.

10.2.2.4.2.4 Network failures and convergence for single-active multihoming

The following figure shows the remote PE (PE3) behavior when there is an Ethernet segment failure.

Figure 167: Single-active multihoming ES failure



28736

- The PE3 behavior for unicast traffic is as follows:
1. PE3 forwards MAC DA = CE2 to PE2 when the MAC advertisement route came from PE2 and the set of Ethernet AD per-ES routes and Ethernet AD per-EVI routes from PE1 and PE2 are active at PE3.
 2. If there is a failure between CE2 and PE2, PE2 withdraws its set of Ethernet AD and ES routes, then PE3 immediately forwards the traffic destined for CE2 to PE1 only (the backup PE). PE3 does not need to wait for the withdrawal of the individual MAC address.

3. After PE2 withdraws its MAC advertisement route, PE3 treats traffic to MAC DA = CE2 as unknown unicast, unless the MAC has been previously advertised by PE1.

Also, a DF election on PE1 is triggered. In general, a DF election is triggered by the same events as for all-active multihoming. In this case, the DF forwards traffic to CE2 when the **es-activation-timer** expiration occurs (the timer activates when there is a transition from non-DF to DF).

10.2.3 EVPN-VPWS for MPLS tunnels

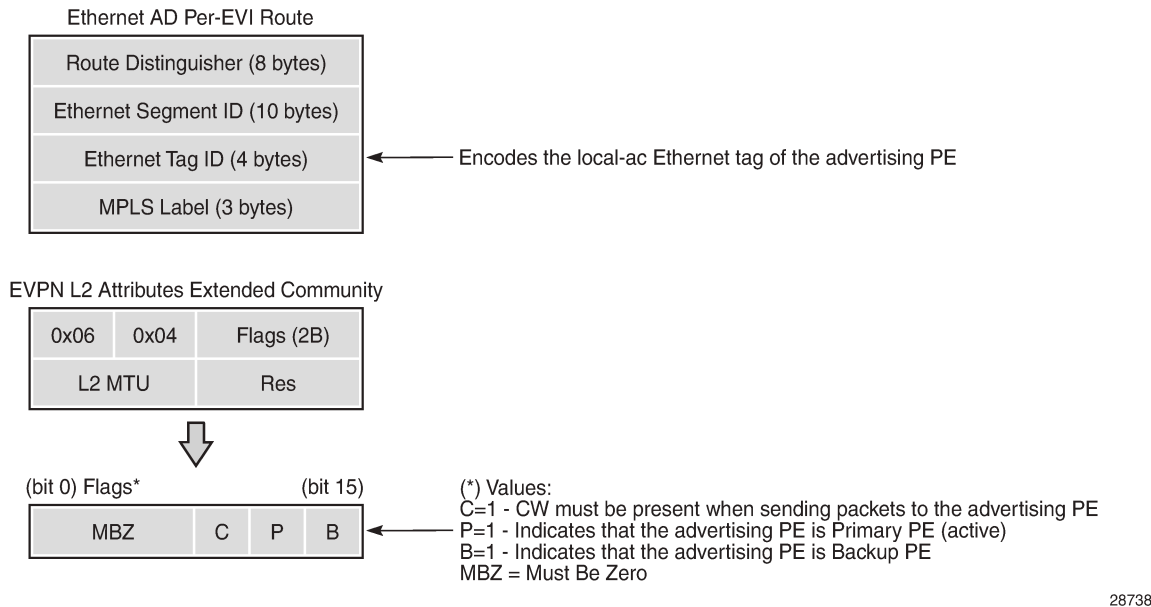
This section contains information about EVPN-VPWS for MPLS tunnels.

- [BGP-EVPN control plane for EVPN-VPWS](#)
- [EVPN for MPLS tunnels in Epipe services](#)
- [Using active/standby PWs and MC-LAG with EVPN-VPWS Epipes](#)
- [EVPN multihoming for EVPN-VPWS services](#)

10.2.3.1 BGP-EVPN control plane for EVPN-VPWS

EVPN-VPWS uses route type 1 and route type 4; it does not use route types 2 or 3. The following figure shows the encoding of the required extensions for the Ethernet AD per-EVI routes. The encoding follows the guidelines described in *draft-ietf-bess-evpn-vpws*.

Figure 168: EVPN-VPWS BGP extensions



If the advertising PE has an access SAP-SDP or spoke SDP that is not part of an Ethernet segment (ES), the PE populates the fields of the AD per-EVI route with the following values:

- Ethernet tag ID field is encoded with the value configured by the user in the **service>epipe>bgp-evpn>local-ac-name>eth-tag value** command

- RD and MPLS label values are encoded as specified in RFC 7432
- ESI is 0
- the route is sent along an EVPN Layer 2 attributes extended community, as specified in IETF *draft-ietf-bess-evpn-vpws*, where:
 - type and subtype are 0x06 and 0x04, as allocated by IANA
 - flag C is set if **control-word** is configured in the service
 - P- and B-flags are zero
 - Layer 2 MTU is encoded with **service-mtu** configured in the Epipe service

If the advertising PE has an access SAP-SDP or spoke SDP that is part of an ES, the AD per-EVI route is sent with the information described in the preceding list, with the following minor differences:

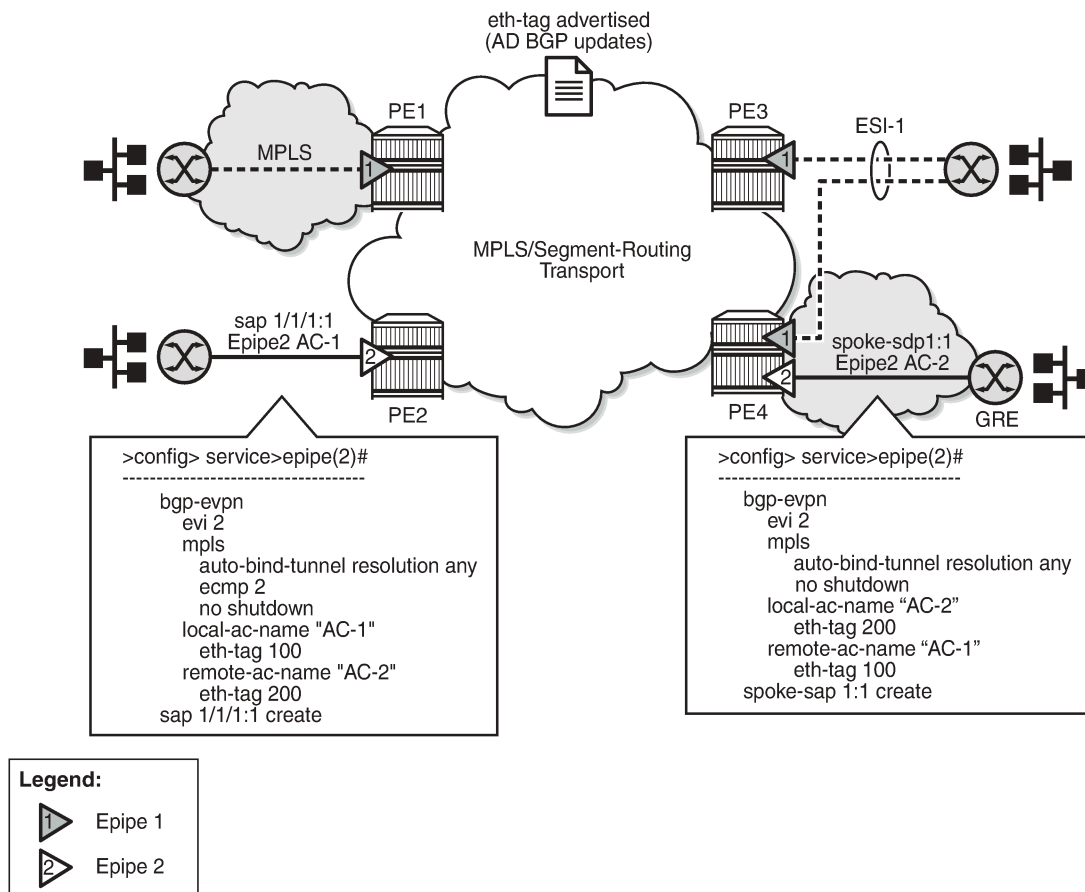
- the ESI encodes the corresponding non-zero value
 - the P- and B-flags are set in the following cases:
 - all-active multihoming
 - all PEs that are part of the ES always set the P-flag
 - the B-flag is never set in the all-active multihoming ES case
 - single-active multihoming
 - only the DF PE sets the P-flag for an EVI and the remaining PEs send it as P=0
 - only the backup DF PE sets the B-flag
- If more than two PEs are present in the same single-active ES, the backup PE is the winner of a second DF election (excluding the DF). The remaining non-DF PEs send B=0.

ES and AD per-ES routes are also advertised and processed for the Ethernet segment, as described in RFC 7432. The ESI label sent with the AD per-ES route is used by BMU traffic on VPLS services; it is not used for Epipe traffic.

10.2.3.2 EVPN for MPLS tunnels in Epipe services

BGP-EVPN can be enabled in Epipe services with either SAPs or spoke SDPs at the access, as shown in the following figure.

Figure 169: EVPN-MPLS in Epipe services



28740

EVPN-VPWS is supported in MPLS networks that also run EVPN-MPLS in VPLS services. From a control plane perspective, EVPN-VPWS is a simplified point-to-point version of RFC 7432 for E-Line services for the following reasons:

- EVPN-VPWS does not use inclusive multicast or MAC/IP routes.
- Ethernet AD per-EVI routes are used to advertise the local attachment circuit identifiers at each side of the VPWS instance. The attachment circuit identifiers are configured as local and remote Ethernet tags. When an AD per-EVI route is imported and the Ethernet tag matches the configured remote Ethernet tag, an EVPN-MPLS destination is created for the Epipe.

In the following configuration example, Epipe 2 is an EVPN-VPWS service between PE2 and PE4 (as shown in the figure):

```
PE2>config>service>epipe(2)#
-----
bgp-evpn
evi 2
mpls
  auto-bind-tunnel resolution any
  ecmp 2
  no shutdown
```

```

    local-ac-name "AC-1"
    eth-tag 100
    remote-ac-name "AC-2"
    eth-tag 200
    sap 1/1/1:1 create

PE4>config>service>epipe(2)#
-----
    bgp-evpn
    evi 2
    mpls
    auto-bind-tunnel resolution any
    no shutdown
    local-ac-name "AC-2"
    eth-tag 200
    remote-ac-name "AC-1"
    eth-tag 100
    spoke-sdp 1:1

```

The following considerations apply to the configuration:

- The **evi** value is used to auto-derive the route target or route distinguisher of the service. The **evi** values must be unique in the system regardless of the type of service they are assigned (Epipe or VPLS).
- Support for the following **bgp-evpn** commands in Epipe services is the same as in VPLS services:
 - **mpls>auto-bind-tunnel**
 - **mpls>control-word**
 - **mpls>entropy-label**
 - **mpls>force-vlan-vc-forwarding**
 - **mpls>route-next-hop**
 - **mpls>shutdown**
- The following **bgp-evpn** commands identify the local and remote attachment circuits, with the configured **eth-tag** values encoded in the advertised and received AD Ethernet per-EVI routes:
 - **local-ac-name** *ac-name*
 - **local-ac-name** *ac-name eth-tag tag-value*, where *tag-value* is 1 to 16777215
 - **remote-ac-name** *ac-name*
 - **remote-ac-name** *ac-name eth-tag tag-value*, where *tag-value* is 1 to 16777215

Changes to the remote **eth-tag** value are allowed without shutting down **bgp-evpn>mpls** or the Epipe service. The local **eth-tag** value cannot be changed without using the **bgp-evpn>mpls>shutdown** command.

Both local and remote **eth-tag** values are mandatory to bring up the Epipe service.

EVPN-VPWS Epipes can also be configured with the following characteristics:

- The BGP-EVPN routes next-hops can be IPv4 or IPv6 addresses and can be resolved to a tunnel in the IPv4 tunnel-table or IPv6 tunnel-table.
- Access attachment circuits can be SAPs or spoke SDPs. Only manually configured spoke SDPs are supported; BGP-VPWS and endpoints are not supported. The **vc-switching** configuration is not supported on **bgp-evpn** enabled Epipes.
- EVPN-VPWS Epipes support **control-word** and **entropy-label**.


When the **bgp-evpn>mpls>control-word** command is configured, the PE sets the C-bit in its AD per-EVI advertisement and sends the control word in the data path. In this case, the PE also expects the control word to be received. If there is a mismatch between the received control word and the configured control word, the system will not set up the EVPN destination. As a result, the service will not come up.

- EVPN-VPWS Epipes can advertise the Layer 2 (service) MTU and check its consistency as follows:
 - The advertised MTU value is taken from the configured **service-mtu** in the Epipe service.
 - The received Layer 2 MTU is checked and compared with the local value. In case of a mismatch between the received MTU and the configured **service-mtu**, the system does not set up the EVPN destination. As a result, the service does not come up.
 - The system does not check the network port MTU value.
 - If the received Layer 2 MTU value is 0, the MTU is ignored.

10.2.3.3 Using active/standby PWs and MC-LAG with EVPN-VPWS Epipes

The use of active/standby PWs (for access spoke SDPs) and MC-LAG (for access SAPs) provides an alternative redundant solution for EVPN-VPWS that does not use the EVPN multihoming procedures described in IETF *draft-ietf-bess-evpn-vpws*. The following figure shows the use of both mechanisms in a single Epipe.

Figure 170: Active/standby PW and MC-LAG support on EVPN-VPWS



28739

In the figure, an active/standby (A/S) PW connects the CE to PE1 and PE2 (left side of the diagram), and an MC-LAG connects the CE to PE3 and PE4 (right side of the diagram). Because EVPN multihoming is not used, there are no AD per-ES routes or ES routes in this example. The redundancy is handled as follows:

- PE1 and PE2 are configured with Epipe 1, where a spoke SDP connects the service in each PE to the access CE. The **local-ac-name eth-tag** is 1 and the **remote-ac-name eth-tag** is 2.
- PE3 and PE4 are configured with Epipe 1, where each PE has a LAG SAP that belongs to a previously configured MC-LAG construct. The **local-ac-name eth-tag** is 2 and the **remote-ac-name eth-tag** is 1.
- An endpoint and A/S PW is configured on the CE on the left side of the diagram. PE1 and PE2 are able to advertise **eth-tag** 1 based on the operational status or the forwarding status of the spoke SDP.

For example, if PE1 receives a standby PW status indication from the CE and the previous status was forward, PE1 withdraws the AD EVI route for **eth-tag** 1. If PE2 receives a forward PW status indication and the previous status was standby or down, PE2 advertises the AD EVI route for **eth-tag** 1.

- The user can configure MC-LAG for access SAPs using the configuration of PE3 and PE4 shown in the figure. In this case, the MC-LAG determines which of the two chassis is active or standby.

If PE4 becomes the standby chassis, the entire LAG port is brought down. As a result, the SAP goes operationally down and PE4 withdraws any previous AD EVI route for **eth-tag** 2.

If PE3 becomes the active chassis, the LAG port becomes operationally up. As a result, the SAP and PE3 advertise the AD per-EVI route for **eth-tag** 2.

10.2.3.4 EVPN multihoming for EVPN-VPWS services

EVPN multihoming is supported for EVPN-VPWS Epipe services with the following considerations:

- Single-active and all-active multihoming is supported for SAPs and spoke SDPs.
- Ethernet segments (ESs) can be shared between the Epipe and VPLS services for LAGs, ports, and SDPs.
- A split-horizon function is not required because there is no traffic between the DF and the non-DF for Epipe services. As a result, the ESI label is never used. For this reason, the **ethernet-segment>multihoming single-active no-esi-label** command does not affect Epipe services.
- The local Ethernet tag values must match on all PEs that are part of the same ES, regardless of the multihoming mode. The PEs in the ES use the AD per-EVI routes from the peer PEs to validate the PEs as DF election candidates for a specific EVI.

The DF election for Epipes that is defined in an all-active multihoming ES is not relevant because all PEs in the ES function in the same way:

- All PEs send P=1 on the AD per-EVI routes.
- All PEs can send upstream and downstream traffic, regardless of whether traffic is unicast, multicast, or broadcast (all traffic is treated as unicast in the Epipe services).

Therefore, the following **tools** command output shows "N/A" when all-active multihoming is configured.

```
*A:PE-2# tools dump service system bgp-evpn ethernet-segment "ESI-12" evi 6000 df
[03/18/2016 20:31:35] All Active VPWS - DF N/A
```

Aliasing is supported for traffic sent to an Ethernet segment destination. If ECMP is enabled on the ingress PE, per-service load balancing is performed to all PEs that advertise P=1. The PEs that advertise P=0 are not considered as next hops for an ES destination.


Although DF election is not relevant for Epipes in an all-active multihoming ES, it is essential for the following forwarding and backup functions in a single-active multihoming ES:

- The PE elected as DF is the primary PE for the Ethernet segment in the Epipe. The primary PE unblocks the SAP or spoke SDP for upstream and downstream traffic. The remaining PEs in the ES bring their ES SAPs or spoke SDPs operationally down.
- The DF candidate list is built from the PEs sending ES routes for the same ES and is pruned for a specific service, depending on the availability of the AD per-ES and per-EVI routes.
- When the SAP or spoke SDPs (part of the ES) come up, the AD per-EVI routes are sent with P=B=0. The remote PEs do not send traffic at this stage.

The remote PEs do not start sending traffic until the DF election process is completed, the **es-activation-timer** has expired, and the PEs advertise AD per-EVI routes with P- and B-bits different from 0.

- The backup PE function is supported as defined in IETF *draft-ietf-bess-evpn-vpws*. The primary PE, backup, or none status is signaled by the PEs (part of the same single-active multihoming ES) in the P- or B-flags of the EVPN Layer 2 attributes extended community. The following figure shows the advertisement and use of the primary, backup, or none indication by the PEs in the ES.

Figure 171: EVPN-VPWS single-active multihoming



28741

As specified in RFC 7432, the remote PEs in VPLS services have knowledge of the primary PE in the remote single-active ES, based on the advertisement of the MAC/IP routes (because only the DF learns and advertises MAC/IP routes).

Because there are no MAC or IP routes in EVPN-VPWS, the remote PEs can forward the traffic based on the P- and B-bits. The process is described in the following list (using the figure as an example):

- The DF PE for an EVI (PE2) sends P=1 and B=0.
- For each ES or EVI, a second DF election is run among the PEs in the backup candidate list to elect the backup PE. The backup PE (PE3) sends P=0 and B=1.

- All remaining multihoming PEs (PE4 and PE5) send P=B=0.
- At the remote PEs (PE1), the P- and B-flags are used to identify the primary and backup PEs within the ES destination. The traffic is then sent to the primary PE, provided that it is active:
 - When a remote PE receives the withdrawal of an Ethernet AD per-ES (or EVI) route from the primary PE, the remote PE immediately switches the traffic to the backup PE for the affected EVIs.
 - The backup PE takes over immediately without waiting for the ES activation timer to expire and bring up its SAP or spoke SDP.
 - The **bgp-evpn mpls ecmp** setting also governs the forwarding in single-active multihoming, regardless of the single-active multihoming bit in the AD per-ES route received at the remote PE (PE1):
 - PE1 always sends the traffic to the primary remote PE (the owner of the P=1 bit). If there are multiple primary PEs and ECMP > 1, PE1 will load-balance the traffic to all the primary PEs, regardless of the multihoming mode.
 - If the last primary PE withdraws its AD per-EVI or ES route, PE1 sends the traffic to the backup PE or PEs. If there are multiple backup PEs and ECMP > 1, PE1 load-balances the traffic to the backup PEs.

10.2.4 EVPN for MPLS tunnels in r-VPLS services

EVPN-MPLS and IP prefix advertisement (enabled by the **ip-route-advertisement** command) are fully supported in routed VPLS services. The following capabilities are supported in a service where **bgp-evpn mpls** is enabled:

- r-VPLS with VRRP support on VPRN interfaces
- r-VPLS support including **ip-route-advertisement** with regular interfaces

This includes the advertisement and process of IP prefix routes defined in IETF *draft-ietf-bess-evpn-prefix-advertisement* with the appropriate encoding for EVPN-MPLS.
- r-VPLS support including **ip-route-advertisement** with **evpn-tunnel** interfaces
- r-VPLS with IPv6 support on the VPRN IP interface

10.2.4.1 Overview

EVPN and MPLS can be enabled on VPLS or r-VPLS services on the 7705 SAR. While the [EVPN-VPLS for MPLS tunnels](#) section focuses on the use of EVPN-MPLS Layer 2 services (that is, how EVPN-MPLS is configured in VPLS services), this section describes how EVPN-MPLS can be used to provide inter-subnet forwarding in r-VPLS and VPRN services. Although inter-subnet forwarding can be provided by regular r-VPLS and VPRN services, EVPN provides an efficient and unified way to populate forwarding databases (FDBs), address resolution protocol (ARP) tables, and routing tables using a single BGP address family. Inter-subnet forwarding in overlay networks would otherwise require data plane learning and the use of routing protocols on a per-VPRN basis.

The 7705 SAR solution for inter-subnet forwarding using EVPN is based on building blocks described in *draft-sajassi-l2vpn-evpn-inter-subnet-forwarding* and the use of the EVPN IP prefix routes (route type 5) as described in *draft-rabadan-l2vpn-evpn-prefix-advertisement*.

The IRB interface refers to an r-VPLS service bound to a VPRN IP interface.

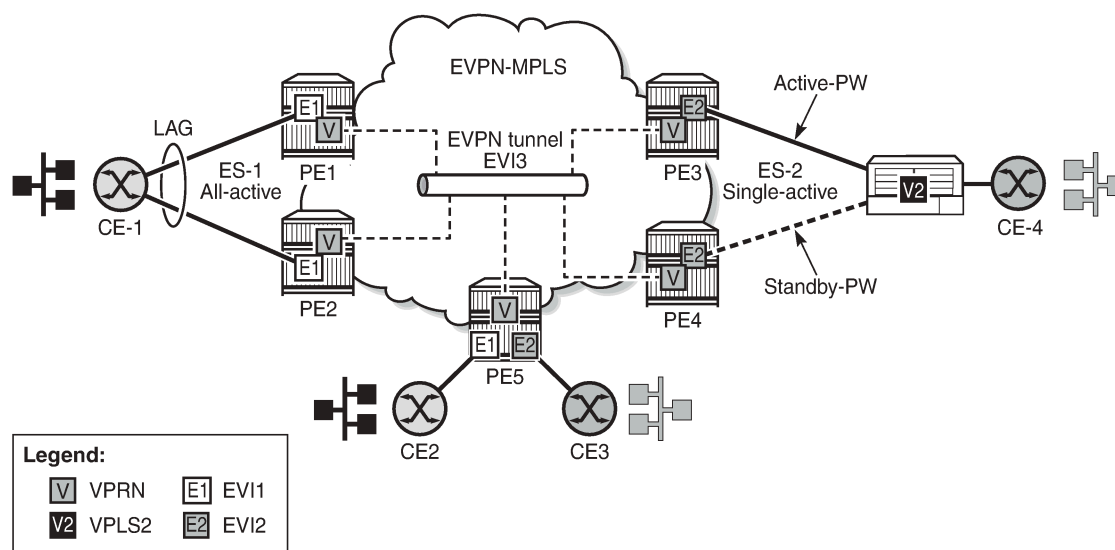
10.2.4.2 EVPN-MPLS multihoming and passive VRRP

SAP-based and spoke SDP-based Ethernet segments are supported on r-VPLS services where **bgp-evpn mpls** is enabled.

The following figure shows an example of EVPN-MPLS multihoming in r-VPLS services, with the following assumptions:

- There are two subnets for a specific customer (for example, EVI1 and EVI2 in the figure), and a VPRN is instantiated in all the PEs for efficient inter-subnet forwarding.
- A backhaul r-VPLS with **evpn-tunnel** mode enabled is used in the core to interconnect all the VPRNs. EVPN IP prefix routes are used to exchange the prefixes corresponding to the two subnets.
- An all-active ES is configured for EVI1 on PE1 and PE2.
- A single-active ES is configured for EVI2 on PE3 and PE4.

Figure 172: EVPN-MPLS multihoming in r-VPLS services



28913

In the figure, the hosts connected to CE1 and CE4 can use regular VRRP for default gateway redundancy; however, this may not be the most efficient way to provide upstream routing.

