



7210 Service Access System

Release 26.3.R1

7210 SAS-Mxp, S, Sx, T Services Guide

3HE 27848 AAAA TQZZA 01

Edition: 01

March 2026

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.

© 2026 Nokia.

Table of contents

List of tables.....	18
List of figures.....	25
1 Getting started.....	30
1.1 About this guide.....	30
1.1.1 Document structure and content.....	30
1.2 7210 SAS modes of operation.....	31
1.3 7210 SAS port modes.....	33
1.4 7210 SAS services configuration process.....	35
1.5 Conventions.....	36
1.5.1 Precautionary and information messages.....	36
1.5.2 Options or substeps in procedures and sequential workflows.....	37
2 Services overview.....	38
2.1 Introduction.....	38
2.1.1 Service types.....	38
2.1.2 Service policies.....	39
2.2 Nokia service model.....	39
2.3 Service entities.....	40
2.3.1 Customers.....	40
2.3.2 SAPs.....	40
2.3.2.1 SAP encapsulation types and identifiers.....	42
2.3.2.2 Ethernet encapsulations.....	42
2.3.2.3 Services and SAP encapsulations.....	43
2.3.2.4 Default SAP on a dot1q port.....	43
2.3.2.5 Default SAPs on a QinQ port (supported only on 7210 SAS devices configured in access-uplink mode).....	44
2.3.2.6 SAP configuration considerations for network mode and access-uplink mode....	46
2.3.2.7 SAP configuration notes when operating 7210 SAS devices in network mode only.....	46
2.3.2.8 QinQ SAP Configuration Restrictions for 7210 SAS platforms in network operating mode.....	47
2.3.2.9 SAP configuration notes for 7210 SAS platforms in access-uplink operating mode.....	48

2.3.3	SDPs.....	49
2.3.3.1	SDP binding.....	50
2.3.3.2	Spoke and mesh SDPs.....	50
2.3.3.3	SDP using BGP route tunnel.....	51
2.3.3.4	SDP keepalives.....	51
2.3.3.5	Mixed-LSP mode of operation.....	52
2.3.4	SAP and service scaling with high SAP scale mode.....	53
2.3.4.1	Guidelines for configuring high SAP scale mode.....	54
2.3.4.2	Guidelines for configuring low SAP scale mode.....	55
2.3.5	G.8032 Ethernet ring protection switching.....	55
2.3.5.1	Overview of G.8032 operation.....	56
2.3.5.2	Ethernet ring sub-rings.....	60
2.3.5.3	OAM considerations.....	65
2.3.5.4	QoS considerations.....	65
2.3.5.5	Support service and solution combinations.....	65
2.3.5.6	Configuration guidelines for G.8032.....	66
2.4	Service creation process overview.....	66
2.5	Deploying and provisioning services.....	67
2.5.1	Phase 1: core network construction.....	67
2.5.2	Phase 2: service administration.....	68
2.5.3	Phase 3: service provisioning.....	68
2.6	Configuration notes.....	68
2.6.1	General.....	68
2.7	Configuring global service entities with CLI.....	69
2.8	Service model entities.....	69
2.9	Basic configuration.....	69
2.10	Common configuration tasks.....	70
2.10.1	Configuring customers accounts.....	70
2.10.1.1	Customer information.....	71
2.10.2	Configuring an SDP.....	71
2.10.2.1	SDP configuration tasks.....	71
2.10.2.2	Configuring an SDP.....	72
2.10.2.3	Configuring a mixed-LSP SDP.....	73
2.11	Ethernet connectivity fault management.....	74
2.11.1	MA, MEP, MIP, and MD levels.....	75
2.11.1.1	Common actionable failures.....	77

2.11.1.2	MEP and MIP support.....	78
2.11.2	Configuring ETH-CFM parameters.....	78
2.11.3	Applying ETH-CFM parameters.....	80
2.12	Layer 2 control processing.....	81
2.13	Dynamic VLAN assignment using dot1x RADIUS authentication with EHS.....	83
2.13.1	Assignment of a VLAN after a host is authenticated using dot1x.....	84
2.13.2	Assigning a VLAN to a device based on dot1x authentication.....	85
2.13.3	Guest and restricted VLAN service support.....	85
2.13.4	Configuration guidelines.....	87
2.14	Service management tasks.....	87
2.14.1	Modifying customer accounts.....	87
2.14.2	Deleting customers.....	88
2.14.3	Modifying SDPs.....	88
2.14.4	Deleting SDPs.....	89
2.15	Global services command reference.....	89
2.15.1	Command hierarchies.....	89
2.15.1.1	Global service configuration commands.....	89
2.15.1.2	Show commands.....	93
2.15.1.3	Tools commands.....	93
2.15.2	Command descriptions.....	94
2.15.2.1	Global service configuration commands.....	94
2.15.2.2	Show commands.....	128
2.15.2.3	Tools commands.....	161
3	VLL services.....	167
3.1	Epipe services.....	167
3.1.1	Epipe services overview.....	167
3.1.2	Epipe with PBB.....	167
3.1.3	Processing packets received with more than two tags on a QinQ SAP in Epipe service.....	168
3.1.3.1	Feature support, configuration notes, and restrictions.....	169
3.1.3.2	Configuration example of Epipe services for processing of packets received with more than two tags on a QinQ SAP.....	170
3.1.4	Epipe operational state decoupling.....	171
3.2	Pseudowire switching.....	172
3.2.1	Pseudowire switching with protection.....	173
3.2.2	Pseudowire switching behavior.....	174

3.2.2.1	Pseudowire switching TLV.....	174
3.2.2.2	Static-to-dynamic pseudowire switching.....	175
3.2.3	Pseudowire redundancy.....	175
3.2.3.1	VLL resilience with two destination PE nodes.....	176
3.2.4	Dynamic multi-segment pseudowire routing.....	177
3.2.4.1	Pseudowire routing.....	181
3.2.4.2	Configuring VLLs using dynamic MS-PW.....	183
3.2.4.3	Pseudowire redundancy.....	184
3.2.4.4	VCCV OAM for dynamic MS-PWs.....	186
3.2.4.5	VCCV-Ping on dynamic MS-PWs.....	186
3.2.4.6	VCCV-Trace on dynamic MS-PWs.....	186
3.2.5	Example dynamic MS-PW configuration.....	187
3.3	Master-slave operation.....	191
3.3.1	Operation of master-slave pseudowire redundancy with existing scenarios.....	192
3.3.1.1	VLL resilience path.....	192
3.3.2	VLL resilience for a switched PW path.....	194
3.3.3	Access mode resilience Using MC-LAG and pseudowire redundancy.....	195
3.3.4	VLL resilience for a switched pseudowire path.....	197
3.4	Pseudowire redundancy service models.....	198
3.4.1	Redundant VLL service model.....	198
3.4.2	T-LDP status notification handling rules.....	200
3.4.2.1	Processing endpoint SAP active/standby status bits.....	200
3.4.2.2	Processing and merging.....	200
3.5	Epipe configuration for MPLS-TP.....	202
3.5.1	SDPs.....	202
3.5.2	VLL spoke-SDP configuration.....	203
3.5.3	Credit-based algorithm.....	205
3.6	VLAN range for SAPs in an Epipe service.....	206
3.6.1	Processing behavior for SAPs using VLAN ranges in access-uplink mode.....	206
3.6.2	VLAN range SAPs feature support and restrictions.....	206
3.6.3	Processing behavior for SAPs using VLAN ranges in network operating mode.....	207
3.7	VLL service considerations.....	208
3.7.1	SDPs.....	208
3.7.2	SAP encapsulations.....	208
3.7.3	QoS policies.....	209
3.7.4	Filter policies.....	209

3.7.5	MAC resources.....	209
3.8	Configuring a VLL service with CLI.....	209
3.8.1	Common configuration tasks.....	210
3.8.2	Configuring VLL components.....	210
3.8.2.1	Creating an Epipe service in network mode.....	210
3.8.2.2	Creating an Epipe service in access-uplink mode.....	210
3.8.2.3	Creating an Epipe service for 7210 SAS-Mxp with range SAPs.....	214
3.8.2.4	Configuring default QinQ SAPs for Epipe transit traffic in a ring scenario in access-uplink mode.....	215
3.8.2.5	Configuring SDP bindings.....	216
3.8.3	Using spoke-SDP control words.....	218
3.8.4	Configuring VLL resilience.....	219
3.8.5	Configuring VLL resilience for a switched pseudowire path.....	220
3.8.6	Service management tasks.....	221
3.8.6.1	Modifying Epipe service parameters.....	221
3.8.6.2	Disabling an Epipe service.....	222
3.8.6.3	Re-enabling an Epipe service.....	222
3.8.6.4	Deleting an Epipe service.....	222
3.9	VLL services command reference.....	223
3.9.1	Command hierarchies.....	223
3.9.1.1	VLL service configuration commands.....	223
3.9.1.2	Show commands.....	228
3.9.1.3	Clear commands.....	229
3.9.2	Command descriptions.....	229
3.9.2.1	VLL service configuration commands.....	229
3.9.2.2	Show commands.....	294
3.9.2.3	Clear commands.....	348
3.9.2.4	Debug commands.....	352
4	Ethernet Virtual Private Networks.....	356
4.1	EVPN applications.....	356
4.1.1	EVPN for MPLS tunnels in E-LAN services.....	356
4.1.2	EVPN for MPLS tunnels in E-Line services.....	357
4.2	EVPN for MPLS tunnels.....	358
4.2.1	BGP-EVPN control plane for MPLS tunnels.....	358
4.2.1.1	EVPN route type 3 — inclusive multicast Ethernet tag route.....	359

4.2.1.2	EVPN route type 2 — MAC/IP advertisement route.....	359
4.2.1.3	EVPN route type 1 — Ethernet Auto-Discovery route.....	360
4.2.1.4	EVPN route type 4 — ES route.....	361
4.2.1.5	BGP tunnel encapsulation extended community.....	361
4.2.2	EVPN for MPLS tunnels in VPLS services.....	362
4.2.2.1	EVPN and VPLS integration.....	365
4.2.2.2	Auto-Derived RD in services with multiple BGP families.....	368
4.2.3	EVPN multi-homing in VPLS services.....	369
4.2.4	EVPN all-active multi-homing.....	369
4.2.4.1	All-active multi-homing procedures.....	370
4.2.4.2	All-active multi-homing service model.....	371
4.2.4.3	ES discovery and DF election procedures.....	373
4.2.4.4	Aliasing.....	378
4.2.4.5	Network failures and convergence for all-active multi-homing.....	380
4.2.4.6	Logical failures on ESs and black holes.....	381
4.2.4.7	Transient issues because of MAC route delays.....	381
4.2.5	EVPN single-active multi-homing.....	382
4.2.5.1	Single-active multi-homing service model.....	383
4.2.5.2	ES and DF election procedures.....	385
4.2.5.3	Backup PE function.....	386
4.2.5.4	Network failures and convergence for single-active multi-homing.....	387
4.2.6	EVPN-VPWS for MPLS tunnels.....	388
4.2.6.1	BGP-EVPN control plane for EVPN-VPWS.....	388
4.2.6.2	EVPN for MPLS tunnels in Epipe services (EVPN-VPWS).....	389
4.2.6.3	Using A/S PW and MC-LAG with EVPN-VPWS Epipes.....	391
4.2.6.4	LAG-based LLF for EVPN-VPWS services.....	392
4.2.6.5	LAG or port standby signaling to the CE on non-DF EVPN PEs (single-active).....	394
4.2.7	BGP and EVPN route selection for EVPN routes.....	395
4.3	General EVPN topics.....	396
4.3.1	ARP and ND snooping and proxy support.....	396
4.3.1.1	Proxy-ARP/ND periodic refresh, unsolicited refresh, and confirm-messages.....	400
4.3.1.2	Proxy-ND and the Router flag in Neighbor Advertisement messages.....	401
4.3.1.3	Procedure to add the R Flag to a specified entry.....	401
4.3.1.4	Configuration guidelines for proxy-ARP and proxy-ND.....	401
4.3.2	BGP-EVPN MAC mobility.....	403
4.3.3	BGP-EVPN MAC duplication.....	403