For example, if PE1 and PE2 are using regular VRRP, the upstream traffic from CE1 may be hashed to the backup IRB VRRP interface instead of being hashed to the master interface. The same thing may occur for single-active multihoming and regular VRRP for PE3 and PE4. The traffic from CE4 will be sent to PE3 although PE4 may be the VRRP master. In that case, PE3 will have to send the traffic to PE4 instead of routing it directly.

In both cases, unnecessary bandwidth between the PEs is used to get to the master IRB interface. In addition, VRRP scaling is limited if aggressive keepalive timers are used.

Because of these issues, passive VRRP is recommended as the best method when EVPN-MPLS multihoming is used in combination with r-VPLS redundant interfaces.

Passive VRRP is a VRRP setting in which the transmission and reception of keepalive messages is completely suppressed, and therefore the VPRN interface always functions as the master. Passive VRRP is enabled by adding the **passive** keyword to the VRRP instance at creation time (passive mode cannot be enabled or disabled while the protocol is running), as shown in the following examples:

```
config service vprn interface vrrp virtual-router-id passive
```

```
config service vprn interface ipv6 vrrp virtual-router-id passive
```

For example, if PE1, PE2, and PE5 in the figure use passive VRRP, then even if each individual r-VPLS interface has a different MAC/IP address, because they share the same VRRP instance 1 and the same backup IP, the three PEs will own the same virtual MAC and virtual IP address (for example, 00-00-5E-00-00-01 and 10.0.0.254). The virtual MAC is auto-derived from 00-00-5E-00-00-VRID, as per RFC 3768. The following is the expected behavior when passive VRRP is used in this example:

- All r-VPLS IRB interfaces for EVI1 have their own physical MAC/IP address; they also own the same default gateway virtual MAC and IP address.
- All EVI1 hosts have a unique configured default gateway; for example, 10.0.0.254.
- When CE1 or CE2 sends upstream traffic to a remote subnet, the packets are routed by the closest PE because the virtual MAC is always local to the PE.

For example, the packets from CE1 hashed to PE1 are routed at PE1. The packets from CE1 hashed to PE2 are routed directly at PE2.

- Downstream packets (for example, packets from CE3 to CE1), are routed directly by the PE to CE1, regardless of the PE to which PE5 routed the packets.

For example, the packets from CE3 sent to PE1 are routed to CE1 at PE1. The packets from CE3 sent to PE2 are routed to CE1 at PE2.

- In case of ES failure in one of the PEs, the traffic is forwarded by the available PE.

For example, if the packets routed by PE5 arrive at PE1 and the link to CE1 is down, PE1 will send the packets to PE2. PE2 will forward the packets to CE1 even if the MAC source address of the packets matches the virtual MAC address of PE2. Virtual MAC addresses bypass the r-VPLS interface MAC protection.

The following list summarizes the advantages of using passive VRRP mode versus regular VRRP for EVPN-MPLS multihoming in r-VPLS services:

- Passive VRRP does not require multiple VRRP instances to achieve default gateway load balancing. Only one instance per r-VPLS—therefore only one default gateway—is needed for all the hosts.
- The convergence time for link or node failures is not impacted by the VRRP convergence because all the nodes in the VRRP instance are master routers.
- Passive VRRP scales better than VRRP because it does not use keepalive or BFD messages to detect failures and allow the backup to take over.

10.2.5 MPLS entropy label

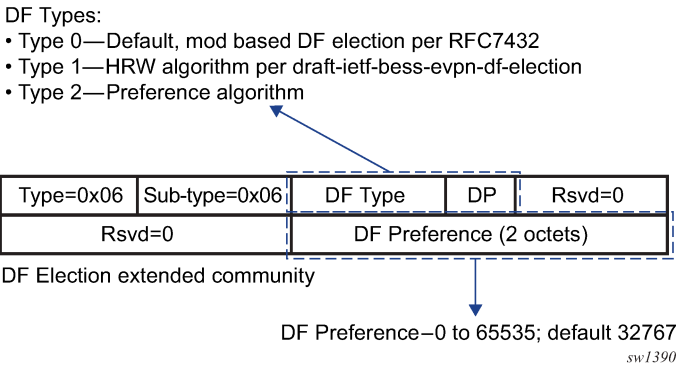
The router supports the MPLS entropy label (RFC 6790), allowing LSR nodes in a network to load-balance labeled packets in a much more granular fashion than allowed by simply hashing on the standard label stack. The entropy label can be enabled on BGP-EVPN services (VPLS and Epipe). See the "MPLS entropy labels" section in the 7705 SAR MPLS Guide for more information.

10.2.6 Preference-based and non-revertive designated forwarder election

The 7705 SAR supports the following service carving modes to elect a PE as the designated forwarder (DF) for a specific ES and service: **auto**, **off**, and **manual**. The **manual** mode supports the preference-based algorithm, which provides a user-controlled method for selecting a DF as described in *draft-rabadan-bess-evpn-pref-df*.

When an ES is configured to use the preference-based algorithm, the ES route is advertised with the DF election extended community attribute (sub-type 0x06). The following figure shows the DF election extended community.

Figure 173: DF election extended community attribute



In the extended community, a DF Type 2 preference algorithm is advertised with a 2-byte preference value (32767) when the **service-carving>manual>preference create** command is configured, triggering preference-based DF election procedures. If any of the PEs exchanging the DF election extended community attribute for a particular ES and EVI do not have a DF Type of 2, all PEs will revert to using the **service-carving mode auto** election procedure. If there are more than four PEs involved in the preference-based DF election for a particular ES and EVI, the system prunes the ES candidate list to the four PEs with the lowest IP addresses.

The configured preference value is used to determine the DF PE when the preference-based DF election algorithm is run for each service EVI. The preferred DF is the PE with the highest or lowest preference value, depending on whether the highest-preference algorithm or the lowest-preference algorithm is used. By default, the highest-preference algorithm is used to elect a DF. If an EVI value or range of values is explicitly configured, the lowest-preference algorithm is used to elect a DF for those EVIs. When there are equal preference values, the DP bit is the first tiebreaker (a DP of 1 wins over a DP of 0) regardless of whether the highest-preference algorithm or lowest-preference algorithm is in use. The lowest PE IP address is the last-resort tiebreaker. If the preference value is not explicitly configured, the default value is used as the DF preference when the DF election extended community attribute is advertised.

The Don't Preempt (DP) bit is set when the **non-revertive** option is enabled. Setting the **non-revertive** option ensures that when a former DF PE comes back up after a failure, it does not take over from an existing active DF PE. By default, the **non-revertive** option is not set and when a former DF PE comes back after a failure, it will take over from an active DF PE.

The following CLI excerpt shows the commands to enable preference-based DF election on a specific ES:

```
config>service>system>bgp-evpn>ethernet-segment#
...
service-carving mode {manual | auto | off}
```

```

service-carving manual
[no] preference [create] [non-revertive]
      value value
      exit
[no] evi evi [to evi]
# value 0..65535; default 32767
...

```

The **service-carving** command must be configured as **manual** to enter the **preference** context.

The **preference** command is supported on ESs regardless of the multihoming mode (single-active or all-active) or the service type (VPLS or Epipe). The preference value can be changed on an active ES without first shutting down the ES, making it possible to force a new DF for maintenance or other reasons.

The Don't Preempt (DP) bit in the DF election extended community attribute is set when the **non-revertive** option is enabled. Setting the **non-revertive** option ensures that when a former DF PE comes back up after a failure, it does not take over from an existing active DF PE. By default, the **non-revertive** option is not set and when a former DF PE comes back after a failure, it will take over from an active DF PE.

If the **non-revertive** option is configured, when the former DF comes back up after a failure and checks existing ES routes, it will advertise an operational preference and DP bit, which does not cause a DF switchover for the ES EVI values.

The following configuration example shows the use of the preference-based algorithm and non-revertive option in an ES defined in PE1 and PE2.

```

*A:PE-1>config>service>system>bgp-evpn# info
-----
ethernet-segment "ES1" create
esi 01:00:00:00:00:12:00:00:00:01
service-carving manual
  preference non-revertive create
    value 10000
  exit
  evi 2001 to 4000
exit
multi-homing single-active
port 1/1/1
no shutdown

/* example of vpls 1 - similar config exists for evi 2-4000 */
*A:PE-1>config>service>vpls# info
-----
vxlan vni 1 create
exit
bgp-evpn
  evi 1
  mpls
    ecmp 2
    auto-bind-tunnel
    resolution any
  exit
sap 1/1/1:1 create
no shutdown
-----
*A:PE-2>config>service>system>bgp-evpn# info
-----
ethernet-segment "ES1" create
esi 01:00:00:00:00:12:00:00:00:01
service-carving manual
  preference non-revertive create
    value 5000

```

```

        exit
        evi 2001 to 4000
    exit
    multi-homing single-active
    port 1/1/1
    no shutdown

*A:PE-2>config>service>vpls# info
-----
vxlan vni 1 create
exit
bgp-evpn
    evi 1
    mpls
        ecmp 2
        auto-bind-tunnel
        resolution any
    exit
    sap 1/1/1:1 create
    no shutdown
-----

```

Based on the configuration in the preceding example, the PE behaves as follows:

- Assuming the ES is **no shutdown** on both PE1 and PE2, the PEs exchange ES routes, including the [Pref, DP-bit] in the DF election extended community attribute.
- For EVIs 1 to 2000, PE2 is immediately promoted to non-designated forwarder (NDF); PE1 becomes the DF and when the **es-activation-timer** expires, PE1 brings up its SAP in EVIs 1 to 2000. For EVIs 2001 to 4000, the result is the opposite and PE2 becomes the DF.
- If port 1/1/1 on PE1 goes down, PE1 withdraws its ES route and PE2 becomes the DF for EVIs 1 to 2000.
- When port 1/1/1 on PE1 comes back up, PE1 compares its ES1 preference with the preferences in the remote PEs in ES1. PE1 advertises the ES route with an "in-use operational" preference of 5000 and DP of 0. Because PE2's preference is the same as PE1's operational value but PE2's DP is 1, PE2 continues to be the DF for EVIs 1 to 4000.



Note: The DP bit is the tiebreaker when there are equal preference values, regardless of the choice of highest-preference or lowest-preference algorithm. The lowest PE-IP is the last-resort tiebreaker.

- PE1's "in-use" preference and DP will continue to be [5000,0] until one of the following conditions is true:
 - PE2 withdraws its ES route, in which case PE1 will readvertise its administrative preference and DP [10000,DP=1]
 - The user changes PE1's preference configuration

10.2.7 IPv6 tunnel resolution for EVPN MPLS Services

EVPN MPLS services can be deployed in a pure IPv6 network infrastructure, where IPv6 addresses are used as next-hops of the advertised EVPN routes, and EVPN routes received with IPv6 next-hops are resolved to tunnels in the IPv6 tunnel-table.

By default, the system IPv4 address is advertised as the next-hop for a local EVPN MPLS service. To configure a system or non-system IPv6 address, use the **route-next-hop** command in the **config>service>epipe>bgp-evpn>mpls** or **config>service>vpls>bgp-evpn>mpls** context.

The configured IP address is used as a next-hop for the MAC/IP, IMET, and AD per-EVI routes advertised for the service. The configured next-hop can be overridden by a policy with the **next-hop-self** command.

10.2.8 EVPN multi-homing support for MPLS tunnels resolved to non-system IPv4/IPv6 addresses

EVPN MPLS multi-homing is supported on PEs that use non-system IPv4 or IPv6 addresses for tunnel resolution. When multi-homing is used in the service, the same IP address must be configured in all three of the following commands to ensure the DF Election candidate list is built correctly.

- The **config>service>system>bgp-evpn>eth-seg>es-orig-ip** command must be configured with the non-system IPv4 or IPv6 address used for the EVPN-MPLS service. This command modifies the originating IP field in the ES routes advertised for the Ethernet segment. The 7705 SAR uses this IP address when adding the local PE as a DF candidate.
- The **config>service>system>bgp-evpn>eth-seg>route-next-hop** command must be configured with the non-system IP address. This command changes the next-hop of the ES and AD per-ES routes to the configured address.
- All the EVPN MPLS services that make use of the Ethernet segment must be configured with the **config>service>vpls|epipe>bgp-evpn>mpls>route-next-hop** command.

10.3 General EVPN topics

This section provides information about general topics related to EVPN:

- [BGP-EVPN MAC mobility](#)
- [BGP-EVPN MAC duplication](#)
- [Conditional static MAC and protection](#)
- [Blackhole MAC](#)
- [CFM interaction with EVPN services](#)
- [BGP and EVPN route selection for EVPN routes](#)
- [Interaction of EVPN and other features](#)

10.3.1 BGP-EVPN MAC mobility

EVPN defines a mechanism to allow the smooth mobility of MAC addresses. The 7705 SAR supports the MAC mobility extended community in MAC advertisement routes as follows:

- The router honors and generates the SEQ (sequence) number in the MAC mobility extended community for MAC moves.
- When a MAC address is EVPN-learned and there is an attempt for it to be learned locally, a BGP update is sent with the SEQ number changed to previous_SEQ+1, except when the MAC duplication *num-moves* value is reached.

- If the sequence number is 0 or if the extended community received in the type 2 MAC/IP advertisement route is not included, then the sequence number is interpreted as 0.
- For MAC mobility, the following MAC address selection procedure is followed:
 - If a PE has two or more active remote EVPN routes for the same MAC address, the highest SEQ number is selected. The tiebreaker is the lowest IP address (BGP next-hop IP address).
 - If a PE has two or more active EVPN routes and the PE is the originator of one of the routes, the highest SEQ number is selected. The tiebreaker is the lowest IP address (BGP next-hop IP address of the remote route is compared to the local system address).



Note: When EVPN multihoming is used in EVPN-MPLS, the ESI is examined to determine whether a MAC address received from two different PEs must be processed within the context of MAC mobility or multihoming. Two MAC routes that are associated with the same remote or local ESI but different PEs are considered reachable through all the PEs. Mobility procedures are not triggered as long as the MAC route still belongs to the same ESI.

10.3.2 BGP-EVPN MAC duplication

EVPN defines a mechanism to protect the EVPN service from control plane churn as a result of loops or accidental duplicated MAC addresses. The 7705 SAR supports an enhanced version of this procedure as described in this section.

A situation may arise where the same MAC address is learned by different PEs in the same VPLS because of two or more hosts being incorrectly configured with the same (duplicate) MAC address. In this situation, the traffic originating from these hosts triggers continuous MAC moves among the PEs attached to these hosts. It is important to recognize this situation and avoid incrementing the sequence number in the MAC mobility attribute to infinity.

To remedy the above situation, a router that detects a MAC mobility event by way of local learning starts a **window minutes** timer (default value is 3 minutes). and if the router detects **num-moves value** before the timer expires (default value is 5 moves), the router concludes that a duplicate MAC situation has occurred. The router then alerts the operator with a trap message. The offending MAC address can be viewed using the **show>service>id n >bgp-evpn command**:

```
10 2018/01/14 01:00:22.91 UTC MINOR: SVCMGR #2331 Base
"VPLS Service 1 has MAC(s)detected as duplicates by EVPN mac-duplication detection."

# show service id 1 bgp-evpn
=====
BGP EVPN Table
=====
MAC Advertisement      : Enabled          Unknown MAC Route      : Disabled
VXLAN Admin Status    : Disabled          Creation Origin        : manual
MAC Dup Detn Moves    : 5                  MAC Dup Detn Window: 3
MAC Dup Detn Retry    : 9                  Number of Dup MACs    : 1
-----
Detected Duplicate MAC Addresses      Time Detected
-----
00:00:00:00:00:12                    01/14/2014 01:00:23
-----
=====
```

After detecting the duplicate address, the router stops sending and processing any BGP MAC advertisement routes for that MAC address until one of the following occurs:

- the MAC address is flushed due to a local event (such as a SAP or SDP binding associated with the MAC address failing) or the reception of a remote update with a higher SEQ number (due to a MAC address flush at the remote router)
- the **retry minutes** timer expires, which flushes the MAC address and restarts the process



Note: The other routers in the VPLS instance will forward the traffic for the duplicate MAC address to the router advertising the best route for the MAC address.

The **num-moves** and **window** values are configurable to allow for the required flexibility in different environments. In scenarios where **bgp>rapid-update evpn** is configured, the operator may want to configure a shorter window timer than in scenarios where BGP updates are sent every default **min-route-advertisement** interval.

The MAC duplication parameters can be configured per VPLS service under the **bgp-evpn>mac-duplication** context, as shown in the following example:

```
*A:DGW1>config>service>vpls>bgp-evpn# info
-----
mac-advertisement
mac-duplication
  detect num-moves 5 window 3
  retry 9
exit
```

10.3.3 Conditional static MAC and protection

RFC 7432 defines the use of the sticky bit in the MAC mobility extended community to signal static MAC addresses. These addresses must be protected in case there is an attempt to dynamically learn them from a different source.

On the 7705 SAR, any conditional static MAC address is advertised by BGP-EVPN as a static address—that is, with the sticky bit set. An example of the configuration of a conditional static MAC address is shown below:

```
*A:PE63>config>service>vpls# info
-----
...
  sap 1/1/1:1000 create
  exit
  static-mac
    mac 00:ca:ca:ca:ca:00 create sap 1/1/1:1000 monitor fwd-status
  exit
  no shutdown

*A:PE64# show router bgp routes evpn mac hunt mac-address 00:ca:ca:ca:ca:00
...
=====
BGP EVPN Mac Routes
=====
Network       : 0.0.0.0/0
Nexthop       : 192.0.2.63
From          : 192.0.2.63
Res. Nexthop  : 192.168.19.1
Local Pref.   : 100
Aggregator AS : None
Atomic Aggr.  : Not Atomic
AIGP Metric   : None
Interface Name : NotAvailable
Aggregator     : None
MED            : 0
```

```

Connector      : None
Community      : target:65000:1000      mac-mobility:Seq: 0/Static
Cluster        : No Cluster Members
Originator Id  : None                    Peer Router Id : 192.0.2.63
Flags          : Used Valid Best IGP
Route Source   : Internal
AS-Path        : No As-Path
EVPN type      : MAC
ESI            : 0:0:0:0:0:0:0:0:0      Tag           : 1063
IP Address     : ::                      RD            : 65063:1000
Mac Address    : 00:ca:ca:ca:ca:00      Mac Mobility   : Seq:0
Neighbor-AS    : N/A
Source Class   : 0                      Dest Class    : 0
-----
Routes : 1
=====

```

Local static MAC addresses or remote MAC addresses with sticky bit are considered to be protected. A packet entering a SAP or SDP binding is discarded if its source MAC address matches one of the protected MAC addresses.

10.3.4 Blackhole MAC

A blackhole MAC is a local FDB record. It is similar to a conditional static MAC; it is associated with a black hole—similar to a VPRN blackhole static route in VPRNs—instead of a SAP or SDP binding. A blackhole MAC can be added by using the following CLI command:

config>service>vpls>static-mac>mac *ieee-address* [create] black-hole

The static blackhole MAC can have security applications for specified MAC addresses (for example, replacement of MAC filters).



Note: A static MAC can only be created as a blackhole MAC if BGP-EVPN is configured first. It is not supported for regular VPLS service, only for an EVPN VPLS service.

For example, when a specified blackhole MAC is added to a service (**static-mac mac 00:00:ca:fe:ca:fe create black-hole**), the following behavior occurs:

- The configured MAC address is created as a static MAC address with a "black-hole" source identifier.

```

*A:PE1# show service id 1 fdb detail
=====
Forwarding Database, Service 1
=====

```

ServId	MAC	Source-Identifier	Type Age	Last Change
1	00:ca:ca:ba:ca:01	eES:	Evpn	06/29/15 23:21:34
		01:00:00:00:00:71:00:00:00:01		
1	00:ca:ca:ba:ca:06	eES:	Evpn	06/29/15 23:21:34
		01:74:13:00:74:13:00:00:74:13		
1	00:ca:00:00:00:00	sap:1/1/1:2	CStatic:P	06/29/15 23:20:58
1	00:ca:fe:ca:fe:00	black-hole	CStatic:P	06/29/15 23:20:00
1	00:ca:fe:ca:fe:69	eMpls:	EvpnS:P	06/29/15 20:40:13
		192.0.2.69:262133		

```

-----
No. of MAC Entries: 5
-----
Legend: L=Learned O=0am P=Protected-MAC C=Conditional S=Static

```

-
- After the blackhole MAC has been successfully added to the FDB, it is treated like any other protected MAC. The blackhole MAC is added as protected (CStatic:P) and advertised in EVPN as static.
 - After the blackhole MAC has been successfully added to the FDB, any frame arriving at any SAP, SDP binding, or EVPN endpoint with a MAC DA that is equal to the blackhole MAC is discarded.

10.3.5 CFM interaction with EVPN services

Ethernet connectivity and fault management (ETH-CFM) allows the operator to validate and measure Ethernet Layer 2 services using standard IEEE 802.1ag and ITU-T Y.1731 protocols. Each tool performs a unique function and adheres to that tool's specific PDU and frame format and the associated rules governing the transmission, interception, and process of the PDU. For more information, see the "ETH OAM capabilities" section in the 7705 SAR OAM and Diagnostics Guide.

EVPN provides powerful solution architectures. ETH-CFM is supported in various Layer 2 EVPN architectures. Because the destination Layer 2 MAC address (unicast or multicast) is ETH-CFM tool-dependent (that is, ETH-CC is sent as a Layer 2 multicast and ETH-DM is sent as a Layer 2 unicast), the ETH-CFM function is allowed to multicast and broadcast to the virtual EVPN connections. The Maintenance Endpoint (MEP) does not populate the local Layer 2 MAC address forwarding database (FDB) with the MAC address related to the MEP. This means that the 48-bit IEEE MAC address is not exchanged with peers and all ETH-CFM frames are broadcast across all virtual connections. To prevent the flooding of unicast packets and allow the remote forwarding databases to learn the remote MEP Layer 2 MAC addresses, the command **cfm-mac-advertisement** must be configured under the **config>service>vpls>bgp-evpn** context. This allows the MEP Layer 2 IEEE MAC addresses to be exchanged with peers. This command tracks configuration changes and sends the required updates via the EVPN notification process related to a change.

Up MEP and Down MEP creation is supported on SAP connections and on spoke and mesh SDP connections in the EVPN service. There is no support for the creation of ETH-CFM MEPs on the virtual connection.

When MEPs are used in combination with EVPN multihoming, the following must be considered:

- Behavior of operationally down MEPs on SAPs or SDPs with EVPN multihoming:
 - all-active multihoming – in this case, no ETH-CFM is expected to be used because the SAPs or SDPs on the PEs are operationally up and active. However, the CE has a single LAG and responds as though it is connected to a single system. In addition, **cfm-mac-advertisement** can lead to traffic loops in all-active multihoming.
 - single-active multihoming – operationally down MEPs defined on single-active Ethernet segment SAPs or SDPs do not send any CCMs when the PE is a non-DF for the Ethernet segment and fault-propagation is configured
- Behavior of operationally up MEPs on Ethernet segment SAPs or SDPs with EVPN multihoming:
 - all-active multihoming – operationally up MEPs defined on non-DF Ethernet segment SAPs or SDPs can send CFM packets. However, they cannot receive CCMs (the SAP or SDP is removed from the default multicast list) or unicast CFM packets (because the MEP MAC address is not installed locally in the FDB, unicast CFM packets are treated as unknown and are not sent to the non-DF MEP).
 - single-active multihoming – operationally up MEPs are able to send or receive CFM packets normally

Because of the above considerations, the use of ETH-CFM in EVPN multihomed SAPs or SDPs is only recommended on operationally down MEPs and single-active multihoming. ETH-CFM is used in this case to notify the CE of the DF or non-DF status.

10.3.6 BGP and EVPN route selection for EVPN routes

When two or more EVPN routes are received at a PE, BGP route selection typically takes place when the route key or the routes are equal. When the route key is different but the PE must make a selection (for instance, the same MAC address is advertised in two routes with different RDs), BGP hands over the routes to EVPN and the EVPN application performs the selection.

EVPN and BGP selection criteria are described below.