4.3.4	Conditional static MAC and protection.....	405
4.3.5	BGP and EVPN route selection for EVPN routes.....	406
4.3.6	EVPN interaction with other features.....	407
4.3.6.1	EVPN-MPLS with existing VPLS features.....	407
4.3.6.2	EVPN with G.8032 in an access ring.....	408
4.3.7	Routing policies for BGP EVPN routes.....	413
4.4	Configuring an EVPN service with CLI.....	413
4.4.1	EVPN-MPLS configuration examples.....	413
4.4.1.1	EVPN single-active multi-homing example.....	413
4.5	EVPN command reference.....	414
4.5.1	Command hierarchies.....	414
4.5.1.1	EVPN configuration commands.....	414
4.5.1.2	EVPN show commands.....	417
4.5.1.3	EVPN clear commands.....	417
4.5.1.4	EVPN tools commands.....	417
4.5.2	Command descriptions.....	417
4.5.2.1	EVPN configuration commands.....	417
4.5.2.2	EVPN show commands.....	459
4.5.2.3	EVPN clear commands.....	471
4.5.2.4	Tools commands.....	472
5	Virtual Private LAN Service.....	474
5.1	VPLS service overview.....	474
5.1.1	VPLS packet walk-through in network mode.....	474
5.1.2	VPLS packet walk-through in access-uplink mode.....	476
5.2	VPLS features.....	478
5.2.1	VPLS enhancements.....	478
5.2.2	VPLS over MPLS in network operating mode.....	478
5.2.3	VPLS over QinQ spokes for 7210 SAS devices configured in access-uplink operating mode.....	479
5.2.4	VPLS MAC learning and packet forwarding.....	479
5.2.5	IGMP snooping in a VPLS service.....	480
5.2.5.1	Configuration guidelines for IGMP snooping in VPLS service.....	481
5.2.6	DHCPv4 snooping.....	482
5.2.7	DHCPv6 snooping.....	482
5.2.7.1	Client and network-facing service objects.....	483

5.2.7.2	DHCPv6 relay agent options.....	484
5.2.7.3	DHCPv6 snooping QoS considerations.....	484
5.2.8	Multicast VLAN Registration (MVR) support in VPLS service.....	485
5.2.8.1	Configuration guidelines for MVR in VPLS services.....	486
5.2.9	Layer 2 forwarding table management.....	486
5.2.9.1	FIB size.....	486
5.2.9.2	FIB size alarms.....	487
5.2.9.3	Local and remote aging timers.....	487
5.2.9.4	Disable MAC aging.....	487
5.2.9.5	Disable MAC learning.....	487
5.2.9.6	Unknown MAC discard.....	488
5.2.9.7	VPLS and rate limiting.....	488
5.2.9.8	MAC move.....	488
5.2.10	VPLS and spanning tree protocol.....	489
5.2.10.1	Spanning tree operating modes.....	489
5.2.10.2	Multiple Spanning Tree.....	490
5.2.10.3	MSTP for QinQ SAPs.....	491
5.2.10.4	Provider MSTP.....	491
5.2.10.5	Enhancements to the Spanning Tree Protocol.....	492
5.2.11	VPLS redundancy.....	494
5.2.11.1	Spoke-SDP redundancy for metro interconnection.....	494
5.2.11.2	Spoke-SDP-based redundant access.....	495
5.2.11.3	Inter-domain VPLS resiliency using multi-chassis endpoints.....	496
5.2.12	VPLS access redundancy.....	496
5.2.12.1	STP-based redundant access to VPLS.....	497
5.2.12.2	Redundant access to VPLS without STP.....	497
5.2.13	MAC flush message processing.....	498
5.2.13.1	MAC Flush with STP.....	499
5.2.13.2	Selective MAC flush.....	500
5.2.13.3	Dual homing to a VPLS service.....	500
5.2.14	VPLS service considerations.....	501
5.2.14.1	SAP encapsulations.....	501
5.2.14.2	VLAN processing.....	501
5.3	BGP Auto-Discovery for LDP VPLS.....	502
5.3.1	BGP AD overview.....	502
5.3.2	Information model.....	502

5.3.3	FEC element for T-LDP signaling.....	503
5.3.4	BGP-AD and Target LDP (T-LDP) interaction.....	504
5.3.5	SDP usage.....	506
5.3.6	Automatic creation of SDPs.....	506
5.3.7	Manually provisioned SDP.....	506
5.3.8	Automatic instantiation of pseudowires (SDP bindings).....	507
5.3.9	Mixing statically configured and auto-discovered pseudowires in a VPLS service.....	507
5.3.10	Resiliency schemes.....	507
5.4	Routed VPLS.....	508
5.4.1	IES or VPRN IP interface binding.....	508
5.4.2	Assigning a service name to a VPLS service.....	508
5.4.3	Service binding requirements.....	509
5.4.3.1	Bound Service Name Assignment.....	509
5.4.3.2	Binding a service name to an IP interface.....	509
5.4.3.3	IP interface attached VPLS service constraints.....	509
5.4.3.4	IP interface and VPLS operational state coordination.....	510
5.4.3.5	IP interface MTU and fragmentation.....	510
5.4.4	ARP and VPLS FIB interactions.....	510
5.4.5	R-VPLS specific ARP cache behavior.....	511
5.4.5.1	The allow-ip-int-binding VPLS flag.....	511
5.4.6	R-VPLS SAP support on standard Ethernet ports.....	512
5.4.6.1	LAG port membership constraints.....	512
5.4.6.2	VPLS feature support and restrictions.....	512
5.4.7	VPLS SAP ingress IP filter override.....	513
5.4.7.1	QoS support for VPLS SAPs and IP interface in a R-VPLS service.....	514
5.4.7.2	R-VPLS and routing protocols support.....	515
5.4.7.3	Spanning tree and split horizon.....	516
5.4.8	R-VPLS MAC ACLs.....	516
5.4.8.1	R-VPLS features supported with MAC ACLs.....	517
5.4.9	R-VPLS and IGMPv3 snooping.....	518
5.4.9.1	Configuration guidelines and restrictions for IGMP snooping in R-VPLS.....	518
5.4.10	R-VPLS supported functionality and restrictions.....	519
5.5	Eppe emulation using dot1q VLAN range SAP in VPLS with G.8032.....	520
5.5.1	Eppe emulation configuration guidelines and restrictions.....	522
5.6	Configuring a VPLS service with CLI.....	523
5.6.1	Basic configuration.....	523

5.6.2	Common configuration tasks.....	525
5.6.3	Configuring VPLS components.....	526
5.6.3.1	Creating a VPLS service.....	526
5.6.3.2	Configuring a VPLS SAP.....	531
5.6.3.3	Configuring SDP bindings.....	539
5.6.4	Configuring VPLS redundancy.....	540
5.6.4.1	Creating a management VPLS for SAP protection.....	540
5.6.4.2	Creating a management VPLS for spoke-SDP protection.....	542
5.6.4.3	Configuring load Balancing with management VPLS.....	544
5.6.4.4	Configuring a BGP-auto-discovery.....	546
5.6.4.5	Configuring load balancing with management VPLS.....	547
5.6.4.6	Configuring selective MAC flush.....	552
5.6.4.7	Configuring load balancing with management VPLS.....	552
5.6.5	Configuring IGMPv3 snooping in RVPLS.....	554
5.6.6	Configuring BGP Auto-Discovery.....	556
5.6.6.1	Configuration steps.....	556
5.6.7	Configuring AS pseudowire in VPLS.....	558
5.7	Service management tasks.....	559
5.7.1	Modifying VPLS service parameters.....	559
5.7.2	Modifying management VPLS parameters.....	559
5.7.3	Deleting a management VPLS.....	559
5.7.4	Disabling a management VPLS.....	560
5.7.5	Deleting a VPLS service.....	560
5.7.6	Disabling a VPLS service.....	560
5.7.7	Re-enabling a VPLS service.....	561
5.8	VPLS services command reference.....	561
5.8.1	Command hierarchies.....	561
5.8.1.1	VPLS service configuration commands.....	562
5.8.1.2	Show commands.....	572
5.8.1.3	Clear commands.....	573
5.8.1.4	Debug commands.....	574
5.8.2	Command descriptions.....	574
5.8.2.1	VPLS configuration commands.....	574
5.8.2.2	VPLS show commands.....	685
5.8.2.3	IGMP snooping show commands.....	749
5.8.2.4	VPLS clear commands.....	771

5.8.2.5	VPLS debug commands.....	780
6	IEEE 802.1ah Provider Backbone Bridging (PBB).....	786
6.1	IEEE 802.1ah PBB overview.....	786
6.2	PBB features.....	787
6.2.1	Integrated PBB-VPLS solution.....	787
6.2.2	PBB technology.....	788
6.2.3	PBB mapping to existing VPLS configurations.....	789
6.2.4	SAP support.....	790
6.2.4.1	PBB B-VPLS.....	790
6.2.4.2	PBB I-VPLS.....	791
6.2.5	PBB packet walk-through.....	791
6.2.6	PBB ELINE service.....	793
6.2.6.1	PBB resiliency for PBB Epipe service.....	793
6.2.6.2	PBB resiliency for B-VPLS.....	793
6.2.7	Access multi-homing for native PBB (B-VPLS over SAP infrastructure).....	793
6.2.8	PBB QoS.....	794
6.2.9	PBB ACL support.....	795
6.2.10	Configuration guidelines.....	795
6.2.11	Configuration guidelines for the 7210 SAS-T.....	796
6.3	Configuration examples.....	796
6.3.1	PBB ELAN and ELINE.....	797
6.3.2	MC-LAG multi-homing for native PBB.....	797
6.4	PBB command reference.....	798
6.4.1	Command hierarchies.....	798
6.4.1.1	PBB service commands.....	798
6.4.1.2	Show commands.....	799
6.4.1.3	Clear commands.....	799
6.4.1.4	Debug commands.....	799
6.4.2	Command descriptions.....	800
6.4.2.1	PBB service configuration commands.....	800
6.4.2.2	PBB show commands.....	805
6.4.2.3	PBB clear commands.....	837
6.4.2.4	PBB debug commands.....	841
7	Internet Enhanced Service.....	846

7.1	IES service overview.....	846
7.2	IES features.....	847
7.2.1	IP interfaces.....	847
7.2.1.1	IPv6 support for IES IP interfaces (access-uplink operating mode).....	847
7.2.1.2	IPv6 support for IES IP interfaces (network operating mode).....	848
7.2.1.3	Encapsulations.....	848
7.2.2	Routing protocols.....	849
7.2.2.1	CPE connectivity check.....	849
7.2.3	QoS policies.....	849
7.2.3.1	CPU QoS for IES interfaces in access-uplink mode.....	850
7.2.3.2	CPU QoS for IES access interfaces in network mode.....	850
7.2.4	Filter policies.....	850
7.2.5	VRRP support for IES IP interfaces in network operating mode.....	851
7.3	Configuring an IES service with CLI.....	851
7.3.1	Basic configuration.....	851
7.3.2	Common configuration tasks.....	852
7.3.3	Configuring IES Components.....	853
7.3.3.1	Configuring an IES service.....	853
7.3.3.2	Configuring IES interface parameters.....	853
7.3.3.3	Configuring IES SAP parameters.....	854
7.3.3.4	Configuring VRRP.....	854
7.3.4	Service management tasks.....	855
7.3.4.1	Modifying IES service parameters.....	855
7.3.4.2	Deleting an IES service.....	855
7.3.4.3	Disabling an IES service.....	856
7.3.4.4	Re-enabling an IES service.....	856
7.4	IES services command reference.....	856
7.4.1	Command hierarchies.....	856
7.4.1.1	IES configuration commands.....	857
7.4.1.2	Show commands.....	862
7.4.2	Command descriptions.....	862
7.4.2.1	IES service configuration commands.....	862
7.4.2.2	Show commands.....	919
8	Virtual Private Routed Network service.....	944
8.1	VP RN service overview.....	944

8.1.1	Routing prerequisites.....	945
8.1.2	BGP support.....	945
8.1.3	Route distinguishers.....	946
8.1.3.1	Route reflector.....	946
8.1.3.2	Customer Edge to Provider Edge route exchange.....	946
8.1.4	Constrained Route Distribution.....	948
8.1.4.1	Constrained VPN route distribution based on route targets.....	948
8.1.4.2	Configuring the route target address family.....	949
8.1.4.3	Originating RT constraint routes.....	949
8.1.4.4	Receiving and re-advertising RT constraint routes.....	949
8.1.4.5	Using RT constraint routes.....	950
8.1.5	BGP fast reroute in a VPRN.....	952
8.1.5.1	BGP fast reroute in a VPRN configuration.....	952
8.2	VPRN features.....	952
8.2.1	IP interfaces.....	952
8.2.1.1	DHCP and DHCPv6.....	953
8.2.2	SAPs.....	955
8.2.2.1	IPv6 support for VPRN IP interfaces.....	955
8.2.2.2	Encapsulations.....	956
8.2.3	QoS policies.....	956
8.2.4	Filter policies.....	956
8.2.4.1	CPU QoS for VPRN interfaces.....	957
8.2.5	CE to PE routing protocols.....	957
8.2.5.1	PE to PE tunneling mechanisms.....	957
8.2.5.2	Per-VRF route limiting.....	958
8.2.6	Exporting MP-BGP VPN routes.....	958
8.2.6.1	Configuration guidelines.....	958
8.2.7	Spoke-SDPs.....	958
8.2.8	Using OSPF in IP-VPNs.....	958
8.2.9	Service label mode of a VPRN.....	959
8.2.10	Multicast in IP-VPN applications.....	959
8.2.10.1	Multicast protocols supported in the provider network.....	960
8.2.10.2	Provider tunnel support.....	962
8.2.11	Inter-AS VPRNs.....	962
8.2.12	Node management using VPRN with GRT leaking.....	964
8.2.12.1	Management with VPRN using GRT leaking.....	965