- **EVPN route selection for MAC routes**

When two or more routes with the same **mac-length** and **mac** but different route keys are received, BGP hands the routes over to EVPN. EVPN selects the route based on the following tiebreaking order:

1. conditional static MAC addresses (local protected MAC addresses)
2. EVPN static MAC addresses (remote protected MAC addresses)
3. data plane learned MAC addresses (regular learning on SAPs and SDP bindings)
4. EVPN MAC addresses with higher SEQ number
5. lowest IP address (next-hop IP address of the EVPN NLRI)
6. lowest Ethernet tag (the tag is 0 for MPLS)
7. Lowest RD

- **BGP route selection:**

The regular BGP route selection is followed.



Note: If BGP runs a route selection operation and a specified—but otherwise valid—EVPN route loses a tiebreaker to another EVPN route, the nonselected route can be displayed, along with a tiebreaker reason, using the **show>router>bgp>routes>evpn evpn-type** command.

10.3.7 Interaction of EVPN and other features

This section contains information about the following topics:

- [Interaction of EVPN-MPLS with existing VPLS features](#)
- [Routing policies for BGP-EVPN IP prefixes](#)

10.3.7.1 Interaction of EVPN-MPLS with existing VPLS features

When enabling existing VPLS features in an EVPN-MPLS service, the following must be considered:

- EVPN-MPLS is only supported in regular VPLS. Other VPLS types, such as management VPLS (mVPLS), are not supported with EVPN-MPLS.
- In general, no router-generated control packets are sent to the EVPN destination bindings, except for ETH-CFM for EVPN-MPLS.

- STP and mVPLS services:
 - STP can be configured in BGP-EVPN services. BPDUs are not sent over the EVPN bindings.
 - The **bgp-evpn** command is blocked in mVPLS services; however, a separate mVPLS service can manage a SAP or spoke SDP in a BGP-EVPN service.
- The **mac-move** command can be used in BGP-EVPN VPLS services for SAPs and SDP bindings; however, the MAC addresses being learned through BGP-EVPN will not be considered.



Note: The MAC duplication function already provides protection against MAC moves between EVPN and SAPs and SDP bindings.

- The **disable-learning** command and other FDB-related tools only work for data-plane learned MAC addresses.
- The **mac-protect** command cannot be used in conjunction with EVPN.



Note: EVPN provides its own protection mechanism for static MAC addresses.

- MAC OAM tools are not supported for BGP-EVPN services, including the following commands: **mac-ping**, **mac-trace**, **mac-populate**, **mac-purge**, and **cpe-ping**.
- SAPs and SDP bindings that belong to a specified ES but are configured on non-BGP-EVPN MPLS-enabled VPLS or Epipe services are kept operationally down with the **StandbyForMHPProtocol** flag.
- CPE ping is not supported for EVPN services. CPE ping packets are not sent over EVPN destinations.
- Other functions not supported in conjunction with the **bgp-evpn** command include:
 - endpoints and attributes
 - MLD snooping and attributes
 - BPDUs translation
 - MAC pinning

10.3.7.2 Routing policies for BGP-EVPN IP prefixes

BGP routing policies are supported for IP prefixes imported or exported through BGP-EVPN.

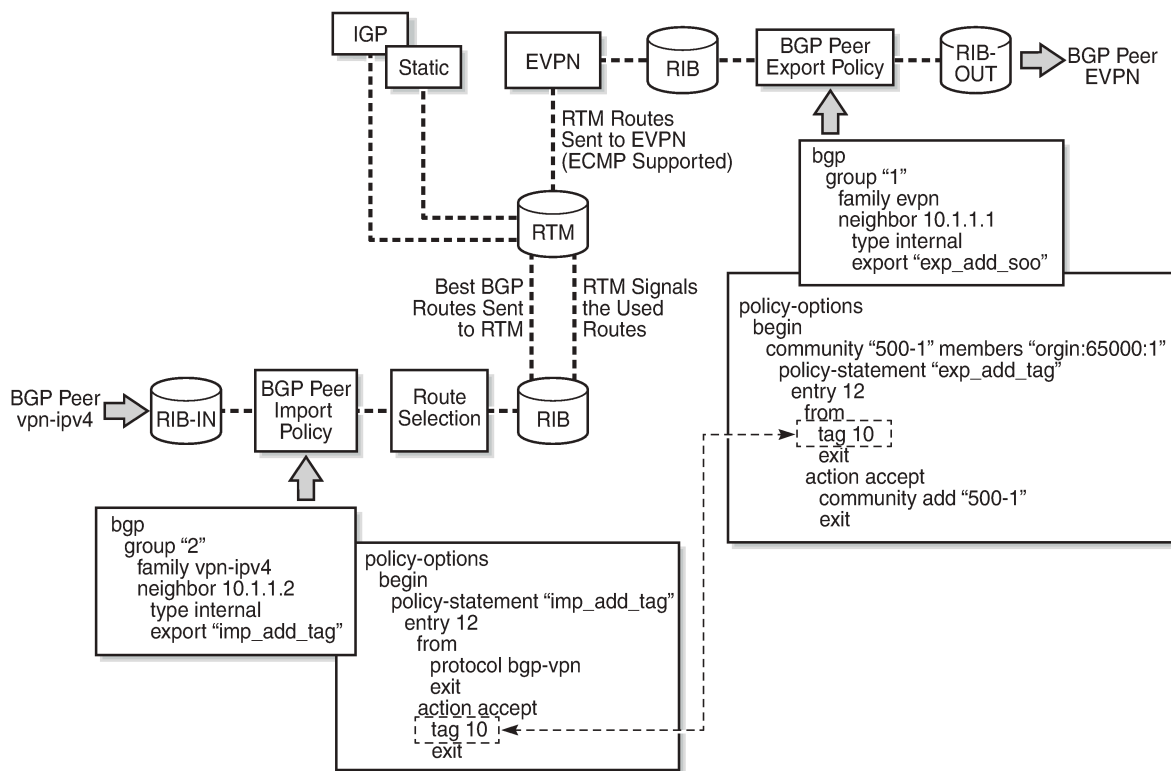
When applying routing policies to control the distribution of prefixes between EVPN and IP-VPN, the user must consider that both families are separate as far as BGP is concerned, and when prefixes are imported in the VPRN routing table, the BGP attributes are lost to the other family. The use of route tags allows the controlled distribution of prefixes across the two families.

The following figure shows an example of how VPN-IPv4 routes are imported into the routing table manager (RTM) and then passed to EVPN for processing.



Note: VPN-IPv4 routes can be tagged at ingress, and that tag is preserved throughout the RTM and EVPN processing so that the tag can be matched at the egress BGP routing policy.

Figure 174: IP-VPN import and EVPN export BGP workflow



28723

Policy tags can be used to match EVPN IP prefixes that were learned not only from BGP VPN-IPv4 but also from other routing protocols. The tag range supported for each protocol is different:


```

<tag> : accepts in decimal or hex
        [0x1..0xFFFFFFFF]H (for OSPF and IS-IS)
        [0x1..0xFFFF]H (for RIP)
        [0x1..0xFF]H (for BGP)

```

The following figure shows an example of the reverse workflow: routes imported from EVPN and exported from RTM to BGP VPN-IPv4.

Figure 175: EVPN import and IP-VPN export BGP workflow



28724

The policy behavior for EVPN IP prefixes is summarized in the following statements:

- for EVPN prefix routes received and imported in RTM:
 - policy entries can match on communities and add tags. This works at the peer level or at the vsi-import level.
 - policy entries can match on **family evpn**
- for exporting RTM to EVPN prefix routes:
 - policy entries can match on tags and based on that, add communities, accept, or reject. This works at the peer level or the virtual switch instance (VSI) export level.

Policy entries can add tags for static routes, RIP, OSPF, IS-IS, and BGP that can then be matched on the BGP peer export policy or VSI export policy for EVPN prefix routes.

10.4 Configuring an EVPN service with CLI

This section provides examples for configuring EVPN multihoming for VPLS services using the CLI:

- [EVPN all-active multihoming configuration example](#)
- [EVPN single-active multihoming configuration example](#)
- [EVPN-MPLS r-VPLS configuration examples](#)

10.4.1 EVPN all-active multihoming configuration example

This section shows a configuration example for three 7705 SAR PEs, with the following assumptions:

- PE-1 and PE-2 are multihomed to CE-12, which uses a LAG to connect to the network. CE-12 is connected to LAG SAPs configured in an all-active multihoming Ethernet segment.
- PE-3 is a remote PE that performs aliasing for traffic destined for the CE-12.

The following configuration excerpt applies to a VPLS-1 on PE-1 and PE-2 and includes the corresponding **ethernet-segment** and **lag** commands.

```
A:PE1# configure lag 1
A:PE1>config>lag# info
-----
mode access
encap-type dot1q
port 1/1/2
lacp active administrative-key 1 system-id 00:00:00:00:69:72
no shutdown
-----

A:PE1>config>lag# /configure service system bgp-evpn
A:PE1>config>service>system>bgp-evpn# info
-----
route-distinguisher 192.0.2.69:0
ethernet-segment "ESI-71" create
esi 0x010000000007100000001
es-activation-timer 10
service-carving
mode auto
exit
multi-homing all-active
lag 1
no shutdown
exit
-----

A:PE1>config>service>system>bgp-evpn# /configure service vpls 1
A:PE1>config>service>vpls# info
-----
bgp
exit
bgp-evpn
cfm-mac-advertisement
evi 1
exit
mpls
ingress-replication-bum-label
auto-bind-tunnel
resolution any
exit
no shutdown
exit
exit
stp
shutdown
exit
sap lag-1:1 create
exit
no shutdown
```



```

-----
A:PE2# configure lag 1
A:PE2>config>lag# info
-----
    mode access
    encap-type dot1q
    port 1/1/3
    lacp active administrative-key 1 system-id 00:00:00:00:69:72
    no shutdown
-----

A:PE2>config>lag# /configure service system bgp-evpn
A:PE2>config>service>system>bgp-evpn# info
-----
    route-distinguisher 192.0.2.72:0
    ethernet-segment "ESI-71" create
        esi 0x010000000007100000001
        es-activation-timer 10
        service-carving
            mode auto
        exit
    multi-homing all-active
    lag 1
    no shutdown
    exit
-----

A:PE2>config>service>system>bgp-evpn# /configure service vpls 1
A:PE2>config>service>vpls# info
-----
    bgp
    exit
    bgp-evpn
        cfm-mac-advertisement
        evi 1
        exit
        mpls
            ingress-replication-bum-label
            auto-bind-tunnel
            resolution any
        exit
        no shutdown
    exit
    exit
    stp
        shutdown
    exit
    sap lag-1:1 create
    exit
    no shutdown
-----

```

The configuration on the remote PE (PE-3), which supports aliasing to PE-1 and PE-2 is shown below. PE-3 does not have an Ethernet segment configured. It only requires the VPLS-1 configuration and ECMP > 1 to perform aliasing.

```

*A:PE3>config>service>vpls# info
-----
    bgp
    exit
    bgp-evpn

```

```

        cfm-mac-advertisement
        evi 1
        exit
        mpls
            ingress-replication-bum-label
            ecmp 4
            auto-bind-tunnel
                resolution any
            exit
            no shutdown
        exit
    exit
    stp
        shutdown
    exit
    sap 1/1/1:1 create
    exit
    spoke-sdp 4:13 create
        no shutdown
    exit
    no shutdown
-----

```

10.4.2 EVPN single-active multihoming configuration example

To use single-active multihoming on PE-1 and PE-2 instead of all-active multihoming, make the following modifications to the configuration of the [EVPN all-active multihoming configuration example](#):

- change the LAG configuration to **single-active**
CE-12 is now configured with two different LAGs; therefore, the *admin-key*, *system-id*, and *system-priority* must be different on PE-1 and PE-2.
- change the Ethernet segment configuration to **single-active**

No changes are needed at the service level on any of the three PEs.

The differences between single-active and all-active multihoming are highlighted in **bold** in the following configuration excerpts:

```

A:PE1# configure lag 1
A:PE1>config>lag# info
-----
    mode access
    encap-type dot1q
    port 1/1/2
    lACP active administrative-key 1 system-id 00:00:00:00:69:69
    no shutdown
-----

A:PE1>config>lag# /configure service system bgp-evpn
A:PE1>config>service>system>bgp-evpn# info
-----
    route-distinguisher 192.0.2.69:0
    ethernet-segment "ESI-71" create
        esi 0x0100000000071000000001
        es-activation-timer 10
        service-carving
            mode auto
        exit
        multi-homing single-active

```

```

        lag 1
        no shutdown
    exit
-----

A:PE2# configure lag 1
A:PE2>config>lag# info
-----
    mode access
    encap-type dot1q
    port 1/1/3
    lacp active administrative-key 1 system-id 00:00:00:00:72:72
    no shutdown
-----

A:PE2>config>lag# /configure service system bgp-evpn
A:PE2>config>service>system>bgp-evpn# info
-----
    route-distinguisher 192.0.2.72:0
    ethernet-segment "ESI-71" create
        esi 0x01000000007100000001
        es-activation-timer 10
        service-carving
            mode auto
        exit
        multi-homing single-active
        lag 1
        no shutdown
    exit
-----

```

10.4.3 EVPN-MPLS r-VPLS configuration examples

This section contains the following configuration examples:


- [EVPN-MPLS r-VPLS without multihoming](#)
- [EVPN-MPLS r-VPLS with all-active multihoming](#)
- [EVPN-MPLS r-VPLS with single-active multihoming](#)

EVPN can be used as the unified control plane VPN technology, not only for providing Layer 2 connectivity, but also for Layer 3 (inter-subnet forwarding). EVPN for MPLS tunnels, along with multihoming and passive VRRP, provides efficient Layer 2 or Layer 3 connectivity to distributed hosts and routers.

10.4.3.1 EVPN-MPLS r-VPLS without multihoming

The first scenario describes r-VPLS support including IP route advertisement (BGP-EVPN route type 5) with EVPN tunnel interfaces without multihoming. VPLS 101 does not have a connected host, but the linked VPRN has SAP 1/2/1:10. The following figure shows an example of topology used for r-VPLS with an EVPN tunnel but without multihoming. IP prefixes are advertised.

Figure 176: R-VPLS with EVPN tunnel, without multihoming



26851

The initial configuration includes the following:

- cards, MDAs, ports
- router interface between PE-2 and PE-3
- IS-IS (or OSPF)
- LDP enabled on the router interface between PE-2 and PE-3

BGP is configured for address family EVPN on PE-2 and PE-3. The BGP configuration on PE-2 is as follows. The BGP configuration on PE-3 is similar.

```
# on PE-2:
configure
router
  autonomous-system 64500
  bgp
    family evpn
      vpn-apply-import
      vpn-apply-export
      enable-peer-tracking
      rapid-withdrawal
      rapid-update evpn
      group "internal"
        peer-as 64500
        neighbor 192.0.2.3
      exit
    exit
  exit
exit
```

The CEs are connected to SAP 1/2/1:10 in VPRN 10. r-VPLS 101 is bound to VPRN 10 and VPRN 10 has a dedicated interface "int-evi-101" for the EVPN tunnel. In general, if only one route target (RT) is used for import and export in the EVPN-VPLS, it is best to add the EVI and have the route distinguisher (RD) and RT auto-derived from the EVI; this simplifies the configuration and reduces the chance of errors. The service configuration on PE-2 is as follows:

```
# on PE-2:
configure
service
  vprn 10 name "VPRN 10" customer 1 create
    route-distinguisher 192.0.2.2:10
    vrf-target target:64500:10
    interface "int-PE-2-CE-20" create
      address 172.16.2.1/24
      sap 1/2/1:10 create
```

```

        exit
    exit
    interface "int-evi-101" create
        vpls "evi-101"
evpn-tunnel
    exit
    exit
    no shutdown
exit
vpls 101 name "evi-101" customer 1 create
    allow-ip-int-bind
    exit
    bgp // RD and RT are not manually configured in BGP context
    exit
    bgp-evpn
        ip-route-advertisement
        evi 101 // RD and RT will be auto-derived from the EVI
        mpls bgp 1
            auto-bind-tunnel
            resolution any
        exit
        no shutdown
    exit
    exit
    no shutdown
exit
exit

```

In the preceding configuration:

- the **allow-ip-int-binding** command is required so that r-VPLS 101 can be bound to VPRN 10
- the service name is required and the configured name "evi-101" must match the name in the VPRN 10 VPLS interface. The service name is configured at service creation time.
- the VPRN 10 VPLS interface is configured with the keyword **evpn-tunnel**. This configuration has the advantage of not having to allocate IP addresses to the r-VPLS interfaces; however, it cannot be used when the r-VPLS interface has local SAPs.

The configuration is similar on PE-3. The RD must be different on PE-2 and PE-3; this is automatically the case when the RD is auto-derived from the configured EVI, as in the example. The RD on PE-2 is 192.0.2.2:101; on PE-3, the RD is 192.0.2.3:101.

PE-3 receives the following BGP-EVPN IP prefix route for prefix 172.16.2.0/24 from PE-2:

```

34 2019/09/27 12:21:38.100 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
    Withdrawn Length = 0
    Total Path Attr Length = 97
    Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
        Address Family EVPN
        NextHop len 4 NextHop 192.0.2.2
        Type: EVPN-IP-Prefix Len: 34 RD: 192.0.2.2:101, tag: 0,
ip_prefix: 172.16.2.0/24 gw_ip 0.0.0.0 Label: 8388464
    Flag: 0x40 Type: 1 Len: 1 Origin: 0
    Flag: 0x40 Type: 2 Len: 0 AS Path:
    Flag: 0x80 Type: 4 Len: 4 MED: 0
    Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
    Flag: 0xc0 Type: 16 Len: 24 Extended Community:
        target:64500:101
        mac-nh:04:0b:ff:ff:ff:a2
    bgp-tunnel-encap:MPLS
"
```

GW IP 0.0.0.0 is an indication that an EVPN tunnel is in use. With EVPN tunnels, no IRB IP address needs to be configured in the VPRN. EVPN tunnels make provisioning easier to automate and save IP addresses from the tenant IP space.

The BGP tunnel encapsulation is MPLS, but the MPLS label in the debug message is not the same as in the service, because the router strips the extra four lowest bits to get the 20-bit MPLS label. In the debug message, the label is 8388464. This is because the debug message is shown before the router can parse the label field and see if it corresponds to an MPLS label (20 bits). The MPLS label is calculated by dividing the label value by 24 (16), as follows: $8388464/16 = 524279$.

The MAC next-hop extended community 04:0b:ff:ff:a2 is the MAC address of the interface "int-evi-101" in VPRN 10 on PE-2, as follows:

```
*A:PE-2# show service id 10 interface "int-evi-101" detail | match MAC
MACSec          : N/A
MAC Address      : 04:0b:ff:ff:ff:a2    Mac Accounting    : Disabled
```

The routing table for VPRN 10 on PE-3 contains the route for prefix 172.16.2.0/24 as the BGP-EVPN route with next-hop "int-evi-101" and interface name "ET-04:0b:ff:ff:a2" (ET stands for EVPN Tunnel), as follows:

```
*A:PE-3# show router 10 route-table

=====
Route Table (Service: 10)
=====
Dest Prefix[Flags]                Type   Proto   Age           Pref
Next Hop[Interface Name]          Metric
-----
172.16.2.0/24                    Remote BGP EVPN 00h06m45s    169
  int-evi-101 (ET-04:0b:ff:ff:a2) 0
172.16.3.0/24                    Local  Local    00h06m48s    0
  int-PE-3-CE-30                  0
-----
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
=====
```

The forwarding database (FDB) for VPLS 101 on PE-3 shows an entry for MAC address 04:0b:ff:ff:a2 that is learned via EVPN. The MAC address is static (S) and protected (P). The MPLS label is 524279.

```
*A:PE-3# show service id 101 fdb detail

=====
Forwarding Database, Service 101
=====
ServId  MAC                Source-Identifier  Type   Last Change
Transport:Tnl-Id
-----
101     04:0b:ff:ff:ff:a2 eMpls:           EvpnS:P 09/27/19 12:55:59
        192.0.2.2:524279
        ldp:65538
101     04:0d:ff:ff:ff:a2 cpm              Intf    09/27/19 12:55:57
-----
No. of MAC Entries: 2
Legend: L=Learned O=0am P=Protected-MAC C=Conditional S=Static Lf=Leaf
```

When the CEs have IPv6 addresses, the VPRN configuration is similar on the PEs, but the **ipv6** context must be enabled on the EVPN tunnel interface so that the router can advertise and process BGP-EVPN type 5 routes with IPv6 prefixes. The configuration of VPLS is identical for IPv4 and IPv6.

```
=====
# on PE-2:
configure
service
  vprn 16 name "VPRN 16" customer 1 create
    route-distinguisher 192.0.2.2:16
    vrf-target target:64500:16
    interface "int-PE-2-CE-26" create
      ipv6
        address 2001:db8:16::2:1/120
      exit
      sap 1/2/1:16 create
      exit
    exit
  interface "int-evi-106" create
    ipv6
    exit
    vpls "evi-106"
      evpn-tunnel
    exit
  exit
  no shutdown
exit
vpls 106 name "evi-106" customer 1 create
  allow-ip-int-bind
  exit
  bgp
  exit
  bgp-evpn
    ip-route-advertisement
    evi 106
    mpls bgp 1
      auto-bind-tunnel
      resolution any
    exit
    no shutdown
  exit
  exit
  no shutdown
exit
```

When advertising IPv6 prefixes, the GW IP field for route type 5 is always populated with the IPv6 address of the r-VPLS interface. In this example, because no specific IPv6 global address is configured, the GW IP will be populated with the auto-created link local address. The following BGP update is received by PE-3 for IPv6 prefix 2001:db8:16::2:0/120:

```
# on PE-3:
36 2019/09/27 12:21:38.123 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Received BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 113
  Flag: 0x90 Type: 14 Len: 69 Multiprotocol Reachable NLRI:
    Address Family EVPN
    NextHop len 4 NextHop 192.0.2.2
    Type: EVPN-IP-Prefix Len: 58 RD: 192.0.2.2:106, tag: 0,
```

```
ip_prefix: 2001:db8:16::2:0/120 gw_ip fe80::60b:1ff:fe02:1 Label: 8388448
  Flag: 0x40 Type: 1 Len: 1 Origin: 0
  Flag: 0x40 Type: 2 Len: 0 AS Path:
  Flag: 0x80 Type: 4 Len: 4 MED: 0
  Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
  Flag: 0xc0 Type: 16 Len: 16 Extended Community:
    target:64500:106
    bgp-tunnel-encap:MPLS
"
```

The IPv6 route table on PE-3 is as follows:

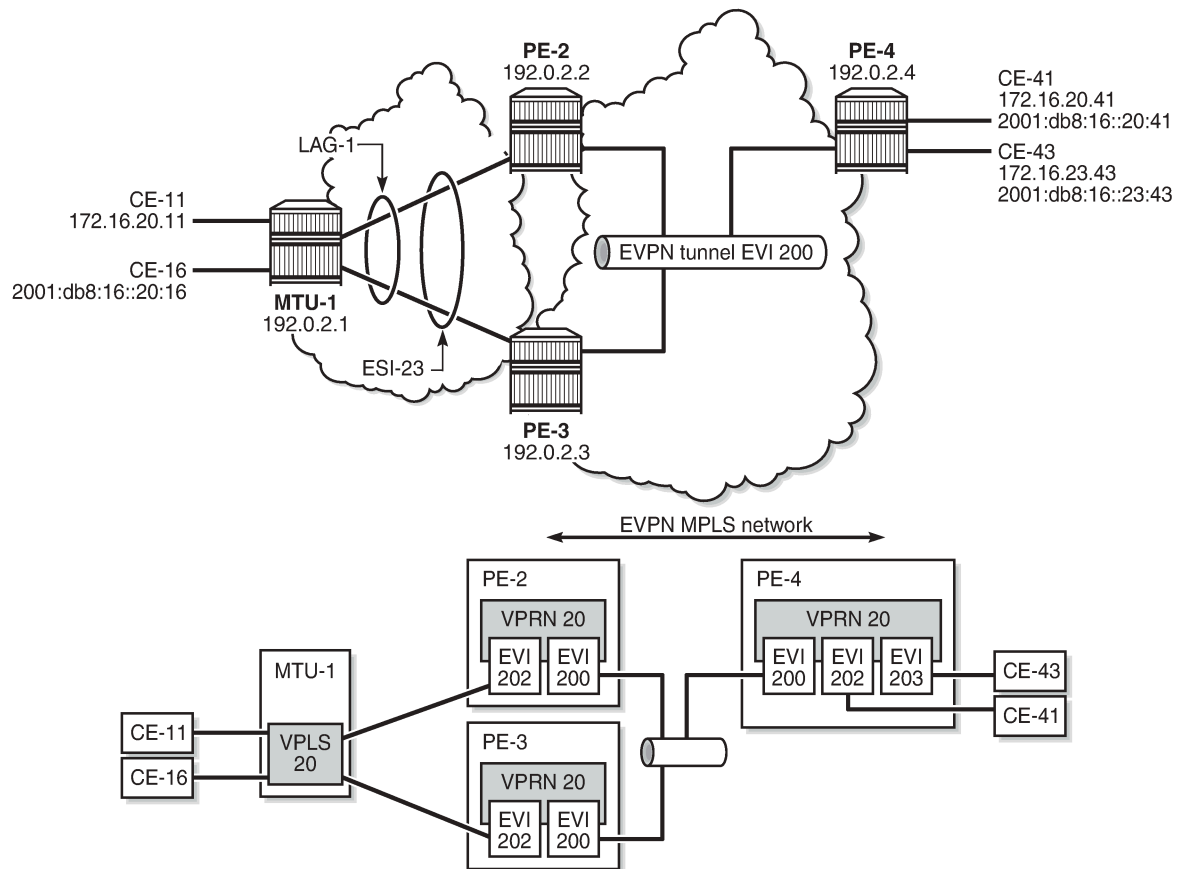
```
*A:PE-3# show router 16 route-table ipv6

=====
IPv6 Route Table (Service: 16)
=====
Dest Prefix[Flags]                Type  Proto  Age      Pref
  Next Hop[Interface Name]          Metric
-----
2001:db8:16::2:0/120              Remote BGP  EVPN  00h17m24s 169
      fe80::60b:1ff:fe02:1-"int-evi-106"
2001:db8:16::3:0/120              Local  Local  00h17m26s  0
      int-PE-3-CE-36
-----
No. of Routes: 2
```

10.4.3.2 EVPN-MPLS r-VPLS with all-active multihoming

The following figure shows an example of the topology with all-active multihoming Ethernet segment (ES) "ESI-23".