8.2.12.2	Route leaking from GRT to VPRN instances.....	965
8.3	Configuring a VPRN service with CLI.....	966
8.3.1	Basic configuration.....	966
8.3.2	Common configuration tasks.....	967
8.3.3	Configuring VPRN components.....	968
8.3.3.1	Creating a VPRN service.....	968
8.3.3.2	Configuring global VPRN parameters.....	968
8.3.4	Configuring VPRN protocols - OSPF.....	973
8.3.4.1	VPRN OSPF CLI syntax.....	973
8.3.5	Service management tasks.....	974
8.3.5.1	Modifying VPRN service parameters.....	974
8.3.5.2	Deleting a VPRN service.....	974
8.3.5.3	Disabling a VPRN service.....	975
8.3.5.4	Re-enabling a VPRN service.....	976
8.4	VPRN services command reference.....	976
8.4.1	Command hierarchies.....	976
8.4.1.1	VPRN configuration commands.....	976
8.4.1.2	Show commands.....	990
8.4.1.3	Clear commands.....	992
8.4.1.4	Debug commands.....	992
8.4.2	Command descriptions.....	993
8.4.2.1	VPRN service configuration commands.....	993
8.4.2.2	Show commands.....	1197
8.4.2.3	Clear commands.....	1292
8.4.2.4	Debug commands.....	1299
9	Common CLI command descriptions.....	1306
9.1	Command descriptions.....	1306
9.1.1	SAP syntax.....	1306
	sap.....	1306
10	Appendix: Port-based split horizon.....	1308
10.1	Overview.....	1308
10.1.1	Topology.....	1308
10.2	Configuration guidelines.....	1309
10.2.1	Verification.....	1310

11	Appendix: DHCP management.....	1312
11.1	DHCP principles.....	1312
11.1.1	DHCP features.....	1313
11.1.1.1	Using Option 82 field.....	1313
11.1.1.2	Trusted and untrusted.....	1314
11.1.2	Common configuration guidelines.....	1314
11.1.2.1	Configuration guidelines for DHCP relay and snooping.....	1314
11.1.2.2	Configuring Option 82 handling.....	1315
12	Standards and protocol support.....	1317
12.1	BGP.....	1317
12.2	Certificate Management.....	1319
12.3	Ethernet.....	1319
12.4	EVPN.....	1320
12.5	Fast Reroute.....	1321
12.6	Internet Protocol (IP) — General.....	1321
12.7	IP — Multicast.....	1324
12.8	IP — Version 4.....	1325
12.9	IP — Version 6.....	1326
12.10	IPsec.....	1327
12.11	IS-IS.....	1328
12.12	Management.....	1329
12.13	MPLS — General.....	1332
12.14	MPLS — GMPLS.....	1332
12.15	MPLS — LDP.....	1333
12.16	MPLS — MPLS-TP.....	1333
12.17	MPLS — OAM.....	1334
12.18	MPLS — RSVP-TE.....	1334
12.19	OSPF.....	1334
12.20	Pseudowire.....	1336
12.21	Quality of Service.....	1336
12.22	RIP.....	1337
12.23	Timing.....	1337
12.24	VPLS.....	1338

List of tables

Table 1: Supported modes of operation and configuration methods.....	32
Table 2: Supported port modes by mode of operation.....	34
Table 3: 7210 SAS platforms supporting port modes.....	35
Table 4: Configuration process.....	36
Table 5: Service and SAP encapsulation.....	43
Table 6: SAP and Service (svc-sap-type) combinations for Layer 2 services when in network mode.....	46
Table 7: SAP types in a service when QinQ SAP is in use (network mode operation).....	48
Table 8: SAP and service combinations for 7210 SAS-T in access-uplink mode.....	48
Table 9: ETH-CFM acronym expansions.....	74
Table 10: Defect conditions and priority settings.....	77
Table 11: L2CP support for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp access-uplink and network mode platforms.....	82
Table 12: Keepalive request results.....	124
Table 13: Output fields: customer.....	130
Table 14: Output fields: FDB MAC.....	131
Table 15: Output fields: SDP.....	134
Table 16: Output fields: SDP-using.....	138
Table 17: Output fields: service-using.....	140
Table 18: Output fields: Ethernet ring.....	142
Table 19: Output fields: Ethernet ring status.....	143
Table 20: Output fields: SDP.....	148
Table 21: Output fields: ETH-CFM association.....	152

Table 22: Output fields: CFM stack table.....	154
Table 23: Output Fields: ETH-CFM domain.....	155
Table 24: Output fields: connection profile.....	161
Table 25: MTU values for specific VC types.....	238
Table 26: Final disposition of the packet based on per-FC and per-SAP policer or meter.....	262
Table 27: Output fields: service ID all.....	305
Table 28: Output fields: base.....	308
Table 29: Output fields: service ID endpoint.....	310
Table 30: Output fields: labels.....	312
Table 31: Output fields: service ID SAP.....	320
Table 32: Output fields: service ID SDP.....	324
Table 33: Output fields: Split Horizon group.....	330
Table 34: Output fields: STP.....	335
Table 35: Output fields: SAP-using.....	340
Table 36: Output fields: SDP.....	342
Table 37: Output fields: SDP-using.....	345
Table 38: Output fields: Service-using.....	347
Table 39: EVPN routes and usage.....	358
Table 40: Proxy-ARP entry combinations.....	400
Table 41: Output fields: EVPN MPLS tunnel endpoints.....	460
Table 42: Output fields: service ID BGP-EVPN.....	461
Table 43: Output fields: EVPN MPLS.....	463
Table 44: Output fields: proxy-ARP.....	465

Table 45: Output fields: proxy-ND.....	467
Table 46: Output fields: system BGP-EVPN.....	469
Table 47: Output Fields: BGP-EVPN multi-homing.....	471
Table 48: Routing behavior in R-VPLS and interaction ARP cache and MAC FIB.....	511
Table 49: ACL lookup behavior with ingress override filter attached to an IES interface in an R-VPLS service.....	513
Table 50: ACL lookup behavior without ingress override filter attached to an IES interface in an R-VPLS service.....	514
Table 51: Routing protocols on IP interfaces bound to a VPLS service.....	515
Table 52: SAP BPDU encapsulation states.....	537
Table 53: MTU Values for VC Types (VPLS).....	602
Table 54: Output fields: FDB information.....	688
Table 55: Output fields: FDB MAC.....	690
Table 56: Output fields: ingress label.....	691
Table 57: Output fields: SAP-using.....	693
Table 58: Output fields: SDP.....	694
Table 59: Output fields: SDP-using.....	696
Table 60: Output fields: service-using.....	700
Table 61: Output fields: service ID all.....	713
Table 62: Output Fields: ARP.....	717
Table 63: Output fields: base.....	718
Table 64: Output fields: FDB.....	720
Table 65: Output fields: labels.....	723
Table 66: Output fields: L2PT.....	724

Table 67: Output fields: MSTP configuration.....	726
Table 68: Output fields: service ID SAP.....	732
Table 69: Output fields: service ID SDP.....	735
Table 70: Output fields: STP.....	742
Table 71: Output fields: DHCP statistics.....	745
Table 72: Output fields: DHCP summary.....	747
Table 73: Output fields: DHCPv6 statistics.....	749
Table 74: Output fields: IGMP snooping all.....	754
Table 75: Output fields: MFIB.....	755
Table 76: Output fields: IGMP-snooping Mrouters.....	757
Table 77: Output fields: MVR.....	758
Table 78: Output fields: port database.....	760
Table 79: Output fields: proxy database.....	762
Table 80: Output fields: querier.....	764
Table 81: Output fields: static.....	765
Table 82: Output Fields: IGMP-snooping statistics.....	768
Table 83: Output fields: service ID endpoint.....	770
Table 84: Output fields: ETH-CFM associations.....	806
Table 85: Output fields: ETH-CFM stack table.....	808
Table 86: ETH-CFM domain field descriptions.....	810
Table 87: Output fields: MEP.....	812
Table 88: Output fields: MIP.....	813
Table 89: Output fields: service ID all.....	821

Table 90: Output fields: base.....	825
Table 91: Output fields: FDB.....	828
Table 92: Output fields: STP.....	830
Table 93: Output fields: ISID-using.....	832
Table 94: Output fields: I-VPLS.....	833
Table 95: Output fields: service ID Epipe.....	834
Table 96: Output fields: ISID-using.....	835
Table 97: Output fields: Service-using.....	836
Table 98: Output fields: MAC name.....	837
Table 99: Output Fields: Customer.....	921
Table 100: Output Fields: SAP-using.....	923
Table 101: Output Fields: Service-using.....	924
Table 102: Output Fields: Service ID All.....	932
Table 103: Output Fields: ARP.....	937
Table 104: Output Fields: Base.....	938
Table 105: Output Fields: Service ID Interface.....	940
Table 106: Output Fields: ID Interface on 7210 SAS-Mxp.....	943
Table 107: BGP fast reroute scenarios (VPRN context).....	952
Table 108: Final disposition of the packet based on per-FC and per-SAP policer or meter.....	1089
Table 109: Route type preference defaults.....	1190
Table 110: Output fields: egress label.....	1198
Table 111: Output fields: ingress label.....	1200
Table 112: Output fields: SAP-using.....	1202

Table 113: Output fields: SDP.....	1204
Table 114: Output Fields: SDP-using.....	1207
Table 115: Output Fields: Service-using.....	1210
Table 116: Output fields: service ID All.....	1213
Table 117: Output fields: ARP.....	1220
Table 118: Output fields: base.....	1222
Table 119: Output Fields: ID base BGP PIC.....	1223
Table 120: Output Fields: DHCP statistics.....	1226
Table 121: Output Fields: Interface.....	1229
Table 122: Output Fields: service ID SAP.....	1232
Table 123: Output Fields: service ID SDP.....	1237
Table 124: Output Fields: Aggregate.....	1239
Table 125: Output fields: ARP.....	1241
Table 126: Output fields: damping.....	1245
Table 127: Output fields: group.....	1248
Table 128: Output fields: neighbor.....	1253
Table 129: Output fields: neighbor received routes.....	1258
Table 130: Output fields: show neighbor add-path.....	1260
Table 131: Output fields: paths.....	1265
Table 132: Output fields: routes.....	1269
Table 133: Output fields: BGP summary.....	1272
Table 134: Output fields: interface.....	1275
Table 135: Output fields: interface detail.....	1276

Table 136: Output fields: interface summary.....	1278
Table 137: Output fields: MVPN.....	1280
Table 138: Output Fields: MVPN list.....	1282
Table 139: Output Fields: Router Table.....	1284
Table 140: Output Fields: Static ARP.....	1286
Table 141: Output Fields: Static Route.....	1289
Table 142: Output Fields: Tunnel Table.....	1291
Table 143: SAP-ID formats.....	1306
Table 144: Encapsulation types.....	1307

List of figures

Figure 1: Service entities for 7210 SAS devices configured in network mode.....	40
Figure 2: SAPs for 7210 SAS configured in network mode.....	41
Figure 3: Multiple SAPs in a service using QinQ uplinks in 7210 SAS configured in access-uplink mode...	41
Figure 4: Multiple SAPs on a single port (7210 SAS configured in network mode).....	43
Figure 5: MPLS SDP pointing from ALA-A to ALA-B.....	50
Figure 6: 7210 SAS using Layer 2 uplinks in an Layer 2 network.....	53
Figure 7: G.8032 Ring in the Initial State.....	56
Figure 8: 0 to 1 G.8032 ring in the protecting state.....	57
Figure 9: 0 to 3 resilient ring service example.....	58
Figure 10: 0 to 4 G.8032 Sub-ring.....	60
Figure 11: 0 to 6 sub-ring homed to VPLS.....	62
Figure 12: Service creation and implementation flow.....	67
Figure 13: Ethernet OAM model for Ethernet access – business.....	76
Figure 14: Ethernet OAM model for Ethernet access – wholesale.....	76
Figure 15: Dynamic VLAN assignment using dot1x in enterprise.....	84
Figure 16: Corporate or guest VPLS assignment based on 802.1x authentication outcome.....	86
Figure 17: Epipe/VLL service.....	167
Figure 18: VLL Resilience with pseudowire redundancy and switching.....	173
Figure 19: Pseudowire switching TLV format.....	174
Figure 20: VLL resilience.....	176
Figure 21: Dynamic MS-PW operation.....	178