Figure 177: EVPN-MPLS r-VPLS with all-active multihoming ES



26852

BGP is configured between PE-2, PE-3, and PE-4 for address family EVPN. The configuration on PE-2 is as follows:

```
# on PE-2:
configure
router
  autonomous-system 64500
  bgp
    family evpn
      vpn-apply-import
      vpn-apply-export
      enable-peer-tracking
      rapid-withdrawal
      rapid-update evpn
      group "internal"
        peer-as 64500
        neighbor 192.0.2.3
      exit
        neighbor 192.0.2.4
      exit
    exit
  exit
exit
```

All-active multihoming Ethernet segment "ESI-23" is configured on PE-2 and PE-3 as follows:

```
configure
  service
    system
      bgp-evpn
        ethernet-segment "ESI-23" create
          esi 01:00:00:00:00:23:00:00:00:01
          es-activation-timer 3
          service-carving
            mode auto
          exit
          multi-homing all-active
lag 1
          no shutdown
        exit
```

The following services are configured on the PEs:

- VPRN 20 has interfaces bound to VPLS 200 and VPLS 202. On PE-4, VPRN 20 also has an interface bound to VPLS 203.
- VPLS 200 is configured as an EVPN tunnel that connects the PEs.
- VPLS 202 and VPLS 203 have attachment circuits to CEs.

The services are configured on PE-2 as follows. The configuration on PE-3 and PE-4 is similar.

```
# on PE-2:
configure
  service
    vprn 20 name "VPRN 20" customer 1 create
      route-distinguisher 192.0.2.2:20
      vrf-target target:64500:20
      interface "int-evi-202" create
        address 172.16.20.2/24
        mac 00:ca:fe:00:02:02
        vrrp 1 passive
backup 172.16.20.254
        ping-reply
        traceroute-reply
      exit
      ipv6
        address 2001:db8:16::20:2/120
link-local-address fe80::16:20:2
vrrp 1 passive
        backup fe80::16:20:fe
        ping-reply
        traceroute-reply
      exit
      exit
      vpls "evi-202"
      exit
    exit
    interface "int-evi-200" create
      ipv6
      exit
      vpls "evi-200"
      evpn-tunnel
      exit
    exit
    router-advertisement
interface "int-evi-202"
use-virtual-mac
```

```

no shutdown
    exit
    exit
    no shutdown
exit
vpls 200 name "evi-200" customer 1 create
    allow-ip-int-bind
    exit
    bgp
    exit
    bgp-evpn
        ip-route-advertisement
        evi 200
        mpls bgp 1
            auto-bind-tunnel
            resolution any
        exit
        no shutdown
    exit
    exit
    no shutdown
exit
vpls 202 name "evi-202" customer 1 create
    allow-ip-int-bind
    exit
    bgp
    exit
    bgp-evpn
        evi 202
        mpls bgp 1
            auto-bind-tunnel
            resolution any
        exit
        no shutdown
    exit
    exit
    sap lag-1:20 create
    exit
    no shutdown
exit

```

The IPv6 VRRP backup address is in the same subnet as the link local address of the interface "int-evi-202". The IPv6 address can be set as preferred. Also for IPv6, router advertisement must be enabled and configured to use the virtual MAC address.

10.4.3.2.1 Passive VRRP

EVI 202 is configured as an r-VPLS interface with passive VRRP. A passive VRRP VRID instance suppresses the transmission and reception of keepalive messages. All PEs configured with passive VRRP become VRRP masters and take ownership of the virtual IP and MAC addresses.

Each individual r-VPLS interface has a different MAC/IP address on each PE. The MAC/IP addresses for "int-evi-202" on PE-2 are MAC 00:ca:fe:00:02:02 and IP 172.16.20.2/24 for IPv4 and the same MAC address with IPv6 2001:db8:16::20:2 and fe80::16:20:2. However, the r-VPLS interfaces on all PEs share the same VRID 1 and backup IP address 172.16.20.254, so the same vMAC/vIP 00:00:5e:00:01:01/172.16.20.254 and vMAC/vIP 00:00:5e:00:02:01/ fe80::16:20:fe are advertised by all PEs. PE-2 advertises the following EVPN MAC routes:

```
82 2019/09/27 12:20:38.600 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.3
```

```

"Peer 1: 192.0.2.3: UPDATE
Peer 1: 192.0.2.3 - Send BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 292
  Flag: 0x90 Type: 14 Len: 240 Multiprotocol Reachable NLRI:
    Address Family EVPN
    NextHop len 4 NextHop 192.0.2.2
    Type: EVPN-MAC Len: 49 RD: 192.0.2.2:202 ESI: ESI-0, tag: 0, mac len: 48
mac: 00:00:5e:00:02:01, IP len: 16, IP: fe80::16:20:fe, label1: 8388384
    Type: EVPN-MAC Len: 37 RD: 192.0.2.2:202 ESI: ESI-0, tag: 0, mac len: 48
mac: 00:00:5e:00:01:01, IP len: 4, IP: 172.16.20.254, label1: 8388384
    Type: EVPN-MAC Len: 49 RD: 192.0.2.2:202 ESI: ESI-0, tag: 0, mac len: 48
    mac: 00:ca:fe:00:02:02, IP len: 16, IP: fe80::16:20:2, label1: 8388384
    Type: EVPN-MAC Len: 49 RD: 192.0.2.2:202 ESI: ESI-0, tag: 0, mac len: 48
    mac: 00:ca:fe:00:02:02, IP len: 16, IP: 2001:db8:16::20:2, label1: 8388384
    Type: EVPN-MAC Len: 37 RD: 192.0.2.2:202 ESI: ESI-0, tag: 0, mac len: 48
    mac: 00:ca:fe:00:02:02, IP len: 4, IP: 172.16.20.2, label1: 8388384
  Flag: 0x40 Type: 1 Len: 1 Origin: 0
  Flag: 0x40 Type: 2 Len: 0 AS Path:
  Flag: 0x80 Type: 4 Len: 4 MED: 0
  Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
  Flag: 0xc0 Type: 16 Len: 24 Extended Community:
    target:64500:202
    bgp-tunnel-encap:MPLS
    mac-mobility:Seq:0/Static
"

```

The three PEs advertise the same (anycast) vMAC/vIP in EVI 202 as protected, but each PE keeps its own MAC entry in the FDB. The following FDB output shows that the source identifier for vMAC 00:00:5e:00:01:01 and vMAC 00:00:5e:00:02:01 is the CPM. These two vMAC entries with source identifier CPM are seen on all PEs.

```
*A:PE-2# show service id 202 fdb detail
```

```
=====
Forwarding Database, Service 202
=====
```

ServId	MAC Transport:Tnl-Id	Source-Identifier	Type Age	Last Change
202	00:00:01:00:00:11	sap:lag-1:20	L/90	09/27/19 12:00:35
202	00:00:01:00:00:16	sap:lag-1:20	L/90	09/27/19 12:00:36
202	00:00:04:00:00:41	eMpls: 192.0.2.4:524279	Evpn	09/27/19 11:57:24
	ldp:65539			
202	00:00:5e:00:01:01	cpm	Intf	09/27/19 12:20:19
202	00:00:5e:00:02:01	cpm	Intf	09/27/19 12:20:19
202	00:ca:fe:00:02:02	cpm	Intf	09/27/19 11:56:56
202	00:ca:fe:00:02:03	eMpls: 192.0.2.3:524274	EvpnS:P	09/27/19 11:57:12
	ldp:65537			
202	00:ca:fe:00:02:04	eMpls: 192.0.2.4:524279	EvpnS:P	09/27/19 11:57:23
	ldp:65539			

```
-----
No. of MAC Entries: 8
-----
```

```
Legend: L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
=====
```

The interface MAC 00:ca:fe:00:02:02 is local, so it also has the CPM as source identifier. MAC 00:ca:fe:00:02:03 is the r-VPLS interface MAC for PE-3 and it is learned via EVPN-MPLS (eMpls) as static (S) and protected (P). MAC address 00:ca:fe:00:02:04 on PE-4 is also static and protected.

PE-4 sends the following IP prefix route (BGP-EVPN route type 5) for prefix 172.16.23.0/24 to the other PEs:

```
35 2019/09/27 12:20:38.600 UTC MINOR: DEBUG #2001 Base Peer 1: 192.0.2.2
"Peer 1: 192.0.2.2: UPDATE
Peer 1: 192.0.2.2 - Send BGP UPDATE:
  Withdrawn Length = 0
  Total Path Attr Length = 97
  Flag: 0x90 Type: 14 Len: 45 Multiprotocol Reachable NLRI:
    Address Family EVPN
    NextHop len 4 NextHop 192.0.2.4
    Type: EVPN-IP-Prefix Len: 34 RD: 192.0.2.4:200, tag: 0,
      ip_prefix: 172.16.23.0/24 gw_ip 0.0.0.0 Label: 8388384
  Flag: 0x40 Type: 1 Len: 1 Origin: 0
  Flag: 0x40 Type: 2 Len: 0 AS Path:
  Flag: 0x80 Type: 4 Len: 4 MED: 0
  Flag: 0x40 Type: 5 Len: 4 Local Preference: 100
  Flag: 0xc0 Type: 16 Len: 24 Extended Community:
    target:64500:200
    mac-nh:04:0f:ff:00:00:05
    bgp-tunnel-encap:MPLS
"
```

The IP prefixes are advertised with the next hop equal to the EVPN-tunnel GW MAC "int-evi-200", as follows:

```
*A:PE-4# show router 20 interface "int-evi-200" detail | match MAC
MACSec          : N/A
MAC Address      : 04:0f:ff:00:00:05   Mac Accounting   : Disabled
```

The routing table for VPRN 20 on PE-2 contains IP prefix 172.16.23.0/24 with next hop 04:0f:ff:00:00:05, as follows:

```
*A:PE-2# show router 20 route-table

=====
Route Table (Service: 20)
=====
Dest Prefix[Flags]                Type   Proto   Age           Pref
  Next Hop[Interface Name]                               Metric
-----
172.16.20.0/24                    Local   Local   00h01m07s    0
   int-evi-202                      0
172.16.23.0/24                    Remote  BGP EVPN 00h00m48s    169
int-evi-200 (ET-04:0f:ff:00:00:05) 0
-----
No. of Routes: 2
```

The following IPv6 routing table for VPRN 20 on PE-2 contains prefix 2001:db8:16::23:0/120, which has also been advertised by PE-4. The next hop is again "int-evi-200", only this time the link local IPv6 address is displayed (GW IP) instead of the MAC address. The next hop is the GW IP value for route type 5, as long as it is a non-zero value. When the GW IP address is 0, route type 5 is expected to contain a **mac-nh** extended community. The MAC encoded in the extended community is used as the next hop in that case.

```
*A:PE-2# show router 20 route-table ipv6
```

```

=====
IPv6 Route Table (Service: 20)
=====
Dest Prefix[Flags]                                Type  Proto  Age      Pref
  Next Hop[Interface Name]                        Metric
-----
2001:db8:16::20:0/120                             Local  Local   00h01m06s  0
      int-evi-202                                0
2001:db8:16::23:0/120                             Remote BGP EVPN 00h00m47s 169
fe80::a3:a899:473e:c489 -"int-evi-200"            0
-----
No. of Routes: 2

```

The EVPN tunnel service VPLS 200 has all the MAC addresses of the EVPN interfaces within VPRN 20 as static (S) and protected (P), as follows:

```

*A:PE-2# show service id "evi-200" fdb detail

=====
Forwarding Database, Service 200
=====
ServId   MAC                Source-Identifier  Type      Last Change
  Transport:Tnl-Id
-----
200      04:0b:ff:00:00:05 cpm               Intf      09/27/19 12:20:31
200      04:0d:ff:00:00:05 eMpls:          EvpnS:P   09/27/19 12:20:39
      192.0.2.3:524275
      ldp:65537
200      04:0f:ff:00:00:05 eMpls:          EvpnS:P   09/27/19 12:20:51
      192.0.2.4:524280
      ldp:65539
-----
No. of MAC Entries: 3
-----
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
=====

```

The VRRP instance in each PE is master, as follows:

```

*A:PE-2# show router 20 vrrp instance

=====
VRRP Instances
=====
Interface Name      VR Id Own Adm State      Base Pri  Msg Int
  IP               IP    Opr Pol Id   InUse Pri  Inh Int
-----
int-evi-202         1    No  Up  Master    100      1
  IPv4              Up      n/a    100      No
  Backup Addr: 172.16.20.254
int-evi-202         1    No  Up  Master    100      1
  IPv6              Up      n/a    100      Yes
  Backup Addr: fe80::16:20:fe
-----
Instances : 2
=====

```

```

*A:PE-3# show router 20 vrrp instance

=====

```

VRRP Instances

Interface Name	VR Id	Own	Adm	State	Base Pri	Msg Int
IP			Opr	Pol Id	InUse Pri	Inh Int
int-evi-202	1	No	Up	Master	100	1
	IPv4		Up	n/a	100	No
Backup Addr: 172.16.20.254						
int-evi-202	1	No	Up	Master	100	1
	IPv6		Up	n/a	100	Yes
Backup Addr: fe80::16:20:fe						
Instances : 2						

```
*A:PE-4# show router 20 vrrp instance
```

VRRP Instances

Interface Name	VR Id	Own	Adm	State	Base Pri	Msg Int
IP			Opr	Pol Id	InUse Pri	Inh Int
int-evi-202	1	No	Up	Master	100	1
	IPv4		Up	n/a	100	No
Backup Addr: 172.16.20.254						
int-evi-203	2	No	Up	Master	100	1
	IPv4		Up	n/a	100	No
Backup Addr: 172.16.23.254						
int-evi-202	1	No	Up	Master	100	1
	IPv6		Up	n/a	100	Yes
Backup Addr: fe80::16:20:fe						
int-evi-203	2	No	Up	Master	100	1
	IPv6		Up	n/a	100	Yes
Backup Addr: fe80::16:23:fe						
Instances : 4						

10.4.3.2.2 Operation

On PE-4, VPRN 20 has one interface bound to VPLS 202 and another interface bound to VPLS 203. CE-41 is attached to VPLS 202 and CE-43 is attached to VPLS 203. When ping messages are sent from CE-41 to CE-43, or vice versa, the messages go via VPRN 20, which has routes to both CEs, as follows:

```
*A:PE-4# show router 20 route-table
```

Route Table (Service: 20)

Dest Prefix[Flags]	Type	Proto	Age	Pref
Next Hop[Interface Name]			Metric	
172.16.20.0/24	Local	Local	04h25m52s	0
int-evi-202			0	
172.16.23.0/24	Local	Local	04h25m51s	0
int-evi-203			0	

No. of Routes: 2

*A:PE-4# show router 20 route-table ipv6

IPv6 Route Table (Service: 20)

Dest Prefix[Flags] Next Hop[Interface Name]	Type	Proto	Age Metric	Pref
2001:db8:16::20:0/120 int-evi-202	Local	Local	00h00m50s 0	0
2001:db8:16::23:0/120 int-evi-203	Local	Local	00h00m50s 0	0

No. of Routes: 2

When traffic is sent between CE-11 and CE-41, which are both associated with VPLS 202, the forwarding is done by VPLS and not via the VPRN. The FDB for VPLS 202 on PE-2 is as follows:

*A:PE-2# show service id 202 fdb detail

Forwarding Database, Service 202

ServId	MAC Transport:Tnl-Id	Source-Identifier	Type Age	Last Change
202	00:00:01:00:00:11	sap:lag-1:20	L/90	09/27/19 12:20:43
202	00:00:01:00:00:16	sap:lag-1:20	L/90	09/27/19 12:20:49
202	00:00:04:00:00:41	eMpls:	Evpn	09/27/19 12:20:38
192.0.2.4:524275	ldp:65539			
202	00:00:5e:00:01:01	cpm	Intf	09/27/19 12:20:19
202	00:00:5e:00:02:01	cpm	Intf	09/27/19 12:20:19
202	00:ca:fe:00:02:02	cpm	Intf	09/27/19 12:20:18
202	00:ca:fe:00:02:03	eMpls:	EvpnS:P	09/27/19 12:20:26
		192.0.2.3:524274		
	ldp:65537			
202	00:ca:fe:00:02:04	eMpls:	EvpnS:P	09/27/19 12:20:37
		192.0.2.4:524275		
	ldp:65539			

No. of MAC Entries: 8

Legend: L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf

MAC 00:00:01:00:00:11 corresponds to CE-11 and is learned on SAP lag-1:20 on PE-2 and advertised via an EVPN MAC route to the BGP peers. MAC 00:00:04:00:00:41 corresponds to CE-41 and was advertised via an EVPN MAC route from PE-4, where the MAC was learned on SAP 1/2/1:41 of VPLS 202, as shown in the following FDB:

*A:PE-4# show service id 202 fdb detail

Forwarding Database, Service 202

ServId	MAC Transport:Tnl-Id	Source-Identifier	Type Age	Last Change
--------	-------------------------	-------------------	-------------	-------------


```

202      00:00:01:00:00:11 eES:      Evpn      09/27/19 12:20:04
                01:00:00:00:00:23:00:00:00:01
202      00:00:01:00:00:16 eES:      Evpn      09/27/19 12:20:10
                01:00:00:00:00:23:00:00:00:01
202      00:00:04:00:00:41 sap:1/2/1:41 L/0      09/27/19 12:19:59
202      00:00:5e:00:01:01 cpm      Intf      09/27/19 12:19:58
202      00:00:5e:00:02:01 cpm      Intf      09/27/19 12:19:58
202      00:ca:fe:00:02:02 eMpls:    EvpnS:P   09/27/19 12:19:59
                192.0.2.2:524274
                ldp:65537
202      00:ca:fe:00:02:03 eMpls:    EvpnS:P   09/27/19 12:19:59
                192.0.2.3:524274
                ldp:65539
202      00:ca:fe:00:02:04 cpm      Intf      09/27/19 12:19:58
-----
No. of MAC Entries: 8
-----
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
=====

```

The MAC address of CE-43 is not present in the FDB of VPLS 202. The FDB of VPLS 203 shows the MAC address of CE-43, but not of CE-41. Traffic between these two VPLS services goes via the VPRN and cannot use Layer 2 forwarding.

```

*A:PE-4# show service id 203 fdb detail

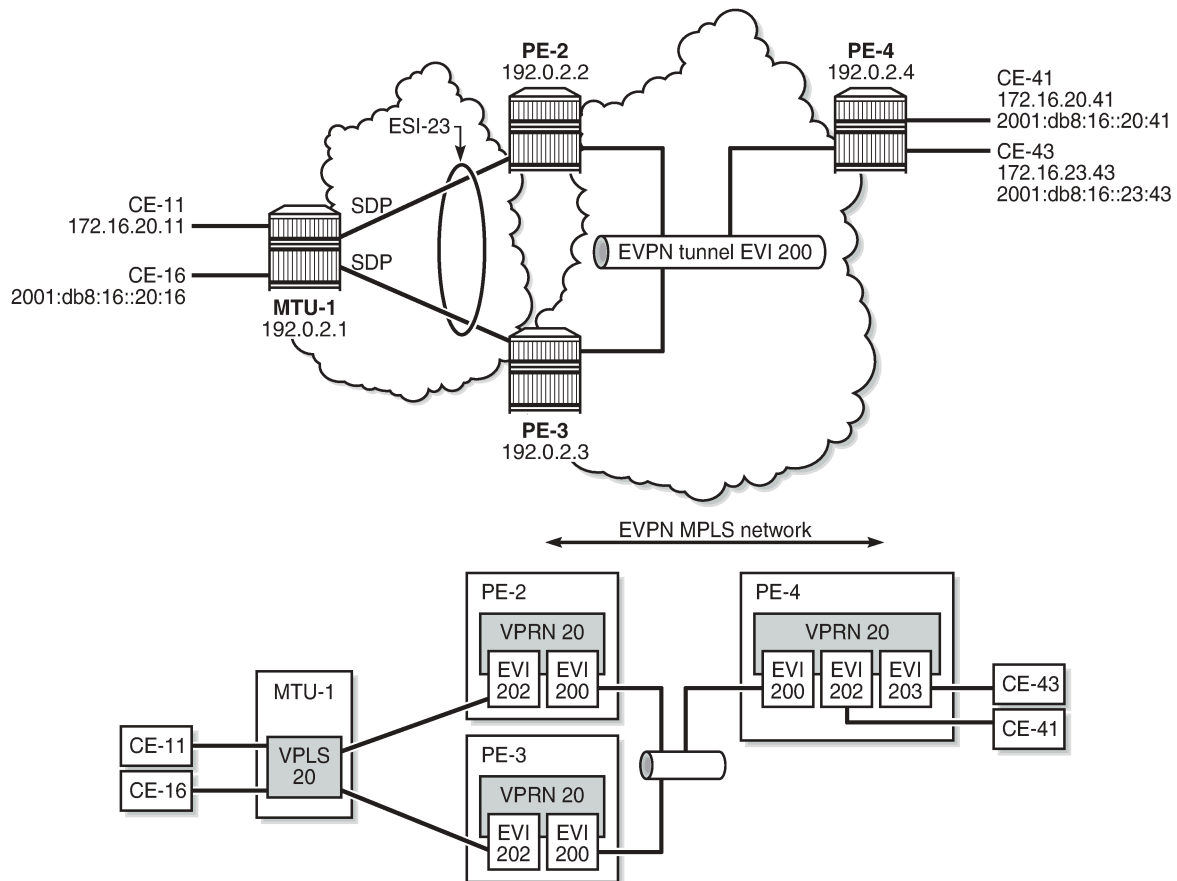
=====
Forwarding Database, Service 203
=====
ServId      MAC              Source-Identifier      Type      Last Change
  Transport:Tnl-Id
-----
203      00:00:04:00:00:43 sap:1/2/1:43      L/0      09/27/19 12:20:32
203      00:00:5e:00:01:02 cpm      Intf      09/27/19 12:20:16
203      00:00:5e:00:02:02 cpm      Intf      09/27/19 12:20:16
203      00:ca:fe:00:23:04 cpm      Intf      09/27/19 12:20:16
-----
No. of MAC Entries: 4
-----
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
=====

```

10.4.3.3 EVPN-MPLS r-VPLS with single-active multihoming

The following figure shows an example of topology with single-active multihoming ES "ESI-23". The difference between this figure and [Figure 177: EVPN-MPLS r-VPLS with all-active multihoming ES](#) is that in this figure, the ES is single-active and SDPs are used instead of a LAG.

Figure 178: EVPN-MPLS r-VPLS with single-active multihoming



26853

The configuration is modified as follows:

- LAG 1 is removed from MTU-1, PE-2, and PE-3
- network interfaces are configured between MTU-1 and PE-2/PE-3 with IS-IS and LDP enabled
- SDPs are configured
- ES "ESI-23" is redefined as single-active multihoming. The SDP is associated with this ES.
- VPLS 202 on PE-2 and PE-3 no longer has a SAP, but has a spoke SDP instead
- no changes are required on VPRN 20 or VPLS 200

The service configuration on PE-2 is as follows. The configuration on PE-3 is similar. No changes are required on PE-4.

```
*A:PE-2# configure service
*A:PE-2>config>service# info
-----
system
  bgp-evpn
    ethernet-segment "ESI-23" create
      esi 01:00:00:00:00:23:00:00:00:01
      es-activation-timer 3
```

```

        service-carving
            mode auto
        exit
        multi-homing single-active
sdp 21
    no shutdown
    exit
    exit
    exit
---snip---
    sdp 21 mpls create
    far-end 192.0.2.1
ldp
    keep-alive
    shutdown
    exit
    no shutdown
    exit
---snip---
    vprn 20 name "VPRN 20" customer 1 create
    route-distinguisher 192.0.2.2:20
    vrf-target target:64500:20
    interface "int-evi-202" create
        address 172.16.20.2/24
        mac 00:ca:fe:00:02:02
        vrrp 1 passive
            backup 172.16.20.254
            ping-reply
            traceroute-reply
        exit
    ipv6
        address 2001:db8:16::20:2/120
        link-local-address fe80::16:20:2 dad-disable
        vrrp 1 passive
            backup fe80::16:20:fe
            ping-reply
            traceroute-reply
        exit
    exit
    vpls "evi-202"
    exit
    exit
    interface "int-evi-200" create
        ipv6
        exit
        vpls "evi-200"
            evpn-tunnel
        exit
    exit
    router-advertisement
        interface "int-evi-202"
            use-virtual-mac
            no shutdown
        exit
    exit
    no shutdown
    exit
    vpls 200 name "evi-200" customer 1 create
        allow-ip-int-bind
        exit
        bgp
        exit
        bgp-evpn
            ip-route-advertisement

```

```

        evi 200
        mpls bgp 1
            auto-bind-tunnel
            resolution any
        exit
        no shutdown
    exit
exit
stp
    shutdown
exit
no shutdown
exit
vpls 202 name "evi-202" customer 1 create
    allow-ip-int-bind
    exit
    bgp
    exit
    bgp-evpn
        evi 202
        mpls bgp 1
            auto-bind-tunnel
            resolution any
        exit
        no shutdown
    exit
exit
stp
    shutdown
exit
spoke-sdp 21:20 create
    no shutdown
exit
no shutdown
exit
```

PE-2 is the designated forwarder (DF) in the single-active ES, as shown in the following output:

```

*A:PE-2# show service id 202 ethernet-segment
No sap entries
=====
SDP Ethernet-Segment Information
=====
SDP                Eth-Seg                Status
-----
21:20              ESI-23                DF
=====

*A:PE-3# show service id 202 ethernet-segment
No sap entries
=====
SDP Ethernet-Segment Information
=====
SDP                Eth-Seg                Status
-----
31:20              ESI-23                NDF
=====
```

When traffic is sent between CE-11 and CE-41, the FDB on PE-2 is as follows, where MAC address 00:00:01:00:00:11 corresponds to CE-11 and is learned on spoke SDP 21:20, and MAC address 00:00:04:00:00:41 corresponds to CE-41 and his advertised by PE-4 in an EVPN-MAC route.

```
*A:PE-2# show service id 202 fdb detail

=====
Forwarding Database, Service 202
=====
ServId      MAC              Source-Identifier  Type      Last Change
      Transport:Tnl-Id
-----
202         00:00:01:00:00:11 sdp:21:20         L/30      09/27/19 12:24:05
202         00:00:01:00:00:16 sdp:21:20         L/30      09/27/19 12:24:10
202         00:00:04:00:00:41 eMpls:           Evpn       09/27/19 12:20:38
                        192.0.2.4:524275
                        ldp:65539
202         00:00:5e:00:01:01 cpm              Intf       09/27/19 12:20:19
202         00:00:5e:00:02:01 cpm              Intf       09/27/19 12:20:19
202         00:ca:fe:00:02:02 cpm              Intf       09/27/19 12:20:18
202         00:ca:fe:00:02:03 eMpls:           EvpnS:P    09/27/19 12:20:26
                        192.0.2.3:524274
                        ldp:65537
202         00:ca:fe:00:02:04 eMpls:           EvpnS:P    09/27/19 12:20:37
                        192.0.2.4:524275
                        ldp:65539
-----
No. of MAC Entries: 8
-----
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static Lf=Leaf
=====
```

When the SDP between MTU-1 and DF PE-2 goes down, traffic from CE-41 to CE-11 is forwarded by PE-4 to DF PE-2. PE-2 cannot forward the packets to CE-11 directly and will forward the packets to its ES peer PE-3. PE-3 will forward to CE-11 even if the MAC SA matches its own vMAC. Virtual MACs bypass the r-VPLS interface protection, so traffic can be forwarded between the PEs without being dropped.

10.5 EVPN command reference

10.5.1 Command hierarchies

- EVPN configuration commands
 - Epipe commands for EVPN
 - VPLS commands for EVPN
 - Epipe and VPLS services-related commands for EVPN
 - VPRN interface commands for EVPN
 - Service system commands for EVPN
 - Redundancy commands for EVPN
 - Router BGP commands for EVPN
- Show commands
- Tools Commands (see the Tools section of the 7705 SAR OAM and Diagnostics Guide)

10.5.1.1 EVPN configuration commands

10.5.1.1.1 Epipe commands for EVPN

```

config
- service
- epipe
- [no] bgp
- route-distinguisher rd
- no route-distinguisher
- route-target {ext-community | export ext-community | import ext-community}
- no route-target
- [no] bgp-evpn
- evi value
- no evi
- local-ac-name ac-name
- no local-ac-name
- eth-tag tag-value
- no eth-tag
- mpls
- auto-bind-tunnel
- ecmp max-ecmp-routes
- resolution-filter
- [no] bgp
- [no] ldp
- [no] rsvp
- [no] sr-isis
- [no] sr-ospf
- [no] sr-te
- [no] weighted-ecmp
- [no] control-word
- ecmp max-ecmp-routes
- [no] entropy-label

```

```

- [no] force-vlan-vc-forwarding
- route-next-hop {ip-address | system-ipv4 | system-ipv6}
- [no] shutdown
- remote-ac-name ac-name
- no remote-ac-name
- eth-tag tag-value
- no eth-tag

```

10.5.1.1.2 VPLS commands for EVPN

```

config
- service
- vpls
- [no] bgp
- route-distinguisher rd
- no route-distinguisher
- route-target ext-community
- route-target export ext-community [import ext-community]
- route-target import ext-community
- no route-target
- vsi-export policy-name [policy-name...(up to 5 max)]
- no vsi-export
- vsi-import policy-name [policy-name...(up to 5 max)]
- no vsi-import
- [no] bgp-evpn
- [no] cfm-mac-advertisement
- evi value
- [no] evi
- incl-mcast-orig-ip ip-address
- no incl-mcast-orig-ip
- [no] ingress-repl-inc-mcast-advertisement
- ip-route-advertisement [incl-host]
- no ip-route-advertisement
- [no] mac-advertisement
- mac-duplication
- detect num-moves num-moves window minutes
- retry minutes
- [no] retry
- mpls
- auto-bind-tunnel
- ecmp max-ecmp-routes
- resolution {disabled | any | filter}
- resolution-filter
- [no] bgp
- [no] ldp
- [no] rsvp
- [no] sr-isis
- [no] sr-ospf
- [no] sr-te
- [no] weighted-ecmp
- [no] control-word
- ecmp max-ecmp-routes
- [no] entropy-label
- [no] force-vlan-vc-forwarding
- [no] ingress-replication-bum-label
- route-next-hop {ip-address | system-ipv4 | system-ipv6}
- [no] shutdown
- split-horizon-group name
- no split-horizon-group
- static-mac
- mac ieee-address [create] black-hole

```

```

- mac ieee-address [create] sap sap-id monitor {fwd-status}
- mac ieee-address [create] spoke-sdp sdp-id:vc-id monitor {fwd-status}
- no mac ieee-address

```

10.5.1.1.3 Epipe and VPLS services-related commands for EVPN

```

config
- service
  - epipe      (see Epipe service configuration commands)
  - vpls      (see VPLS command reference)

```

10.5.1.1.4 VPRN interface commands for EVPN

```

config
- service
  - vprn
    - interface
      - vpls
        - [no] evpn-tunnel

```

10.5.1.1.5 Service system commands for EVPN

```

config
- service
  - system
    - [no] bgp-evpn
      - ad-per-es-route-target evi-rt
      - ad-per-es-route-target evi-rt-set route-distinguisher ip-address
      - ethernet-segment name [create]
      - no ethernet-segment name
      - es-activation-timer seconds
      - no es-activation-timer
      - es-orig-ip {ip-address | ipv6-address}
      - no es-orig-ip
      - esi esi
      - no esi
      - lag lag-id
      - no lag
      - multi-homing single-active [no-esi-label]
      - multi-homing all-active
      - no multi-homing
      - port port-id
      - no port
      - route-next-hop {ip-address | ipv6-address}
      - no route-next-hop
      - sdp sdp-id
      - no sdp
      - service-carving
        - manual
          - evi start [to to]
          - no evi start
          - preference [create] [non-revertive]
          - no preference
            - value value
        - mode {manual | auto | off}

```



```

- [no] shutdown
- route-distinguisher rd
- no route-distinguisher

```

10.5.1.1.6 Redundancy commands for EVPN

```

config
- redundancy
- bgp-evpn-multi-homing
- boot-timer seconds
- es-activation-timer seconds

```

10.5.1.1.7 Router BGP commands for EVPN

See the 7705 SAR Routing Protocols Guide for router BGP command descriptions.

```

config
- router
- [no] bgp
- def-recv-evpn-encap mpls
- family [evpn]
- group name
- no group
- def-recv-evpn-encap mpls
- family [evpn]
- neighbor ip-address
- def-recv-evpn-encap mpls
- family [evpn]
- rapid-update [mvpn-ipv4] [evpn]
- no rapid-update

```

10.5.1.2 Show commands

```

show
- service
- evpn-mpls [tep-ip-address]
- id service-id
- bgp
- bgp-evpn
- evpn-mpls
- evpn-mpls esi esi
- ethernet-segment [ethernet-segment-name]
- system
- bgp-evpn
- bgp-evpn ethernet-segment
- bgp-evpn ethernet-segment name name [all]
- bgp-evpn ethernet-segment name name evi [evi]
- bgp-route-distinguisher [vpls] [epipe]
- bgp-route-distinguisher svc
- bgp-route-distinguisher ad-evi-rt-set
- bgp-route-distinguisher system

```

```

show
- redundancy

```

- [bgp-evpn-multi-homing](#)

10.5.2 Command descriptions

- [EVPN configuration commands](#)
 - [BGP commands](#)
 - [BGP-EVPN commands](#)
- [Routed VPLS EVPN commands](#)
- [EVPN service system commands](#)
- [EVPN redundancy commands](#)
- [Show commands](#)

10.5.2.1 EVPN configuration commands

10.5.2.1.1 BGP commands

bgp

Syntax

[no] bgp

Context

config>service>epipe

config>service>vpls

Description

This command enables the context to configure the BGP-related parameters for a BGP Epipe or VPLS.

The **no** version of the command removes the BGP instance.

route-distinguisher

Syntax

route-distinguisher *rd*

no route-distinguisher

Context

config>service>epipe>bgp

config>service>vpls>bgp

Description

This command configures the route distinguisher (RD) component that is signaled in the MP-BGP NLRI for Layer 2 VPN and EVPN families. This value is used for the BGP Epipe NLRI and BGP VPLS NLRI when this command is configured.

Alternatively, for BGP-EVPN-enabled Epipe and VPLS services, the **route-distinguisher** value can be auto-derived from the **evi evi** value (**config>service>epipe>bgp-evpn>evi** or **config>service>vpls>bgp-evpn>evi**) if this command is not configured.

Default

no route-distinguisher

Parameters

rd
specifies the route distinguisher address

Values	<i>ip-addr:comm-val 2byte-asnumber:ext-comm-val 4byte-asnumber:comm-val</i>	
	<i>ip-addr:</i>	a.b.c.d
	<i>comm-val:</i>	0 to 65535
	<i>2byte-asnumber:</i>	1 to 65535
	<i>ext-comm-val:</i>	0 to 4294967295
	<i>4byte-asnumber:</i>	1 to 4294967295

route-target

Syntax

For Epipe:
route-target {*ext-community* | **export** *ext-community* | **import** *ext-community*}
no route-target

For VPLS:
route-target *ext-community*
route-target export *ext-community* [**import** *ext-community*]
route-target import *ext-community*
no route-target

Context

config>service>epipe>bgp
config>service>vpls>bgp

Description

This command configures the route target (RT) component that is signaled in the related MP-BGP attribute that is used for the BGP Epipe, BGP VPLS, and EVPN when this command is configured in this Epipe or VPLS service.

If this command is not used, the RT is formed automatically using the Epipe or VPLS ID. The extended community can have the same two formats as the Epipe or VPLS ID, that is, a two-octet AS-specific extended community or an IPv4-specific extended community. The extended community can also have a four-octet AS-specific community format. For BGP-EVPN-enabled Epipe and VPLS services, the route target can also be auto-derived from the **evi value** (**config>service>vpls>bgp-evpn>evi** or **config>service>epipe>bgp-evpn>evi**) if this command is not configured.

Default

no route-target

Parameters

- export ext-community

specifies extended communities allowed to be sent to remote PE neighbors
- import ext-community

specifies extended communities allowed to be accepted from remote PE neighbors
- ext-community

specifies the extended community
- Values

target:{ip-addr:comm-val | 2byte-asnumber:ext-comm-val | 4byte-asnumber:comm-val}
- ip-addr:

a.b.c.d
- comm-val:

0 to 65535
- 2byte-asnumber:

1 to 65535
- ext-comm-val:

0 to 4294967295
- 4byte-asnumber:

1 to 4294967295

vsi-export

Syntax

- vsi-export policy-name [policy-name ... (up to 5 max)]
- no vsi-export

Context

config>service>vpls>bgp

Description

This command specifies the name of the virtual switch instance (VSI) export policies to be used for BGP VPLS when this command is configured in this VPLS service. If multiple policy names are configured, the policies are evaluated in the order in which they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

Default

no vsi-export

Parameters

policy-name

specifies a VSI export policy (32 characters maximum)

vsi-import

Syntax

vsi-import *policy-name* [*policy-name* ... (up to 5 max)]

no vsi-import

Context

config>service>vpls>bgp

Description

This command specifies the name of the virtual switch instance (VSI) import policies to be used for BGP VPLS when this command is configured in this VPLS service. If multiple policy names are configured, the policies are evaluated in the order in which they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

Default

no vsi-import

Parameters

policy-name

specifies a VSI import policy (32 characters maximum)

10.5.2.1.2 BGP-EVPN commands

bgp-evpn

Syntax

[no] bgp-evpn

Context

```
config>service>epipe
```

```
config>service>vpls
```

Description

This command enables the context to configure the BGP-EVPN parameters in the specified service.

cfm-mac-advertisement**Syntax**

```
[no] cfm-mac-advertisement
```

Context

```
config>service>vpls>bgp-evpn
```

Description

This command enables the advertisement and withdrawal, as appropriate, of the IEEE MAC address associated with the MEP created on a SAP in an EVPN service.

The update occurs each time a MEP is added or deleted or an IEEE MAC address is changed for an MEP on a SAP within the service. The size of the update depends on the number of MEPs in the service affected by the modification.

Only enable this functionality for services that require a resident MAC address to properly forward unicast traffic and that do not perform Layer 2 MAC learning as part of the data plane.

Local MEP IEEE MAC addresses are not stored in the local FDB and therefore cannot be advertised through a control plane to a peer without this command.

The **no** version of the command disables the functionality and withdraws all previously advertised MEP IEEE MAC addresses.

Default

```
no cfm-mac-advertisement
```

evi**Syntax**

```
evi value
```

```
[no] evi
```

Context

```
config>service>epipe>bgp-evpn
```

```
config>service>vpls>bgp-evpn
```

Description

This command specifies a 2-byte EVPN instance (EVI) that is unique in the system. It is used for the service-carving algorithm for multihoming and for auto-deriving route targets and route distinguishers.

If not specified, the **evi value** is 0 and there are no route distinguishers or route targets auto-derived from the value. If the **evi value** is specified and no other route distinguishers or route targets are configured in the service, the following rules apply:

- the route distinguisher is derived from *system-ip:evi value*
- the route target is derived from *autonomous-system:evi value*

If VSI import and VSI export policies are configured, the route target must be configured in the policies and the policy values take preference over the auto-derived route targets. The operational route target for a service is shown using the **show>service>id>bgp** command.

The **no** version of the command sets the **evi value** back to 0.

Default

no evi

Parameters

value

specifies the EVPN instance

Values 1 to 65535

incl-mcast-orig-ip

Syntax

incl-mcast-orig-ip *ip-address*

no incl-mcast-orig-ip

Context

config>service>vpls>bgp-evpn

Description

This command specifies that the IP address configured with the command is encoded in the **originating-ip** field of EVPN Inclusive Multicast Routes with tunnel type Ingress Replication (IR) (value 6).

The configured address does not need to be reachable in the base router or have an interface in the base router. The originating IP address is used solely for BGP route-key selection.

The **no** version of the command withdraws the affected Inclusive Multicast Routes and readvertises the routes with the default system IP address in the originating IP field.

Default

system IP address

Parameters

ip-address

specifies the IPv4 address

ingress-repl-inc-mcast-advertisement

Syntax

[no] ingress-repl-inc-mcast-advertisement

Context

config>service>vpls>bgp-evpn

Description

This command enables the advertisement of the inclusive multicast Ethernet tag route with tunnel type ingress-replication in the PMSI tunnel attribute.

Default

ingress-repl-inc-mcast-advertisement

ip-route-advertisement

Syntax

ip-route-advertisement [incl-host]

no ip-route-advertisement

Context

config>service>vpls>bgp-evpn

Description

This command enables or disables the advertisement of IP prefixes in EVPN. When enabled, any active route in the r-VPLS VPRN route table will be advertised in EVPN using the VPLS BGP configuration. The interface host addresses are not advertised in EVPN unless the **incl-host** parameter is specified.

Default

no ip-route-advertisement

Parameters

incl-host

specifies to advertise the interface host addresses in EVPN

local-ac-name

Syntax

local-ac-name *ac-name*

no local-ac-name

Context

config>service>epipe>bgp-evpn

Description

This command enables the context and specifies the attachment circuit name in which the local Ethernet tag value is configured.

Default

no local-ac-name

Parameters

ac-name

specifies the name of the local attachment circuit

eth-tag

Syntax

eth-tag *tag-value*

no eth-tag

Context

config>service>epipe>bgp-evpn>local-ac-name

config>service>epipe>bgp-evpn>remote-ac-name

Description

This command configures the Ethernet tag value. When configured in the **local-ac-name** context, the system uses the value in the advertised AD per-EVI route sent for the attachment circuit. When configured in the **remote-ac-name context**, the system compares that value with the **eth-tag tag-value** of the imported AD per-EVI routes for the service. If there is a match, the system creates an EVPN-MPLS destination for the Epipe.

Default

n/a

Parameters

tag-value

specifies the Ethernet tag value of the attachment circuit

Values 1 to 16777215

mac-advertisement

Syntax

[no] **mac-advertisement**

Context

config>service>vpls>bgp-evpn

Description

This command enables the advertisement in BGP of the learned MAC addresses on SAPs and SDP bindings. When the **mac-advertisement** command is disabled, the local MAC addresses are withdrawn in BGP.

Default

mac-advertisement

mac-duplication

Syntax

mac-duplication

Context

config>service>vpls>bgp-evpn

Description

This command enables the context to configure the BGP-EVPN MAC duplication parameters.

detect

Syntax

detect num-moves *num-moves* **window** *minutes*

Context

config>service>vpls>bgp-evpn>mac-duplication

Description

This command modifies the default behavior of the **mac-duplication** command. The **mac-duplication** feature is always enabled by default and monitors the number of moves of a MAC address for a period of time (**window**).

Default

num-moves 5 window 3

Parameters

num-moves

specifies the number of MAC moves in a VPLS service. The counter is incremented when a specified MAC is locally relearned in the FDB or flushed from the FDB due to the reception of a better remote EVPN route for that MAC.

Values 3 to 10

Default 3

minutes

specifies the duration of the window

Values 1 to 15

Default 3

retry

Syntax

retry *minutes*

no retry

Context

config>service>vpls>bgp-evpn>mac-duplication

Description

This command specifies the timer after which the MAC in hold-down state is automatically flushed and the MAC duplication process starts again. This value is expected to be equal to or greater than two times that of the **window** *minutes*.

The **no** form of the command implies that when MAC duplication is detected, MAC updates for that MAC will be held down until the user intervenes or a network event that flushes the MAC occurs.

Default

9

Parameters

minutes

specifies the BGP EVPN MAC duplication retry

Values 2 to 60

mpls

Syntax

mpls

Context

config>service>epipe>bgp-evpn

config>service>vpls>bgp-evpn

Description

This command enables the context to configure the BGP-EVPN MPLS parameters.

auto-bind-tunnel

Syntax

auto-bind-tunnel

Context

config>service>epipe>bgp-evpn>mpls

config>service>vpls>bgp-evpn>mpls

Description

This command enables the context to configure automatic binding of a BGP-EVPN MPLS service to tunnels to MP-BGP peers.

ecmp

Syntax

ecmp *max-ecmp-routes*

Context

config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel

config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel

Description

This command enables ECMP for Layer 2 unicast traffic on Epipe or VPLS services for EVPN-MPLS destinations. The command allows the resolution of an EVPN-MPLS next hop to a group of ECMP tunnels of type RSVP-TE or SR-TE and determines the number of Traffic Engineering (TE) tunnels that the EVPN-MPLS next hop can be resolved to. For shortest path tunnels, such as LDP, SR-ISIS, and SR-OSPF, the number of tunnels in the ECMP group is determined by the **config>router>ecmp** command.

Default

1

Parameters

max-ecmp-routes

specifies the number of TE tunnels that an EVPN-MPLS next hop can be resolved to

Values 1 to 8

resolution

Syntax

resolution {disabled | any | filter}

Context

config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel

config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel

Description

This command configures the resolution mode in the automatic binding of a BGP-EVPN MPLS service to tunnels to MP-BGP peers.

The user must configure the **resolution** command to enable autobind resolution to tunnels in the tunnel table manager (TTM). The following options are available:

- **disabled** (explicitly set): autobinding to the tunnel is removed
- **any**: any supported tunnel type in the EVPN context will be selected according to TTM preference
- **filter**: when **filter** and one or more tunnel types are explicitly specified under the **resolution-filter** command, only those tunnel types are selected according to TTM preference

The user must set the **resolution** command to **filter** in order to activate the list of tunnel types configured under the **resolution-filter** command.

Default

disabled

Parameters

any

enables the binding to any supported tunnel type in a BGP-EVPN MPLS context following TTM preference

disabled

disables the automatic binding of a BGP-EVPN MPLS service to tunnels to MP-BGP peers

filter

enables the binding to the subset of tunnel types configured under **resolution-filter**

resolution-filter

Syntax

resolution-filter

Context

```
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel
```

```
config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel
```

Description

This command enables the context that allows the configuration of the subset of tunnel types that can be used in the resolution of BGP-EVPN routes within the automatic binding of BGP-EVPN MPLS service to tunnels to MP-BGP peers.

The user must set the **resolution** command to **filter** in order to activate the list of tunnel types configured under the **resolution-filter** command.

The following tunnel types are supported in a BGP-EVPN MPLS context in order of preference: RSVP, SR-TE, LDP, SR-ISIS, SR-OSPF, and BGP.

bgp

Syntax

[no] bgp

Context

```
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
```

```
config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
```

Description

This command selects the BGP tunnel type.

The **bgp** command instructs BGP-EVPN to search for a BGP LSP to the address of the BGP next hop. If the user does not enable the BGP tunnel type, inter-area or inter-as prefixes will not be resolved.

ldp

Syntax

[no] ldp

Context

config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Description

This command selects the LDP tunnel type.