Figure 22: MS-PW addressing using FEC 129 all Type 2.....	178
Figure 23: Advertisement of PE addresses by PW routing.....	179
Figure 24: Signaling of dynamic MS-PWs using T-LDP.....	179
Figure 25: Mapping of All to SAP.....	180
Figure 26: VLL using dynamic MS-PWs, Inter-AS scenario.....	181
Figure 27: Pseudowire redundancy.....	185
Figure 28: Dynamic MS-PW example.....	187
Figure 29: Master-slave pseudowire redundancy.....	191
Figure 30: VLL resilience.....	193
Figure 31: VLL resilience with pseudowire switching.....	194
Figure 32: Access node resilience using MC-LAG and PW redundancy.....	196
Figure 33: VLL resilience with PW redundancy and PW switching.....	197
Figure 34: Redundant VLL endpoint objects.....	199
Figure 35: Default QinQ SAP for transit traffic in a ring scenario.....	215
Figure 36: SDPs — unidirectional tunnels.....	216
Figure 37: VLL resilience.....	219
Figure 38: VLL resilience with pseudowire switching.....	220
Figure 39: EVPN for MPLS in VPLS services.....	357
Figure 40: EVPN routes type 1 and 4.....	359
Figure 41: EVPN-VPLS integration.....	366
Figure 42: DF election.....	370
Figure 43: Split-horizon.....	371
Figure 44: Aliasing.....	371

Figure 45: ES discovery and DF election.....	373
Figure 46: All-active multi-homing ES failure.....	380
Figure 47: Black hole caused by SAP/SVC shutdown.....	381
Figure 48: Transient issues caused by "slow" MAC learning.....	381
Figure 49: Backup PE.....	383
Figure 50: Single-active multi-homing ES failure.....	387
Figure 51: EVPN-VPWS BGP extensions.....	388
Figure 52: EVPN-MPLS VPWS.....	389
Figure 53: A/S PW and MC-LAG support on EVPN-VPWS.....	391
Figure 54: Link loss forwarding for EVPN-VPWS.....	392
Figure 55: LACP standby signaling from the non-DF.....	394
Figure 56: Proxy-ARP example usage in an EVPN network.....	397
Figure 57: Network topology of an access ring.....	408
Figure 58: VPLS service architecture.....	475
Figure 59: Access port ingress packet format and lookup.....	475
Figure 60: Network port egress packet format and flooding.....	476
Figure 61: Access port ingress packet format and lookup.....	477
Figure 62: Network port egress packet format and flooding.....	477
Figure 63: DHCPv6 snooping for Layer 2 services.....	483
Figure 64: MVR and MVR by proxy.....	486
Figure 65: Access resiliency.....	491
Figure 66: H-VPLS with spoke redundancy.....	495
Figure 67: H-VPLS resiliency based on AS pseudowires.....	496

Figure 68: Dual-homed MTUs in two-tier hierarchy H-VPLS.....	497
Figure 69: H-VPLS with SAP redundancy.....	499
Figure 70: Dual homed CE connection to VPLS.....	500
Figure 71: BGP AD NLRI versus IP VPN NLRI.....	503
Figure 72: Generalized pseudowire-ID FEC element.....	504
Figure 73: BGP-AD and T-LDP interaction.....	505
Figure 74: Secure access to network using MAC filter.....	517
Figure 75: Epipe Emulation in a Ring using VPLS with G.8032.....	521
Figure 76: Example configuration for protected VPLS SAP.....	541
Figure 77: Example configuration for protected VPLS spoke-SDP.....	543
Figure 78: Example configuration for load balancing with management VPLS.....	544
Figure 79: Example Configuration for Load Balancing Across Two Protected VPLS Spoke-SDPs.....	548
Figure 80: Example configuration for load balancing with management VPLS.....	553
Figure 81: BGP AD configuration example.....	557
Figure 82: Sample topology-AS pseudowire in VPLS.....	558
Figure 83: Large H-VPLS deployment.....	787
Figure 84: Large PBB-VPLS Deployment.....	788
Figure 85: QinQ payload in provider header example.....	789
Figure 86: PBB mapping to VPLS constructs.....	790
Figure 87: PBB packet walk-through.....	792
Figure 88: Access dual-homing into PBB BEBs - topology view.....	794
Figure 89: Internet Enhanced Service.....	846
Figure 90: Virtual Private Routed Network.....	945

Figure 91: Route distinguisher.....	946
Figure 92: Directly connected IP target.....	947
Figure 93: Multiple hops to IP target.....	948
Figure 94: Multicast in IP-VPN applications.....	960
Figure 95: Inter-AS Option-A: VRF-to-VRF model.....	963
Figure 96: Inter-AS Option-B.....	963
Figure 97: Option C example.....	964
Figure 98: OSPF areas.....	1182
Figure 99: Split horizon group example.....	1308
Figure 100: IP address assignment with DHCP.....	1312

1 Getting started

This chapter provides process flow information to configure and provision services. It also provides an overview of the document organization and content, and describes the terminology used in this guide.

1.1 About this guide

**Note:**

Unless explicitly noted otherwise, this guide uses 7210 SAS-Dxp to refer to the 7210 SAS-Dxp 12p, 7210 SAS-Dxp 16p, and 7210 SAS-Dxp 24p platforms.

This guide describes the subscriber services support provided by the following 7210 SAS platforms, operating in one of the modes described in [Table 1: Supported modes of operation and configuration methods](#). If multiple modes of operation apply, they are explicitly noted in the topic:

- 7210 SAS-Mxp
- 7210 SAS-Sx/S 1/10GE
- 7210 SAS-Sx 10/100GE
- 7210 SAS-T

See [7210 SAS modes of operation](#) for information about the modes of operation supported by the 7210 SAS product family.