The **ldp** command instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next hop.

rsvp

Syntax

[no] rsvp

Context

config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Description

This command selects the RSVP-TE tunnel type.

The **rsvp** command instructs BGP to search for the best metric RSVP-TE LSP to the address of the BGP next hop. This address can correspond to the system interface or to another loopback address used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. If there are multiple RSVP-TE LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel ID.

sr-isis

Syntax

[no] sr-isis

Context

config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Description

This command selects the segment routing (SR) tunnel type programmed by an IS-IS instance in the TTM.

When the **sr-isis** or **sr-ospf** command is enabled, a segment routing (SR) tunnel to the BGP next hop is selected in the TTM from the lowest numbered IS-IS or OSPF instance.

sr-ospf

Syntax

[no] **sr-ospf**

Context

```
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
```

```
config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
```

Description

This command selects the segment routing (SR) tunnel type programmed by an OSPF instance in the TTM.

When the **sr-isis** or **sr-ospf** command is enabled, a segment routing (SR) tunnel to the BGP next hop is selected in the TTM from the lowest numbered IS-IS or OSPF instance.

sr-te

Syntax

[no] **sr-te**

Context

```
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
```

```
config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter
```

Description

This command selects the segment routing traffic engineered LSP programmed in the TTM.

The **sr-te** command instructs the code to search for the best metric SR-TE LSP to the address of the BGP next hop. The LSP metric is provided by MPLS in the tunnel table. If there are multiple SR-TE LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel ID.

weighted-ecmp

Syntax

[no] **weighted-ecmp**

Context

```
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel  
config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel
```

Description

This command enables weighted ECMP for packets using tunnels that an Epipe or VPLS automatically binds to. When weighted ECMP is enabled, packets are sprayed across LSPs in the ECMP set according to the outcome of the hash algorithm and the configured **load-balancing-weight** of each LSP. See the 7250 IXR MPLS Guide, "MPLS Commands" for more information about the **load-balancing-weight** command.

The **no** form of this command disables weighted ECMP for next-hop tunnel selection.

Default

no weighted-ecmp

control-word

Syntax

[no] control-word

Context

```
config>service>epipe>bgp-evpn>mpls  
config>service>vpls>bgp-evpn>mpls
```

Description

This command enables the transmission and reception of the control word. As defined in RFC 7432, the use of the control word helps avoid frame disordering.

The **control-word** command is enabled or disabled for all EVPN-MPLS destinations at the same time.

Default

no control-word

ecmp

Syntax

ecmp *max-ecmp-routes*

Context

```
config>service>epipe>bgp-evpn>mpls  
config>service>vpls>bgp-evpn>mpls
```

Description

This command controls the number of paths allowed to reach a specified MAC address when that MAC in the FDB is associated with a remote all-active multihomed Ethernet segment.

The configuration of two or more ECMP paths to a specified MAC enables the aliasing function described in RFC 7432.

Default

0

Parameters

max-ecmp-routes

specifies the number of paths allowed to the same multihomed MAC address, assuming the MAC is located in an all-active multihomed Ethernet segment

Values 0 to 8

entropy-label

Syntax

[no] entropy-label

Context

config>service>epipe>bgp-evpn>mpls

config>service>vpls>bgp-evpn>mpls

Description

This command inserts the entropy label (EL) and entropy label indicator (ELI) in packets for which at least one LSP in the stack for the far end of the tunnel used by the service has advertised entropy-label capability. If the tunnel is the **rsvp** type, **entropy-label** can also be controlled under the **config>router>mpls** or **config>router>mpls>lsp** context.

Default

no entropy-label

force-vlan-vc-forwarding

Syntax

[no] force-vlan-vc-forwarding

Context

config>service>epipe>bgp-evpn>mpls

config>service>vpls>bgp-evpn>mpls

Description

This command allows the system to preserve the VLAN ID and 802.1p bits of the service-delimiting qtag in a new tag added in the customer frame before sending it to the EVPN-MPLS destinations.

This command may be used in conjunction with the **sap ingress vlan-translation** command. If it is used, the configured translated VLAN ID will be the VLAN ID sent to the EVPN-MPLS destinations as opposed to the service-delimiting tag VLAN ID. If the ingress SAP or SDP binding is null-encapsulated, the output VLAN ID and P-bits will be 0.

Default

no force-vlan-vc-forwarding

ingress-replication-bum-label

Syntax

[no] **ingress-replication-bum-label**

Context

config>service>vpls>bgp-evpn>mpls

Description

This command allows the user to configure the system so that a separate label is sent for BMU (broadcast, multicast, and unknown unicast) traffic in a specified service. By default, the same label is used for unicast and flooded BMU packets when forwarding traffic to remote PEs.

When saving label space, this may cause transient packet duplication for all-active multihoming. By enabling **ingress-replication-bum-label**, the system will advertise two labels per EVPN VPLS instance, one for unicast and one for BMU traffic. The ingress PE will use the BMU label for flooded traffic to the advertising egress PE so that the egress PE can determine if the unicast traffic has been flooded by the ingress PE. Depending on the scale required in the network, the user may choose between saving label space or avoiding transient packet duplication sent to an all-active multihomed CE for certain MAC addresses.

Default

no ingress-replication-bum-label

route-next-hop

Syntax

route-next-hop {*ip-address* | **system-ipv4** | **system-ipv6**}

Context

config>service>epipe>bgp-evpn>mpls

config>service>vpls>bgp-evpn>mpls

Description

This command configures the next hop of the EVPN routes.

Default

system-ipv4

Parameters

ip-address

the non-system IP address to use as the next hop for the EVPN routes

system-ipv4

specifies the system IPv4 address as the next hop for the service EVPN

system-ipv6

specifies the system IPv6 address as the next hop for the service EVPN

shutdown

Syntax

[no] shutdown

Context

config>service>epipe>bgp-evpn>mpls

config>service>vpls>bgp-evpn>mpls

Description

This command controls the administrative state of EVPN-MPLS in the service.

split-horizon-group

Syntax

split-horizon-group *name*

no split-horizon-group

Context

config>service>vpls>bgp-evpn>mpls

Description

This command allows the user to configure an explicit split horizon group for all BGP-EVPN MPLS destinations that can be shared by other SAPs and spoke SDPs. The use of explicit split horizon groups for EVPN-MPLS and spoke SDPs allows the integration of VPLS and EVPN-MPLS networks.

If the **split-horizon-group** command under **bgp-evpn>mpls** is not used, the default split horizon group that contains all the EVPN destinations is still used, but it is not possible to refer to it on SAPs or

spoke SDPs. Split horizon groups can be configured within the **service** context. The same group name can be associated with SAPs, spoke SDPs, and EVPN-MPLS destinations. The configuration of **bgp-evpn>mpls>split-horizon-group** will only be allowed if **bgp-evpn>mpls** is **shutdown**; no changes are allowed when **bgp-evpn>mpls** is **no shutdown**.

When the SAPs or spoke SDPs are configured within the same split horizon group as the EVPN-MPLS endpoints, MAC addresses are still learned on them but they are not advertised in BGP-EVPN.

Default

no split-horizon-group

Parameters

name

specifies the split horizon group name

remote-ac-name

Syntax

remote-ac-name *ac-name*

no remote-ac-name

Context

config>service>epipe>bgp-evpn

Description

This command enables the context and specifies the attachment circuit name in which the remote Ethernet tag value is configured.

Default

no remote-ac-name

Parameters

ac-name

specifies the name of the remote attachment circuit

static-mac

Syntax

static-mac

Context

config>service>vpls

Description

This command enables the context that allows the assignment of a set of conditional static MAC addresses to a SAP, spoke SDP, or **black-hole**. In the FDB, the static MAC address is then associated with the active SAP or spoke SDP.

A set of conditional static MAC addresses can be created within a VPLS supporting BGP-EVPN. Unless they are configured as **black-hole**, conditional static MAC addresses are dependent on the SAP or SDP state.

Static MAC addresses configured in a BGP-EVPN service are advertised as protected (EVPN signals the MAC address as protected).

mac

Syntax

mac *ieee-address* [**create**] **black-hole**

mac *ieee-address* [**create**] **sap** *sap-id* **monitor** {**fwd-status**}

mac *ieee-address* [**create**] **spoke-sdp** *sdp-id:vc-id* **monitor** {**fwd-status**}

no mac *ieee-address*

Context

config>service>vpls>static-mac

Description

This command assigns a conditional static MAC address entry to an EVPN VPLS SAP, spoke SDP, or a black hole on the 7705 SAR.

Parameters

ieee-address

specifies the static MAC address

sap-id

specifies the SAP ID

sdp-id

specifies the SDP ID

vc-id

specifies the virtual circuit ID

create

mandatory keyword while creating a static MAC

black-hole

creates a static FDB entry for the MAC address to blackhole traffic

fwd-status

specifies that this static MAC address will be installed in the FDB based on the forwarding status of the SAP or spoke SDP

10.5.2.2 Routed VPLS EVPN commands

```
evpn-tunnel
```

Syntax

```
[no] evpn-tunnel
```

Context

```
config>service>vprn>if>vpls
```

Description

This command enables or disables EVPN tunnel mode for the attached r-VPLS interface. When enabled, no IP address is required for the interface.

The **no** version of the command disables EVPN tunnel mode.

Default

```
no evpn-tunnel
```

10.5.2.3 EVPN service system commands

```
bgp-evpn
```

Syntax

```
[no] bgp-evpn
```

Context

```
config>service>system
```

Description

This command enables the context to configure the BGP-EVPN parameters in the base instance.

```
ad-per-es-route-target
```

Syntax

```
ad-per-es-route-target evi-rt
```

```
ad-per-es-route-target evi-rt-set route-distinguisher ip-address
```

Context

```
config>service>system>bgp-evpn
```


Description

This command controls how Ethernet auto-discovery (AD) per-ES routes are generated.

The system can send either a separate Ethernet AD per-ES route per service or several Ethernet AD per-ES routes aggregating the route targets for multiple services. While the two methods can interoperate, RFC 7432 states that the EVPN AD per-ES route must be sent with a set of route targets corresponding to all the EVIs defined on the Ethernet segment. Either method can be enabled using the **evi-rt** and **evi-rt-set** options.

The default option, **evi-rt**, configures the system to send a separate Ethernet AD per-ES route per service.

The **evi-rt-set** option, when enabled, allows the aggregation of routes; that is, a single AD per-ES route with the associated RD (*ip-address:1*) and a set of EVI route targets are advertised (to a maximum of 128 route targets). If the number of EVIs defined in the Ethernet segment is significantly large for the packet size, the system will send more than one route. For example:

- AD per-ES route for EVI-route-set 1 is sent with RD *ip-address:1*
- AD per-ES route for EVI-route-set 2 is sent with RD *ip-address:2*

Default

ad-per-es-route-target evi-rt

Parameters

evi-rt

specifies the option to advertise a separate AD per-ES route per service

evi-rt-set

specifies the option to advertise a set of AD per-ES routes aggregating the route targets for all the services in the Ethernet segment

ip-address

specifies the IP address part of the route distinguisher (RD) being used in the **evi-rt-set** option

ethernet-segment

Syntax

ethernet-segment *name* [**create**]

no ethernet-segment *name*

Context

config>service>system>bgp-evpn

Description

This command configures an Ethernet segment (ES) instance and its corresponding name.

Default

n/a

Parameters

- name*

specifies the ES name (28 character s maximum)
- create**

mandatory keyword for creating an ES

es-activation-timer

Syntax

- es-activation-timer** *seconds*
- no es-activation-timer**

Context

config>service>system>bgp-evpn>ethernet-segment

Description

This command configures the activation timer for a specified Ethernet segment. This timer delays the activation of the Ethernet segment on a specified PE that has been elected as the designated forwarder (DF). Only when the timer has expired can the SAP associated with an Ethernet segment be activated (for single-active multihoming) or added to the default multicast list (for all-active multihoming).

If this command is not configured, the system uses the value configured in the **config>redundancy>bgp-evpn-multi-homing>es-activation-timer** context, if it has been configured. Otherwise, the system uses the default value of 3 seconds.

Default

no es-activation-timer

Parameters

- seconds*

specifies the delay time for the ES activation timer
- Values**

0 to 100
- Default**

3

es-orig-ip

Syntax

- es-orig-ip** {*ip-address* | *ipv6-address*}
- no es-orig-ip**

Context

```
config>service>system>bgp-evpn>ethernet-segment
```

Description

This command configures the originating IP address advertised in the ES route. By default, the originating IP address is the system IP of the PE. However, this value can be changed to the IPv4 or IPv6 address configured with this command.

When the **es-orig-ip** is configured, an ES shutdown is required when adding or advertising local ES routes.

When adding local ES routes, this command changes how the ES routes are added to the candidate list; the configured IP address is added, instead of the system IP address.

When advertising local ES routes, the configured IP address is used for the orig-ip of the route.

The **no** form of this command reverts the originating IP address to the system IP address.

Default

```
no es-orig-ip
```

Parameters

ip-address

specifies an IPv4 originating address

ipv6-address

specifies an IPv6 originating address

esi

Syntax

```
esi esi
```

```
no esi
```

Context

```
config>service>system>bgp-evpn>ethernet-segment
```

Description

This command configures the 10-byte Ethernet segment identifier (ESI) associated with the Ethernet segment that is signaled in the BGP-EVPN routes. The *esi* value cannot be changed unless the Ethernet segment is **shutdown**. Reserved *esi* values (ESI-0 and ESI-MAX) are not allowed.

Default

```
no esi
```

Parameters

esi

specifies the 10-byte ESI

Values `xx-xx-xx-xx-xx-xx-xx-xx-xx-xx`, where `xx` is a hexadecimal number (00 to ff) and the separator is a dash ("-"), colon (":"), or space (" ")
For example, 00-11-22-33-44-55-66-77-88-99

lag

Syntax

lag *lag-id*

no lag

Context

config>service>system>bgp-evpn>ethernet-segment

Description

This command configures a LAG ID associated with the Ethernet segment. If the Ethernet segment is configured as **all-active**, only a LAG can be associated with the Ethernet segment. If the Ethernet segment is configured as **single-active**, a LAG, port, or SDP can be associated with the Ethernet segment. A specified LAG can be part of only one Ethernet segment.

Default

no lag

Parameters

lag-id

specifies the LAG ID associated with the Ethernet segment

Values 1 to 800

multi-homing

Syntax

multi-homing single-active [**no-esi-label**]

multi-homing all-active

no multi-homing

Context

config>service>system>bgp-evpn>ethernet-segment

Description

This command configures the multihoming mode for the Ethernet segment as single-active or all-active multihoming, as defined in RFC 7432.

By default, the use of an ESI label is enabled for **all-active** and **single-active** as defined in RFC 7432 (for single-active multihoming, the ESI label is used to avoid transient loops).

When **single-active no-esi-label** is specified, the system does not allocate a label for the ESI and advertises ESI label 0 to peers. Even if the ESI is configured not to send the ESI label, upon receiving an ESI label from a peer, the PE always sends traffic to that peer using the received ESI label.

Default

no multi-homing

Parameters

single-active

configures single-active mode for the Ethernet segment

all-active

configures single-active mode for the Ethernet segment

no-esi-label

configures single-active mode without adding an ESI label for the Ethernet segment

port

Syntax

port port-id

no port

Context

config>service>system>bgp-evpn>ethernet-segment

Description

This command configures a port ID associated with the Ethernet segment. If the Ethernet segment is configured as **all-active**, only a LAG can be associated with the Ethernet segment. If the Ethernet segment is configured as **single-active**, a LAG, port, or SDP can be associated with the Ethernet segment. A specified port can be part of only one Ethernet segment. Only Ethernet ports can be added to an Ethernet segment.

Default

no port

Parameters

port-id

specifies the port ID associated with the Ethernet segment

port-id	slot/mda/port
	mw-link-id mw-link-link-num

link-num

1 to 24

route-next-hop

Syntax

route-next-hop {*ip-address* | *ipv6-address*}**no route-next-hop**

Context

config>service>system>bgp-evpn>eth-seg

Description

This command configures the next hop of the EVPN routes.

Default

no route-next-hop

Parameters

ip-address

the IPv4 address to use as the next hop for the EVPN routes

ipv6-address

the IPv6 address to use as the next hop for the EVPN routes

sdp

Syntax

sdp *sdp-id***no sdp**

Context

config>service>system>bgp-evpn>ethernet-segment

Description

This command configures an SDP ID associated with the Ethernet segment. If the Ethernet segment is configured as **all-active**, only a LAG can be associated with the Ethernet segment. If the Ethernet segment is configured as **single-active**, a LAG, port, or SDP can be associated with the Ethernet segment. A specified SDP can be part of only one Ethernet segment. Only user-configured SDPs can be added to an Ethernet segment.

Default

no sdp

Parameters

sdp-id

specifies the SDP identifier

Values 1 to 17407

service-carving

Syntax

service-carving

Context

config>service>system>bgp-evpn>ethernet-segment

Description

This command enables the context to configure service carving in an Ethernet segment. The service carving algorithm determines which PE is the designated forwarder (DF) in a specific ES and for a specific service.

manual

Syntax

manual

Context

config>service>system>bgp-evpn>ethernet-segment>service-carving

Description

This command enables the context to manually configure the service-carving algorithm, that is, to configure the EVIs for which the PE is the DF and configure the DF preference election information.

There are two service-carving manual algorithms for DF election:

- manual non-preference

The **preference** command is not configured for this algorithm. The PE is the DF for the EVIs specified with the **evi** command. The manual non-preference algorithm only supports two PEs in the Ethernet segment.

- manual preference-based

When the **preference** command is configured, the algorithm uses the configured value to determine the DF election. The highest-preference algorithm is used for EVIs that are not defined with the **evi** command. The lowest-preference algorithm is used for EVIs that are explicitly defined with the **evi** command. The preference-based DF election algorithm supports up to four PEs. If there are more than four PEs using the preference-based algorithm for an ESI and EVI, the ES candidate list for the DF preference election will be pruned to the four lowest PE IP addresses.

evi

Syntax

evi start [*to to*]

no evi start

Context

config>service>system>bgp-evpn>ethernet-segment>service-carving>manual

Description

This command configures the EVI values for which the PE is primary (DF) or the EVI values that use the lowest-preference algorithm.

When the service-carving manual non-preference algorithm is used, the configured EVI values indicate the EVIs for which the PE is the DF.

When the service-carving manual preference-based algorithm is used, the configured EVI values indicate the EVIs that are using the lowest-preference algorithm whereas the undefined EVIs use the highest-preference algorithm.



Note: Multiple individual **evi** values and ranges can be configured. A separate **evi** command must be used for each configuration.

Default

n/a

Parameters

start

specifies the initial **evi value** of the range

Values 1 to 65535

to

specifies the end **evi value** of the range. If not configured, only the individual start value is considered.

Values 1 to 65535

preference

Syntax

preference [*create*] [*non-revertive*]

no preference

Context

```
config>service>system>bgp-evpn>ethernet-segment>service-carving>manual
```

Description

This command creates the preference context for the ES. The preference context ensures that the PE will run the preference-based DF election algorithm. The **non-revertive** option ensures that when a former DF PE comes back up after a failure, it does not take over from an existing active DF PE.

Default

no preference

Parameters

create

mandatory keyword to create the preference context in an ES

non-revertive

configures a non-revertive ES

value

Syntax

value *value*

Context

```
config>service>system>bgp-evpn>ethernet-segment>service-carving>manual>preference
```

Description

This command modifies the default preference value used for the PE in the ES. An ES shutdown is not required to modify this value.

The value determines the DF PE when the preference-based DF election algorithm is run for each service EVI. The preferred DF is the PE with the highest or lowest preference value, depending on whether the EVIs are configured. By default, the highest-preference algorithm is used to elect a DF. If an EVI value or range of values is explicitly configured, the lowest-preference algorithm is used to elect a DF for those EVIs. If the **value** is not explicitly configured, the default value is used when the DF preference extended community attribute is advertised.

Default

32767

Parameters

value

the preference value used in the preference-based DF election algorithm

Values 0 to 65535

mode

Syntax

mode {**manual** | **auto** | **off**}

Context

config>service>system>bgp-evpn>ethernet-segment>service-carving

Description

This command configures the service-carving mode. This determines how the DF is elected for a specified Ethernet segment and service.

Default

mode auto

Parameters

auto

specifies that the service-carving algorithm is as defined in RFC 7432. The DF for the service is calculated based on the modulo function of the service (identified by the **evi value**) and the number of PEs.

manual

specifies that the DF is elected based on the configuration of the **manual** command in the **service-carving>manual** context

off

specifies that all the services elect the same DF PE, assuming that the same PEs are active for all the configured services. The PE with the lowest IP address is elected as DF for the Ethernet segment.

shutdown

Syntax

[no] **shutdown**

Context

config>service>system>bgp-evpn>ethernet-segment

Description

This command changes the administrative status of the Ethernet segment.

The user can configure **no shutdown** only when the **esi**, **multi-homing**, and **lag**, **port**, or **sdp** commands are configured. If the Ethernet segment or the corresponding **lag**, **port**, or **sdp** is **shutdown**, the Ethernet segment route and the AD per-ES routes will be withdrawn. No changes are allowed when the Ethernet segment is **no shutdown**.

Default

shutdown

route-distinguisher

Syntax

route-distinguisher *rd*
no route-distinguisher

Context

config>service>system>bgp-evpn

Description

This command configures the route distinguisher (RD) component that is signaled in the MP-BGP NLRI for EVPN corresponding to the base EVPN instance (route type 4, Ethernet Segment routes). If the route distinguisher component is not configured, the system uses the system IP address as the default route distinguisher for the IP address portion of the RD.

Default

system-ip:0

Parameters

rd
specifies the IP address
Values *ip-addr:* a.b.c.d
 comm-val: 0 to 65535

10.5.2.4 EVPN redundancy commands

redundancy

Syntax

redundancy

Context

config

Description

This command enables the context to configure the global redundancy parameters.

bgp-evpn-multi-homing

Syntax

bgp-evpn-multi-homing

Context

config>redundancy

Description

This command enables the context to configure the BGP-EVPN global timers.

boot-timer

Syntax

boot-timer *seconds*

Context

config>redundancy>bgp-evpn-multi-homing

Description

This command allows the necessary time during PE boot-up for the control plane protocols to come up before bringing up the Ethernet segments and running the DF algorithm.

The following considerations apply to the functionality.

- The configured **boot-timer** value must provide enough time to allow the IOM and BGP sessions to come up before exchanging ES routes and running the DF election for each EVI.
- The boot timer is synchronized across CSMs and is relative to the system up time; therefore, it is not subject to change or reset upon CSM switchover.
- The boot timer is never interrupted. However, the **es-activation-timer** can be interrupted if a new event triggers the DF election.
- The boot timer runs per EVI on the ESs in the system. While the timer is running, the system does not run the DF election for any EVI. When the boot timer expires, the DF election for the EVI is run, and if the system is elected DF for the EVI, the **es-activation-timer** activates.
- The system does not advertise ES routes until the boot timer has expired. This guarantees that the peer ES PEs do not run the DF election until the PE is ready to become the DF, if required.

Default

boot-timer 10

Parameters

seconds

specifies the delay time used for the boot timer

Values 0 to 600

es-activation-timer

Syntax

es-activation-timer *seconds*

Context

config>redundancy>bgp-evpn-multi-homing

Description

This command configures the global Ethernet segment activation timer. The timer delays the activation of an Ethernet segment on a specified PE that has been elected as DF. Only when the **es-activation-timer** has expired can the SAP or SDP binding associated with an Ethernet segment be activated (for single-active multihoming) or added to the default multicast list (for all-active multihoming).

The **es-activation-timer** configured at the **ethernet-segment** level supersedes this global **es-activation-timer**.

Default

es-activation-timer 3

Parameters

seconds

specifies the delay time for the ES activation timer

Values 0 to 100

10.5.2.5 Show commands

evpn-mpls

Syntax

evpn-mpls [*tep-ip-address*]

Context

show>service

Description

This command displays the remote EVPN-MPLS tunnel endpoints in the system.

Parameters

tep-ip-address
specifies the IP address of a tunnel endpoint

Output

The following output is an example of EVPN-MPLS tunnel endpoint information, and [Table 211: Service EVPN-MPLS field descriptions](#) describes the fields.

Output example

```
*A:PE70(4)# show service evpn-mpls
=====
EVPN MPLS Tunnel Endpoints
=====
EvpnMplsTEP Address EVPN-MPLS Dest      ES Dest      ES BMac Dest
-----
192.0.2.69      3          1          1
192.0.2.71      2          0          0
192.0.2.72      3          1          1
192.0.2.73      2          1          0
192.0.2.254     1          0          0
-----
Number of EvpnMpls Tunnel Endpoints: 5
=====

*A:PE70(4)# show service evpn-mpls 192.0.2.69
=====
BGP EVPN-MPLS Dest
=====
Service Id      Egr Label
-----
1               262140
1               262141
20000           262138
-----

=====
BGP EVPN-MPLS Ethernet Segment Dest
=====
Service Id      Eth Seg Id      Egr Label
-----
1               01:00:00:00:00:71:00:00:00:01  262141
-----

=====
BGP EVPN-MPLS ES BMac Dest
=====
Service Id      ES BMac      Egr Label
-----
20000           00:00:00:00:71:71  262138
-----
=====
```

Table 211: Service EVPN-MPLS field descriptions

Label	Description
EVPN MPLS Tunnel Endpoints	
EvpnMplsTEP Address	The IP address of the remote EVPN MPLS tunnel endpoint
EVPN-MPLS Dest	The number of EVPN MPLS destinations
ES Dest	The number of Ethernet segment destinations
ES BMac Dest	Not applicable
BGP EVPN-MPLS Dest	
Service Id	The local service ID for the specified EVPN MPLS tunnel endpoint
Egr Label	The egress label for the specified EVPN MPLS tunnel endpoint
BGP EVPN-MPLS Ethernet Segment Dest	
Service Id	The local service ID for the specified EVPN MPLS Ethernet segment destination
Eth Seg Id	The Ethernet segment ID for the specified EVPN MPLS Ethernet segment destination
Egr Label	The egress label for the specified EVPN MPLS Ethernet segment destination
BGP EVPN-MPLS ES BMac Dest	
Service Id	Not applicable
ES BMac	Not applicable
Egr Label	Not applicable

bgp

Syntax

bgp

Context

show>service>id

Description

This command displays all the information for a specified BGP instance in a service.

Output

The following output is an example of BGP information in a service, and [Table 212: Service ID BGP field descriptions](#) describes the fields.

Output example

```
*A:PE-1# show service id 7000 bgp
=====
BGP Information
=====
Vsi-Import      : None
Vsi-Export      : None
Route Dist      : 1:1
Oper Route Dist : 1:1
Oper RD Type    : configured
Rte-Target Import : None           Rte-Target Export: None
Oper RT Imp Origin : derivedEvi    Oper RT Import  : 64500:7000
Oper RT Exp Origin : derivedEvi    Oper RT Export  : 64500:7000
PW-Template Id   : None
-----
=====
```

Table 212: Service ID BGP field descriptions

Label	Description
BGP Information	
Vsi-Import	The name of the virtual switch instance import policy for the specified service
Vsi-Export	The name of the virtual switch instance export policy for the specified service
Route Dist	The route distinguisher identifier for the specified service
Oper Route Dist	The operational route distinguisher identifier for the specified service
Oper RD Type	The type of operational route distinguisher for the specified service
Rte-Target Import	The route target component that is signaled in the related MP-BGP attribute, specifying communities allowed to be accepted from remote PE neighbors
Rte-Target Export	The route target component that is signaled in the related MP-BGP attribute, specifying communities allowed to be sent to remote PE neighbors
Oper RT Imp Origin	The operational route target component specifying the originating community accepted from remote PE neighbors

Label	Description
Oper RT Import	The operational route target component that is signaled in the related MP-BGP attribute, specifying communities allowed to be accepted from remote PE neighbors
Oper RT Exp Origin	The operational route target component specifying the originating community sent to remote PE neighbors
Oper RT Export	The operational route target component that is signaled in the related MP-BGP attribute, specifying communities allowed to be sent to remote PE neighbors
PW-Template Id	Not applicable

bgp-evpn

Syntax

bgp-evpn

Context

show>service>id

Description

This command displays the BGP-EVPN parameters configured for a specified service, including the configuration of the **mac-advertisement** command, as well as the **mac-duplication** parameters. The command shows the duplicate MAC addresses that **mac-duplication** has detected.

This command also shows whether the **ip-route-advertisement** command (and the **incl-host** parameter) is enabled. If the service is BGP-EVPN MPLS, the command also shows the parameters corresponding to EVPN-MPLS.

Output

The following output is an example of service BGP-EVPN information, and [Table 213: Service BGP-EVPN field descriptions](#) describes the fields.

Output example

```
*A:DutA# show service id 1 bgp-evpn
=====
BGP EVPN Table
=====
MAC Advertisement   : Enabled           Unknown MAC Route   : Disabled
CFM MAC Advertise   : Enabled
VXLAN Admin Status  : Disabled           Creation Origin      : manual
MAC Dup Detn Moves   : 3                  MAC Dup Detn Window : 3
MAC Dup Detn Retry   : 9                  Number of Dup MACs   : 0
IP Route Advertise*  : Disabled

EVI                  : 1
Ing Rep Inc McastAd  : Enabled
```

```

-----
Detected Duplicate MAC Addresses          Time Detected
-----
=====
* indicates that the corresponding row element may have been truncated.
=====

BGP EVPN MPLS Information
=====
Admin Status      : Enabled
Force Vlan Fwding : Disabled          Control Word      : Disabled
Route NextHop Type: system-ipv6
Split Horizon Group: (Not Specified)
Ingress Rep BUM Lbl: Disabled          Max Ecmp Routes    : 4
Ingress Ucast Lbl : 262142             Ingress Mcast Lbl  : 262142
Entropy Label     : Disabled
=====

BGP EVPN MPLS Auto Bind Tunnel Information
=====
Resolution        : any
Filter Tunnel Types: (Not Specified)
=====

```

Table 213: Service BGP-EVPN field descriptions

Label	Description
BGP EVPN Table	
MAC Advertisement	The state of MAC advertising for the service
Unknown MAC Route	Not applicable
CFM MAC Advertise	The state of the IEEE MAC address associated with the MEP created on a SAP in an EVPN service
VXLAN Admin Status	Not applicable
Creation Origin	Indicates whether the route creation was manual or automatic
MAC Dup Detn Moves	The number of detected moves of a MAC address for the period of time defined by the window parameter
MAC Dup Detn Window	The period of time defined by the window parameter for MAC duplication detection
MAC Dup Detn Retry	The period of time after which the MAC in hold-down state is automatically flushed and the MAC duplication process starts again
Number of Dup MACs	The number of duplicate MAC addresses detected
IP Route Advertise*	Not applicable

Label	Description
EVI	The EVPN instance
Ing Rep Inc McastAd	The state of the advertisement of the inclusive multicast Ethernet tag route with tunnel type ingress-replication in the PMSI tunnel attribute
Detected Duplicate MAC Addresses	The MAC address of detected duplicate MACs
Time Detected	The time that the duplicated MAC address was detected
BGP EVPN MPLS Information	
Admin Status	The administrative status of the transport tunnel
Force Vlan Fwding	Indicates whether forced VLAN forwarding is enabled or disabled
Route NextHop Type	Indicates the type of BGP-EVPN route next-hop: system-ipv4, system-ipv6, or explicit if a non-system IP address is configured with route-next-hop
NextHop Addr	The IP address used as the next hop of the BGP-EVPN route This field is only displayed if a non-system IP address is configured with route-next-hop .
Control Word	Indicates whether the control word for the service is enabled or disabled
Split Horizon Group	The name of the split horizon group for all BGP-EVPN MPLS destinations that can be shared by other SAPs and spoke SDPs
Ingress Rep BUM Lbl	Indicates whether sending a separate label for BMU (broadcast, multicast, and unknown unicast) traffic in a specified service—in addition to a label for unicast traffic—is enabled or disabled
Max Ecmp Routes	The number of paths allowed to reach a specified MAC address when that MAC address in the FDB is associated with a remote all-active multihomed Ethernet segment
Ingress Ucast Lbl	The number of ingress unicast labels received
Ingress Mcast Lbl	The number of ingress multicast labels received
Entropy Label	Indicates whether the insertion of the entropy label (EL) and entropy label indicator (ELI) in packets for which at least one LSP in the stack for the far end of the tunnel used by the service has advertised entropy-label capability is enabled or disabled

Label	Description
BGP EVPN MPLS Auto Bind Tunnel Information	
Resolution	The configuration of the auto-bind-tunnel command (disabled, any, or filter)
Filter Tunnel Types	The tunnel types that can be used in the resolution of BGP-EVPN routes within the automatic binding of BGP-EVPN MPLS service to tunnels to MP-BGP peers

ethernet-segment

Syntax

ethernet-segment [*ethernet-segment-name*]

Context

show>service>id

Description

This command displays the SAP and SDP Ethernet segment information.

Parameters

ethernet-segment-name

specifies the name of the Ethernet segment for which the information is displayed

Output

The following output is an example of service Ethernet segment information, and [Table 214: Service ID Ethernet segment field descriptions](#) describes the fields.

Output example

```
*A:Sar18 Dut-B>show# service id 7 ethernet-segment
=====
SAP Ethernet-Segment Information
=====
SAP                Eth-Seg
-----
lag-1:1            eslag1
=====
No sdp entries
*A:Sar18 Dut-B>show#
```

Table 214: Service ID Ethernet segment field descriptions

Label	Description
SAP Ethernet-Segment Information	

Label	Description
SAP	The SAP ID
Eth-Seg	The Ethernet segment name associated with the specified SAP
SDP Ethernet-Segment Information	
SDP	The SDP ID
Eth-Seg	The Ethernet segment name associated with the specified SDP

evpn-mpls

Syntax

evpn-mpls
evpn-mpls esi esi

Context

show>service>id

Description

This command displays the existing EVPN-MPLS destinations for a specified service and all related information. The command allows filtering based on the **esi esi** value (for EVPN multihoming) to display the EVPN-MPLS destinations associated with an ESI.

Parameters

- esi

specifies a 10-byte Ethernet segment identifier (ESI) by which to filter the displayed information

Values

ESI-0, ESI-MAX, or xx-xx-xx-xx-xx-xx-xx-xx-xx, where xx is a hexadecimal number (00 to ff) and the separator is a dash ("-"), colon (":"), or space (" ")

For example, 00-11-22-33-44-55-66-77-88-99
- ieee-address

specifies a 48-bit MAC address by which to filter information. The parameter is entered in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx is a hexadecimal number.

Output

The following output is an example of service EVPN-MPLS information, and [Table 215: Service ID EVPN-MPLS field descriptions](#) describes the fields.

Output example

```

*A:PE1# show service id 1 evpn-mpls
=====
BGP EVPN-MPLS Dest
=====
TEP Address      Egr Label      Num. MACs      Mcast          Last Change
      Transport
-----
192.0.2.69      262140         0              Yes            07/15/2015 19:44:07
                  ldp
192.0.2.69      262141         2              No             07/15/2015 19:44:07
                  ldp
192.0.2.70      262139         0              Yes            07/15/2015 19:44:07
                  ldp
192.0.2.70      262140         1              No             07/15/2015 19:44:07
                  ldp
192.0.2.72      262140         0              Yes            07/15/2015 19:44:07
                  ldp
192.0.2.72      262141         1              No             07/15/2015 19:44:07
                  ldp
192.0.2.73      262139         0              Yes            07/15/2015 19:44:09
                  ldp
192.0.2.254     262142         1              Yes            07/15/2015 19:44:31
                  bgp
-----
Number of entries : 8
=====

BGP EVPN-MPLS Ethernet Segment Dest
=====
Eth SegId              Num. Macs          Last Change
-----
01:00:00:00:00:71:00:00:00:01  2              07/15/2015 20:41:09
01:74:13:00:74:13:00:00:74:13  1              07/15/2015 20:41:07
-----
Number of entries: 2
=====

BGP EVPN-MPLS ES BMAC Dest
=====
vBmacAddr              Num. Macs          Last Change
-----
No Matching Entries
=====

*A:PE1# show service id 1 evpn-mpls esi 01:00:00:00:00:71:00:00:00:01
=====
BGP EVPN-MPLS Ethernet Segment Dest
=====
Eth SegId              Num. Macs          Last Change
-----
01:00:00:00:00:71:00:00:00:01  2              07/15/2015 20:41:09
=====

BGP EVPN-MPLS Dest TEP Info
=====
TEP Address      Egr Label      Last Change
      Transport
-----

```

192.0.2.69	262141	07/15/2015 20:41:09
	ldp	
192.0.2.72	262141	07/15/2015 20:41:09
	ldp	

Number of entries : 2		

=====		

Table 215: Service ID EVPN-MPLS field descriptions

Label	Description
BGP EVPN-MPLS Dest	
TEP Address	The tunnel endpoint address
Egr Label	The egress label for the tunnel endpoint
Transport	The type of transport tunnel for the tunnel endpoint
Num. MACs	The number of MAC addresses associated with the transport tunnel
Mcast	Indicates whether multicast is used
Last Change	The time of the last change to the transport tunnel
BGP EVPN-MPLS Ethernet Segment Dest	
Eth SegId	The Ethernet segment ID of the tunnel endpoint
Num. Macs	The number of MAC addresses associated with the transport endpoint
Last Change	The time of the last change to the transport endpoint
BGP EVPN-MPLS ES BMAC Dest	
vBmacAddr	Not applicable
Num. Macs	Not applicable
Last Change	Not applicable

system

Syntax
system

Context
show>service

Description

This command enables the context to display the service system parameters.

bgp-evpn

Syntax

- bgp-evpn
- bgp-evpn ethernet-segment
- bgp-evpn ethernet-segment name *name* [all]
- bgp-evpn ethernet-segment name *name* evi [*evi*]

Context

show>service>system

Description

This command shows all the information related to the base EVPN instance, including the base RD used for ES routes, Ethernet segments, or individual Ethernet segment information.

Parameters

- ethernet-segment**
 - displays information for BGP-EVPN Ethernet segments
 - name*
 - specifies the name of an Ethernet segment for which to show information
 - all**
 - displays all available information for the specified Ethernet segment
 - evi*
 - displays information for the specified EVI
- Values** 1 to 65535

Output

The following outputs are examples of service system BGP-EVPN information:

- service system BGP-EVPN ([Output example](#) , [Table 216: Service system BGP-EVPN field descriptions](#))
- service system BGP-EVPN Ethernet segment name all ([Output example](#) , [Table 217: Service system BGP-EVPN Ethernet segment name field descriptions](#))
- service system BGP-EVPN Ethernet segment name ([Output example](#), [Table 218: Service system BGP route distinguisher field descriptions](#))

Output example

```
*A:PE1# show service system bgp-evpn
=====
```



```
System BGP EVPN Information
=====
Eth Seg Route Dist.           : <none>
Eth Seg Oper Route Dist.      : <none>
Eth Seg Oper Route Dist Type   : none
Ad Per ES Route Target        : evi-rt
=====

*A:PE1# show service system bgp-evpn ethernet-segment
=====
Service Ethernet Segment
=====
Name                           ESI                           Admin    Oper
-----
ESI-71                         01:00:00:00:00:71:00:00:00:01 Enabled    Up
-----
Entries found: 1
=====
```

Table 216: Service system BGP-EVPN field descriptions

Label	Description
System BGP EVPN Information	
Eth Seg Route Dist.	The IP address of the Ethernet segment route distinguisher
Eth Seg Oper Route Dist.	The IP address of the operational Ethernet segment route distinguisher
Eth Seg Oper Route Dist Type	The operational Ethernet segment route distinguisher type
Ad Per ES Route Target	The configured option for the ad-per-es-route-target command: evi-rt or evi-rt-set
Service Ethernet Segment	
Name	The name of the Ethernet segment (ES)
ESI	The Ethernet segment identifier
Admin	The administrative state of the ES
Oper	The operational state of the ES

Output example

```
*A:7705:Dut-B# show service system bgp-evpn ethernet-segment name "eslag1" all =====
=====Service Ethernet Segment=====
=====Name=====
UpESI           : eslag1Admin State           : Enabled           Oper State           :
Oper Multi-homing : allActiveES SHG Label           : 131046           Source BMAC
LSB             : <none>           Lag Id             : 1           ES Activation
Timer           : 0 secs           Exp/Imp Route-Target : target:bc:01:00:00:00:00Svc Carving
```

: manual		Oper Svc Carving		: manualCfg Range Type		: lowest-	
pref							

DF Pref Election Information							

Preference Mode	Preference Value	Last Admin Change		Oper Pref Value	Do No Preempt		

non-revertive	200	01/11/2023 18:38:49		300	Disabled		

EVI Ranges: <none>							
ISID Ranges: <none>							
=====							
EVI Information							
=====							
EVI	SvcId		Actv Timer Rem		DF		

1	0	1	no3	0	3	no2	0

Number of entries: 4							
=====							
DF Candidate list							

EVI	DF Address						

1	192.0.2.69						
1	10.20.1.3						
2	10.20.1.2						
2	10.20.1.3						
3	10.20.1.2						
3	10.20.1.3						
4	10.20.1.2						
4	10.20.1.3						

Number of entries: 8							

=====							
ISID Information							
=====							
ISID	SvcId		Actv Timer Rem		DF		

No entries found							
=====							

DF Candidate list							

ISID	DF Address						

No entries found							

=====							
BMAC Information							
=====							
SvcId	BMacAddress						

No entries found

=====

Table 217: Service system BGP-EVPN Ethernet segment name field descriptions

Label	Description
Service Ethernet Segment	
Name	The Ethernet segment name
Admin State	The administrative state of the Ethernet segment
Oper State	The operational state of the Ethernet segment
ESI	The Ethernet segment identifier
Multi-homing	The configured multihoming type
Oper Multi-homing	The operational multihoming type
ES SHG Label	The split horizon group label for the Ethernet segment
Source BMAC LSB	Not applicable
Lag Id	The LAG ID for the Ethernet segment
ES Activation Timer	The configured value for the Ethernet segment activation timer
Exp/Imp Route-Target	The export or import route target
Svc Carving	The service-carving algorithm to determine which PE is the designated forwarder (DF)
Oper Svc Carving	The operational service-carving mode, either manual or auto
Cfg Range Type	Specifies the mode that the configured EVI values are in for the Ethernet segment, either lowest-preference (for manual preference-based service-carving) or primary (for manual non-preference based service-carving)
EVI Information	
EVI	The EVPN instance (EVI) for the Ethernet segment
SvcId	The service ID for the Ethernet segment
Actv Timer Rem	The time remaining on the activation timer for the EVI
DF	The designated forwarder for the EVI
DF Candidate list	

Label	Description
EVI	The EVPN instance for the Ethernet segment
DF Address	The IP address of the designated forwarder for the EVI
ISID Information	
ISID	Not applicable
SvcId	Not applicable
Actv Timer Rem	Not applicable
DF	Not applicable
DF Candidate list	
ISID	Not applicable
DF Address	Not applicable
BMAC Information	
SvcId	Not applicable
BMacAddress	Not applicable

Output example

```
*A:Sar18 Dut-B>show>service>system# bgp-evpn ethernet-segment name "ETH_segment_2"
=====
Service Ethernet Segment
=====
Name                : ETH_segment_2
Admin State         : Disabled          Oper State          : Down
ESI                 : 00:00:55:55:55:55:55:55:55
Multi-homing        : all-active          Oper Multi-homing    : all-active
ES SHG Label        : None
Source BMAC LSB     : <none>
ES Activation Timer  : 3 secs (default)
Exp/Imp Route-Target : <none>
Svc Carving         : auto
=====
```

Output example (BGP-EVPN Ethernet segment name EVI)

```
*A:PE-2# show service system bgp-evpn ethernet-segment name "ETH_segment_2" evi
=====
EVI Information
=====
EVI          SvcId          Actv Timer Rem    DF
-----
5            5              0                yes
30           30              0                yes
-----
Number of entries: 2
=====
```

```
*A:7705:Dut-B# show service system bgp-evpn ethernet-segment name "eslag1" evi 1
=====
EVI DF and Candidate List
=====
EVI          SvcId          Actv Timer Rem          DF  DF Last Change
-----
1             1             0             no  01/11/2023 18:38:514
=====

=====
DF Candidates                               Time Added           Oper Pref  Do Not
                                   Value           Preempt
-----
10.20.1.2                               01/11/2023 18:39:58  300        Disabl*
10.20.1.3                               01/11/2023 18:38:51  300        Enabled
-----
Number of entries: 2
=====
* indicates that the corresponding row element may have been truncated.
*A:7705:Dut-B#
```

bgp-route-distinguisher

Syntax

- bgp-route-distinguisher [vpls] [epipe]
- bgp-route-distinguisher svc
- bgp-route-distinguisher ad-evi-rt-set
- bgp-route-distinguisher system

Context

show>service>system

Description

This command shows the information related to BGP route distinguishers for BGP-EVPN service.

Parameters

- vpls**
displays VPLS-related service BGP route distinguisher information
- epipe**
displays Epipe-related service BGP route distinguisher information
- svc**
displays service BGP route distinguisher information
- ad-evi-rt-set**
displays BGP-EVPN Ethernet segment AD EVI RT set route distinguishers
- system**
displays service system BGP route distinguisher information

Output

The following output is an example of BGP route distinguisher information, and [Table 218: Service system BGP route distinguisher field descriptions](#) describes the fields.

Output example (BGP Route Distinguisher VPLS)

```
A:7705:Dut-C# show service system bgp-route-distinguisher vpls
=====
Service Route Distinguishers
=====
Svc Id      Type  Oper Route-Distinguisher      Route-Distinguisher
-----
1           vpls  10.20.1.3:1                  configured
-----
Number of RD Entries: 1
=====
Service System BGP Route Distinguisher Information
=====
Oper Route Distinguisher      Type
-----
Auto-rd                      <none>
Ethernet-segment             10.20.1.3:0                  default
EVI RT Set RD Range          10.10.10.3:1-10.10.10.3:512  configured
=====
BGP EVPN Ethernet Segment AD EVI RT Set Route Distinguishers
=====
Eth Seg      EVI      Svc ID      Route Distinguisher
-----
eslag1       1        1           10.10.10.3:1
-----
Number of Entries: 1
```

Table 218: Service system BGP route distinguisher field descriptions

Label	Description
Service Route Distinguishers	
Svc Id	The service ID
Type	The service type
Oper Route-Distinguisher	The operational route distinguisher
Route-Distinguisher	The type of route distinguisher
Service System BGP Route Distinguisher Information	
Oper Route Distinguisher	The operational route distinguisher for the BGP route distinguisher for the service system
Type	The type of BGP route distinguisher for the service system
BGP EVPN Ethernet Segment AD EVI RT Set Route Distinguishers	
Eth Seg	The Ethernet segment name for the AD EVI RT set

Label	Description
EVI	The EVI ID for the AD EVI RT set
Svc ID	The service ID associated with the AD EVI RT set
Route Distinguisher	The route distinguisher for the AD EVI RT set

redundancy

Syntax

redundancy

Context

show

Description

This command enables the context for the display of global redundancy parameters.

bgp-evpn-multi-homing

Syntax

bgp-evpn-multi-homing

Context

show>redundancy

Description

This command shows the information related to the EVPN global timers.

Output

The following output is an example of BGP-EVPN multihoming information, and [Table 219: Redundancy BGP-EVPN multihoming field descriptions](#) describes the fields.

Output example

```
*A:PE2# show redundancy bgp-evpn-multi-homing
=====
Redundancy BGP EVPN Multi-homing Information
=====
Boot-Timer           : 10 secs
Boot-Timer Remaining : 0 secs
ES Activation Timer   : 3 secs
=====
```

Table 219: Redundancy BGP-EVPN multihoming field descriptions

Label	Description
Redundancy BGP EVPN Multi-homing Information	
Boot-Timer	The value of the BGP EVPN boot timer
Boot-Timer Remaining	The time remaining on the BGP EVPN boot timer
ES Activation Timer	The value of the Ethernet segment activation timer

11 List of acronyms

Table 220: Acronyms

Acronym	Expansion
2G	second-generation wireless telephone technology
3DES	triple DES (data encryption standard)
3G	third-generation mobile telephone technology
6VPE	IPv6 on virtual private edge router
7705 SAR	7705 Service Aggregation Router
7750 SR	7750 Service Router
8 PSK	eight phase shift keying
16 QAM	16-state quadrature amplitude modulation
32 QAM	32-state quadrature amplitude modulation
64 QAM	64-state quadrature amplitude modulation
128 QAM	128-state quadrature amplitude modulation
256 QAM	256-state quadrature amplitude modulation
ABR	area border router available bit rate
AC	alternating current attachment circuit
ACK	acknowledge
ACL	access control list
ACR	adaptive clock recovery
AD	auto-discovery
ADM	add/drop multiplexer
ADP	automatic discovery protocol
AES	advanced encryption standard
AFI	authority and format identifier

Acronym	Expansion
AIGP	accumulated IGP
AIS	alarm indication signal
ALG	application level gateway
AMP	active multipath
AN	association number
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
AU	administrative unit
AUG	administrative unit group
B3ZS	bipolar with three-zero substitution
Batt A	battery A
B-bit	beginning bit (first packet of a fragment)
BBE	background block errors
Bc	committed burst size
Be	excess burst size
BECN	backward explicit congestion notification
Bellcore	Bell Communications Research

Acronym	Expansion
BFD	bidirectional forwarding detection
BGP	border gateway protocol
BGP-LS	border gateway protocol link state
BGP-LU	border gateway protocol labeled unicast
BITS	building integrated timing supply
BTCA	best timeTransmitter clock algorithm
BMU	broadcast, multicast, and unknown traffic Traffic that is not unicast. Any nature of multipoint traffic: <ul style="list-style-type: none"> • 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)
BNM	bandwidth notification message
BOF	boot options file
BoS	bottom of stack
BPDU	bridge protocol data unit
BRAS	Broadband Remote Access Server
BSC	Base Station Controller
BSM	bootstrap message
BSR	bootstrap router
BSTA	Broadband Service Termination Architecture
BTS	base transceiver station
CA	certificate authority connectivity association
CAK	connectivity association key
CAS	channel associated signaling

Acronym	Expansion
CBN	common bonding networks
CBS	committed buffer space
CC	continuity check control channel
CCM	continuity check message
CCTV	closed-circuit television
CE	circuit emulation customer edge
CEM	circuit emulation
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
CKN	connectivity association key name
CLI	command line interface
CLP	cell loss priority
CMP	certificate management protocol
C-multicast	customer multicast
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.
CPROTO	C prototype
CPU	central processing unit

Acronym	Expansion
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
CV	connection verification customer VLAN (tag)
CW	control word
CWDM	coarse wavelength-division multiplexing
DA/FAN	distribution automation and field area network
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 designated forwarder
DH	Diffie-Hellman

Acronym	Expansion
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
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
DVMRP	distance vector multicast routing protocol
e911	enhanced 911 service

Acronym	Expansion
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)
EIR	excess information rate
EJBCA	Enterprise Java Beans Certificate Authority
E-LAN	Ethernet local area network
E-Line	Ethernet virtual private line
EL	entropy label
eLER	egress label edge router
ELI	entropy label indicator
E&M	ear and mouth earth and magneto exchange and multiplexer
eMBMS	evolved MBMS
EOP	end of packet
EPC	evolved packet core
EPD	early packet discard
Epip	Ethernet VLL
EPL	Ethernet private line
EPON	Ethernet Passive Optical Network
EPS	equipment protection switching

Acronym	Expansion
ERO	explicit route object
ES	Ethernet segment errored seconds
ESD	electrostatic discharge
ESI	Ethernet segment identifier
ESMC	Ethernet synchronization message channel
ESN	extended sequence number
ESP	encapsulating security payload
ESPI	encapsulating security payload identifier
ETE	end-to-end
ETH-BN	Ethernet bandwidth notification
ETH-CFM	Ethernet connectivity fault management (IEEE 802.1ag)
EVC	Ethernet virtual connection
EVDO	evolution - data optimized
EVI	EVPN instance
EVPL	Ethernet virtual private link
EVPN	Ethernet virtual private network
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
FEAC	far-end alarm and control
FEC	forwarding equivalence class
FECN	forward explicit congestion notification
FeGW	far-end gateway
FEP	front-end processor

Acronym	Expansion
FF	fixed filter
FFD	fast fault detection
FIB	forwarding information base
FIFO	first in, first out
FIPS-140-2	Federal Information Processing Standard publication 140-2
FM	fault management
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
GPRS	general packet radio service
GPS	Global Positioning System
GRE	generic routing encapsulation
GRT	global routing table
GSM	Global System for Mobile Communications (2G)

Acronym	Expansion
GTP-U	GPRS tunneling protocol user plane
GW	gateway
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
H-VPLS	hierarchical virtual private line service
IANA	Internet Assigned Numbers Authority
IBN	isolated bonding networks
ICB	inter-chassis backup
ICK	integrity connection value key
ICMP	Internet control message protocol
ICMPv6	Internet control message protocol for IPv6
ICP	IMA control protocol cells
ICV	integrity connection value
IDS	intrusion detection system
IDU	indoor unit
IED	intelligent end device
IEEE	Institute of Electrical and Electronics Engineers
IEEE 1588v2	Institute of Electrical and Electronics Engineers standard 1588-2008

Acronym	Expansion
IES	Internet enhanced service
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
IMET-IR	inclusive multicast Ethernet tag—ingress replication
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
I-PMSI	inclusive PMSI
IPoATM	IP over ATM
IPS	intrusion prevention system
IPSec	Internet protocol security
IR	ingress replication
IRB	integrated routing and bridging
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

Acronym	Expansion
IW	interworking
JP	join prune
KEK	key encryption key
KG	key group
LB	loopback
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
LFA	loop-free alternate
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)
LSPA	LSP attributes

Acronym	Expansion
LSR	label switching router link-state request
LSU	link-state update
LT	linktrace
LTE	long term evolution line termination equipment
LTM	linktrace message
LTN	LSP ID to NHLFE
LTR	link trace reply
MA	maintenance association
MAC	media access control
MACsec	media access control security
MA-ID	maintenance association identifier
MBB	make-before-break
MBGP	multicast BGP multiprotocol BGP multiprotocol extensions for BGP
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	multiclass multilink point-to-point protocol
MCS	multicast server multi-chassis synchronization
MCT	MPT craft terminal
MD	maintenance domain

Acronym	Expansion
MD5	message digest version 5 (algorithm)
MDA	media dependent adapter
MDDDB	multidrop data bridge
MDL	maintenance data link
MDT	multicast distribution tree
ME	maintenance entity
MED	multi-exit discriminator
MEF	Metro Ethernet Forum
MEG	maintenance entity group
MEG-ID	maintenance entity group identifier
MEN	Metro Ethernet network
MEP	maintenance association endpoint
MFC	multi-field classification
MHD	multi-homed device
MHF	MIP half function
MHN	multi-homed network
MI	member identifier
MIB	management information base
MI-IS-IS	multi-instance IS-IS
MIR	minimum information rate
MKA	MACsec key agreement
MLD	multicast listener discovery
mLDP	multicast LDP
MLPPP	multilink point-to-point protocol
mLSP	multicast LSP
MoFRR	multicast-only fast reroute
MP	merge point multilink protocol

Acronym	Expansion
	multipoint
MP-BGP	multiprotocol border gateway protocol
MPLS	multiprotocol label switching
MPLSCP	multiprotocol label switching control protocol
MPP	MPT protection protocol
MPR	see Wavence
MPR-e	Microwave Packet Radio (standalone mode)
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)
MRAI	minimum route advertisement interval
MRRU	maximum received reconstructed unit
MRU	maximum receive unit
MSDP	Multicast Source Discovery Protocol
MSDU	MAC Service Data Unit
MSK	master session key
MSO	multi-system operator
MS-PW	multi-segment pseudowire
MSS	maximum segment size Microwave Service Switch
MTIE	maximum time interval error
MTSO	mobile trunk switching office
MTU	maximum transmission unit multi-tenant unit

Acronym	Expansion
M-VPLS	management virtual private line service
MVPN	multicast VPN
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
NFM-P	Network Functions Manager - Packet (formerly 5620 SAM)
NGE	network group encryption
NG-MVPN	next generation MVPN
NH	next hop
NHLFE	next hop label forwarding entry
NHOP	next-hop
NLOS	non-line-of-sight
NLPID	network level protocol identifier
NLRI	network layer reachability information
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)
NOC	network operations center
NPAT	network port address translation
NRC-F	Network Resource Controller - Flow

Acronym	Expansion
NRC-P	Network Resource Controller - Packet
NRC-T	Network Resource Controller - Transport
NRC-X	Network Resource Controller - Cross Domain
NSAP	network service access point
NSD	Network Services Director
NSP	native service processing Network Services Platform
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
OSPF-TE	OSPF-traffic engineering (extensions)

Acronym	Expansion
OSS	operations support system
OSSP	organization specific slow protocol
OTP	one time password
OWAMP	one-way active measurement protocol
P2MP	point to multipoint
PADI	PPPoE active discovery initiation
PADR	PPPoE active discovery request
PAE	port authentication entities
PSB	path state block
PBO	packet byte offset
PBR	policy-based routing
PBX	private branch exchange
PCAP	packet capture
PCC	path computation client
PCE	path computation element
PCEP	Path Computation Element Communication Protocol
PCM	pulse code modulation
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

Acronym	Expansion
PHB	per-hop behavior
PHP	penultimate hop popping
PHY	physical layer
PIC	prefix independent convergence
PID	protocol ID
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
PLSP	path LSP
PM	performance monitoring
PMSI	P-multicast service interface
P-multicast	provider multicast
PN	packet number
PoE	power over Ethernet
PoE+	power over Ethernet plus
POH	path overhead
POI	purge originator identification
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

Acronym	Expansion
PRS	primary reference source
PRTC	primary reference time clock
PSE	power sourcing equipment
PSK	pre-shared key
PSN	packet switched network
PSNP	partial sequence number PDU
PTA	PMSI tunnel attribute
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
PWE	pseudowire emulation
PWE3	pseudowire emulation edge-to-edge
Q.922	ITU-T Q-series Specification 922
QL	quality level
QoS	quality of service
QPSK	quadrature phase shift keying
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

Acronym	Expansion
RIB	routing information base
RIP	routing information protocol
RJ45	registered jack 45
RMON	remote network monitoring
RNC	radio network controller
RP	rendezvous point
RPF RTM	reverse path forwarding RTM
RPS	radio protection switching
RPT	rendezvous-point tree
RR	route reflector
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
RSVP-TE	resource reservation protocol - traffic engineering
RT	receive/transmit
RTC	route target constraint
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 source-active

Acronym	Expansion
SAA	service assurance agent
SAFI	subsequent address family identifier
SAK	security association key
SAP	service access point
SAToP	structure-agnostic TDM over packet
SCADA	supervisory control and data acquisition
SC-APS	single-chassis automatic protection switching
SCI	secure channel identifier
SCP	secure copy
SCTP	Stream Control Transmission Protocol
SD	signal degrade space diversity
SDH	synchronous digital hierarchy
SDI	serial data interface
SDN	software defined network
SDP	service destination point
SE	shared explicit
SeGW	secure gateway
SES	severely errored seconds
SETS	synchronous equipment timing source
SF	signal fail
SFP	small form-factor pluggable (transceiver)
SFTP	SSH file transfer protocol
(S,G)	(source, group)
SGT	self-generated traffic
SHA-1	secure hash algorithm
SHG	split horizon group
SIR	sustained information rate

Acronym	Expansion
SL	short length
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
SPF	shortest path first
SPI	security parameter index
S-PMSI	selective PMSI
SPT	shortest path tree
SR	Service Router (7750 SR) segment routing
SRLG	shared risk link group
SRP	stateful request parameter
SRRP	subscriber routed redundancy protocol
SR-ISIS	segment routing IS-IS
SR-OSPF	segment routing OSPF
SR-TE	segment routing traffic engineering
SSH	secure shell
SSM	source-specific multicast synchronization status messaging
SSU	system synchronization unit
S-tag	service VLAN tag

Acronym	Expansion
STM	synchronous transport module
STM1	synchronous transport module, level 1
STP	spanning tree protocol
STS	synchronous transport signal
SVC	switched virtual circuit
SVEC	synchronization vector
SYN	synchronize
TACACS+	Terminal Access Controller Access-Control System Plus
TC	traffic class (formerly known as EXP bits)
TCI	tag control information
TCP	transmission control protocol
TCP-AO	TCP Authentication Option
TDA	transmit diversity antenna
TDEV	time deviation
TDM	time division multiplexing
TE	traffic engineering
TEDB	traffic engineering database
TEID	tunnel endpoint identifier
TEP	tunnel endpoint
TFTP	trivial file transfer protocol
T-LDP	targeted LDP
TLS	transport layer security
TLV	type length value
TM	traffic management
ToD	time of day
ToS	type of service
T-PE	terminating provider edge router
TPID	tag protocol identifier

Acronym	Expansion
TPIF	IEEE C37.94 teleprotection interface
TPMR	two-port MAC relay
TPS	transmission protection switching
TSoP	transparent SDH/SONET over packet
TTL	time to live
TTLS	tunneled transport layer security
TTM	tunnel table manager
TU	tributary unit
TUG	tributary unit group
TWAMP	two-way active measurement protocol
U-APS	unidirectional automatic protection switching
UAS	unavailable seconds
UBR	unspecified bit rate
UDP	user datagram protocol
UFD	unidirectional forwarding detection
UMH	upstream multicast hop
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
VCB	voice conference bridge
VCC	virtual channel connection
VCCV	virtual circuit connectivity verification
VCI	virtual circuit identifier

Acronym	Expansion
VID	VLAN ID
VLAN	virtual LAN
VLL	virtual leased line
VM	virtual machine
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
VPWS	virtual private wire service
VRF	virtual routing and forwarding table
VRRP	virtual router redundancy protocol
V-SAP	virtual service access point
VSE	vendor-specific extension
VSI	virtual switch instance
VSO	vendor-specific option
VT	virtual trunk virtual tributary
VTG	virtual tributary group
Wavence	formerly 9500 MPR (Microwave Packet Radio)
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

Acronym	Expansion
XOR	exclusive-OR
XRO	exclude route object

12 Supported standards and protocols

This chapter lists the 7705 SAR compliance with security and telecom standards, the protocols supported, and proprietary MIBs.

12.1 Security standards

FIPS 140-2—Federal Information Processing Standard publication 140-2, Security Requirements for Cryptographic Modules

12.2 Telecom standards

ANSI/TIA/EIA-232-C—Interface Between Data Terminal Equipment and Data Circuit-Terminating Equipment Employing Serial Binary Data Interchange

IEEE 802.1AB-2016—IEEE Standard for Local and metropolitan area networks - Station and Media Access Control Connectivity Discovery

IEEE 802.1ad—IEEE Standard for Local and Metropolitan Area Networks—Virtual Bridged Local Area Networks

IEEE Std 802.1AE-2006 Media Access Control (MAC) Security

IEEE Std 802.1AEbw-2013—Media Access Control (MAC) Security Amendment 2: Extended Packet Numbering

IEEE 802.1ag—Service Layer OAM

IEEE 802.1p/q—VLAN Tagging

IEEE 802.1x-2010—IEEE Standard for Local and Metropolitan Area Networks—Port-based Network Access Control

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 C37.94-2017—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.826—End-to-end error performance parameters and objectives for international, constant bit-rate digital paths and connections

ITU-T G.8032 — Ethernet Ring Protection Switching

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

12.3 Protocol support

12.3.1 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

12.3.2 BFD

RFC 7130—Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces

RFC 7881—Seamless Bidirectional Forwarding Detection (S-BFD) for IPv4, IPv6, and MPLS

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

12.3.3 BGP

RFC 1397—BGP Default Route Advertisement
RFC 1997—BGP Communities Attribute
RFC 2385—Protection of BGP Sessions via the TCP MD5 Signature Option
RFC 2439—BGP Route Flap Dampening
RFC 2545—Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
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 4684—Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)
RFC 4724—Graceful Restart Mechanism for BGP - GR Helper
RFC 4760—Multi-protocol Extensions for BGP (previously RFC 2858)
RFC 4893—BGP Support for Four-octet AS Number Space
RFC 4798—Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)
RFC 5549—Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop
RFC 5925—The TCP Authentication Option
RFC 5926—Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)
RFC 6513—Multicast in MPLS/BGP IP VPNs
RFC 6514—BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs
RFC 7311—The Accumulated IGP Metric Attribute for BGP
RFC 7606—Revised Error Handling for BGP UPDATE Messages
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

12.3.4 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

12.3.5 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

12.3.6 Digital data network management

V.35

RS-232 (also known as EIA/TIA-232)

X.21

12.3.7 ECMP

RFC 2992—Analysis of an Equal-Cost Multi-Path Algorithm

12.3.8 Ethernet VPN (EVPN)

RFC 7432—BGP MPLS-Based Ethernet VPN

draft-ietf-bess-evpn-vpls-seamless-integ—(PBB-)EVPN Seamless Integration with (PBB-)VPLS

draft-ietf-bess-evpn-vpws—Virtual Private Wire Service support in Ethernet VPN

draft-ietf-rabadan-bess-evpn-pref-pdf—Preference-based EVPN DF Election

12.3.9 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

RFC 2427—Multiprotocol Interconnect over Frame Relay

12.3.10 GRE

RFC 2784—Generic Routing Encapsulation (GRE)

12.3.11 Internet protocol (IP) – version 4

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
RFC 3021—Using 31-Bit Prefixes on IPv4 Point-to-Point Links

12.3.12 Internet protocol (IP) – version 6

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

12.3.13 IPSec

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 2409—The Internet Key Exchange (IKE)
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 4301—Security Architecture for the Internet Protocol
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)
RFC 7383—Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation

12.3.14 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 4971—Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information
RFC 5120—M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)
RFC 5304—IS-IS Cryptographic Authentication
RFC 5305—IS-IS Extensions for Traffic Engineering
RFC 5307—IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)
RFC 5308—Routing IPv6 with IS-IS
RFC 5309—Point-to-Point Operation over LAN in Link State Routing Protocols
RFC 5310—IS-IS Generic Cryptographic Authentication
RFC 6232—Purge Originator Identification TLV for IS-IS

12.3.15 LDP

RFC 5036—LDP Specification

RFC 5283—LDP Extension for Inter-Area Label Switched Paths

RFC 5350—IANA Considerations for the IPv4 and IPv6 Router Alert Options

RFC 5443—LDP IGP Synchronization

RFC 5561—LDP Capabilities

RFC 6388—Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths

RFC 6512—Using Multipoint LDP When the Backbone Has No Route to the Root

RFC 6829—Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6

RFC 7552—Updates to LDP for IPv6

draft-ietf-mpls-ldp-ip-pw-capability—Controlling State Advertisements Of Non-negotiated LDP Applications

draft-ietf-mpls-oam-ipv6-rao—IPv6 Router Alert Option for MPLS OAM

draft-pdutta-mpls-ldp-adj-capability-00—LDP Adjacency Capabilities

draft-pdutta-mpls-ldp-v2-00—LDP Version 2

draft-pdutta-mpls-mlldp-up-redundancy-00.txt—Upstream LSR Redundancy for Multi-point LDP Tunnels

12.3.16 LDP and IP FRR

RFC 5286—Basic Specification for IP Fast Reroute: Loop-Free Alternates

RFC 7490—Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)

12.3.17 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)

RFC 5440—Path Computation Element (PCE) Communication Protocol (PCEP)

RFC 6790—The Use of Entropy Labels in MPLS Forwarding

RFC 8253—PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)

RFC 8697—Path Computation Element Communication Protocol (PCEP) Extensions for Establishing Relationships between Sets of Label Switched Paths (LSPs)

RFC 8745—Path Computation Element Communication Protocol (PCEP) Extensions for Associating Working and Protection Label Switched Paths (LSPs) with Stateful PCE

RFC 8800—Path Computation Element Communication Protocol (PCEP) Extension for Label Switched Path (LSP) Diversity Constraint Signaling

draft-dhody-pce-pceps-tls13-02—Updates for PCEPS

draft-ietf-pce-stateful-pce—PCEP Extensions for Stateful PCE

draft-ietf-pce-segment-routing—PCEP Extensions for Segment Routing

draft-alvarez-pce-path-profiles—PCE Path Profiles

12.3.18 MPLS – OAM

RFC 6424—Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels

RFC 8029—Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures

12.3.19 Multicast

RFC 3956—Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address

RFC 3973—Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)

RFC 4610—Anycast-RP Using Protocol Independent Multicast (PIM), which is similar to RFC 3446—Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)

RFC 6514—BGP Encodings and Procedures for Multicast in MPLS/IP VPNs

RFC 6826—Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths

cisco-ipmulticast/pim-autorp-spec—Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast, which is similar to RFC 5059—Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)

draft-ietf-l2vpn-vpls-pim-snooping-07—Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)

draft-ietf-mboned-msdp-deploy-nn.txt—Multicast Source Discovery Protocol (MSDP) Deployment Scenarios

12.3.20 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
RFC 3954—Cisco Systems NetFlow Services Export Version 9
RFC 5101—Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information
RFC 5102—Information Model for IP Flow Information Export
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

12.3.21 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 3623—Graceful OSPF Restart (support for Helper mode)
RFC 3630—Traffic Engineering (TE) Extensions to OSPF
RFC 4203 for Shared Risk Link Group (SRLG) sub-TLV
RFC 4577—OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) (support for basic OSPF at PE-CE links)
RFC 4915—Multi-Topology (MT) Routing in OSPF
RFC 4970—Extensions to OSPF for Advertising Optional Router Capabilities
RFC 5185—OSPF Multi-Area Adjacency

12.3.22 OSPFv3

RFC 4552—Authentication/Confidentiality for OSPFv3

12.3.23 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

12.3.24 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

12.3.25 RIP

RFC 1058—Routing Information Protocol
RFC 2453—RIP Version 2

12.3.26 RADIUS

RFC 2865—Remote Authentication Dial In User Service
RFC 2866—RADIUS Accounting

12.3.27 RSVP-TE and FRR

RFC 2430—A Provider Architecture for DiffServ & TE
RFC 2702—Requirements for Traffic Engineering over MPLS
RFC 2747—RSVP Cryptographic Authentication
RFC 2961—RSVP Refresh Overhead Reduction Extensions
RFC 3097—RSVP Cryptographic Authentication - Updated Message Type Value
RFC 3209—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

12.3.28 Segment routing (SR)

draft-francois-rtgwg-segment-routing-ti-lfa-04—Topology Independent Fast Reroute using Segment Routing

draft-gredler-idr-bgp-ls-segment-routing-ext-03—BGP Link-State extensions for Segment Routing

draft-ietf-isis-segment-routing-extensions-04—IS-IS Extensions for Segment Routing

draft-ietf-mpls-spring-lsp-ping-02—Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane

draft-ietf-ospf-segment-routing-extensions-04—OSPF Extensions for Segment Routing

draft-ietf-spring-segment-routing-15—Segment Routing Architecture

12.3.29 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

12.3.30 SSH

RFC 4253—The Secure Shell (SSH) Transport Layer Protocol

draft-ietf-secsh-architecture.txt—SSH Protocol Architecture

draft-ietf-secsh-userauth.txt—SSH Authentication 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

12.3.31 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

IEC/IEEE 61850-9-3—Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation

IEEE C37.238-2017 - IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications

IEEE Std 1588-2008—IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

IEEE Std 1588-2008—IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, Annex E – Transport of PTP over User Datagram Protocol over Internet Protocol Version 6

IEEE Std 1588-2019—IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, Annex J

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

ITU-T G.8275.2—Telecommunication Standardization Section of ITU, Precision time protocol telecom profile for time/phase synchronization with partial timing support from the network, issued 06/2016

RFC 5905—Network Time Protocol Version 4: Protocol and Algorithms Specification

RFC 8573—Message Authentication Code for the Network Time Protocol

12.3.32 TACACS+

IETF draft-grant-tacacs-02.txt—The TACACS+ Protocol

12.3.33 TLS

RFC 5246—The Transport Layer Security (TLS) Protocol Version 1.2

RFC 5425—Transport Layer Security (TLS) Transport Mapping for Syslog

RFC 5922—Domain Certificates in the Session Initiation Protocol (SIP)

RFC 6460—Suite B Profile for Transport Layer Security (TLS)

RFC 8446—The Transport Layer Security (TLS) Protocol Version 1.3

12.3.34 TWAMP

RFC 5357—A Two-Way Active Measurement Protocol (TWAMP)

12.3.35 VPLS

RFC 4762—Virtual Private LAN Services Using LDP

12.3.36 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

12.4 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 document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)