**Note:**

Unless explicitly noted otherwise, the phrase “Supported on all 7210 SAS platforms as described in this document” is used to indicate that the topic and CLI commands apply to all the 7210 SAS platforms in the following list, when operating in the specified modes only:

- **network mode of operation**
7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T.
- **standalone mode of operation**
7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE.
- **standalone-VC mode of operation**
7210 SAS-Sx/S 1/10GE

If the topic and CLI commands are supported on the 7210 SAS-T operating in the access-uplink mode, they are explicitly indicated, where applicable.

1.1.1 Document structure and content

This guide uses the following structure to describe routing protocols and route policies content:

**Note:**

This guide generically covers Release 26.x.Rx content and may include some content that will be released in later maintenance loads. See the *7210 SAS Software Release Notes 26.x.Rx*, part number 3HE 27852 000x TQZZA, for information about features supported in each load of the Release 26.x.Rx software.

- This guide is organized into functional chapters and provides concepts and descriptions of the implementation flow. Each chapter describes a software area and provides CLI syntax and command usage to configure parameters for the functional area.
- Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.
- Unless explicitly noted, the CLI commands and their configuration is similar for both [network](#) and [access-uplink](#) operating modes for features applicable to both modes of operation.

1.2 7210 SAS modes of operation

Unless explicitly noted, the phrase “mode of operation” and “operating mode” refers to the current operating mode of the 7210 SAS router. Each operating mode provides configuration access to a specific group of CLI commands.

**Note:**

Not all CLI commands are supported on all 7210 SAS platforms in all modes of operation. Users can only configure CLI commands supported by the current operating mode of the router. See the *7210 SAS Software Release Notes 26.x.Rx*, part number 3HE 27852 000x TQZZA, and to the appropriate 7210 SAS software user guide for information about features and capabilities supported by a 7210 SAS platform when operating in a specific mode.

The following modes of operation are supported by the 7210 SAS product family:

- **access-uplink**

In the access-uplink operating mode, the 7210 SAS router uplinks to the network using Layer 2 Ethernet VLAN switching (without IP/MPLS).

Platforms Supported: 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C, and 7210 SAS-T.

- **network**

In the network operating mode, the 7210 SAS router uses IP/MPLS uplinks to the network. The IP routing protocols and MPLS functionality is available; see the appropriate 7210 SAS software user guide for more information about supported features.

Platforms Supported: 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C, 7210 SAS-Mxp, 7210 SAS-R6, 7210 SAS-R12, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T.

- **satellite**

In the satellite operating mode, the 7210 SAS platform uses high-capacity uplinks (for example, 10GE ports on the 7210 SAS-Mxp and 100GE ports on the 7210 SAS-Sx 10/100GE) to connect to the 7750 SR host. The 7210 SAS router is managed by the 7750 SR host. There is no direct CLI access to the satellite node, and all services and protocols are configured on the host.

See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for boot options to configure the [satellite](#) mode of operation on the router. See the 7750 SR software user guides for information about service and protocol provisioning, and operating the 7210 SAS router in [satellite](#) mode.

Platforms Supported: 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE.

- **standalone**

In the standalone operating mode, the 7210 SAS platform supports IP/MPLS uplinks. It is operated and managed independently.

The functionality and features available on the standalone 7210 SAS platform are similar to the [network](#) operating mode. The standalone mode is primarily used to differentiate between a node being managed by the 7750 SR host (in the [satellite](#) operating mode), and a node managed independently (standalone operating mode).

Platforms Supported: 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE.

- **standalone-VC**

In the standalone-VC operating mode, a set of 7210 SAS devices are stacked to provide larger 1GE/10GE port density and control-plane redundancy. The stack of nodes is provisioned and managed as a single chassis, and not as individual nodes.

The functionality and features available on the 7210 SAS platform are similar to the [network](#) operating mode, with additional capabilities, such as control-plane redundancy with non-stop routing and non-stop services.

Platforms Supported: 7210 SAS-Sx/S 1/10GE

For 7210 SAS platforms that support multiple explicit modes of operation ([Table 1: Supported modes of operation and configuration methods](#)), the operating mode must be configured in the Boot Option File (BOF) to ensure that the router boots up in the specified mode. For example, the 7210 SAS-T supports access-uplink and network modes of operation, and the 7210 SAS-Sx/S 1/10GE supports satellite, standalone, and standalone-VC mode of operations. In some cases, the 7210 SAS router operates in a specific mode implicitly, and explicit configuration is not required.

See the appropriate *Basic System Configuration Guide* for boot options and information about how to boot the 7210 SAS platform in a specific operating mode.

The following table lists the supported modes of operation and the configuration methods for the 7210 SAS platforms. Unless explicitly noted otherwise, the operating mode is supported on all variants of the specific 7210 SAS platform.

Table 1: Supported modes of operation and configuration methods

7210 SAS platform	Mode of operation and configuration method				
	Network	Access-uplink	Standalone	Standalone-VC	Satellite
7210 SAS-D		Implicit	Implicit		
7210 SAS-Dxp		Implicit	Implicit		
7210 SAS-K 2F1C2T		Implicit	Implicit		

7210 SAS platform	Mode of operation and configuration method				
	Network	Access-uplink	Standalone	Standalone-VC	Satellite
7210 SAS-K 2F6C4T ¹	Port Mode Configuration ²	Port Mode Configuration ²	Implicit		
7210 SAS-K 3SFP+ 8C ¹	Port Mode Configuration ²	Port Mode Configuration ²	Implicit		
7210 SAS-Mxp	Implicit ³		Explicit BOF Configuration		Explicit BOF Configuration
7210 SAS-R6 ⁴	Implicit		Implicit		
7210 SAS-R12 ⁴	Implicit		Implicit		
7210 SAS-Sx/S 1/10GE	Implicit ³		Explicit BOF Configuration	Explicit BOF Configuration	Explicit BOF Configuration
7210 SAS-Sx 10/100GE	Implicit ³		Explicit BOF Configuration		Explicit BOF Configuration
7210 SAS-T	Explicit BOF Configuration	Explicit BOF Configuration	Implicit		

1.3 7210 SAS port modes

Unless explicitly noted, the phrase “port mode” refers to the current port configuration of the 7210 SAS node. The 7210 SAS platform supports the configuration of the following port modes:

- **access port mode**

Access ports are configured for customer-facing traffic if Service Access Points (SAPs) are required. The appropriate encapsulation type must be configured to distinguish the services on the port; services are configured on the port based on the encapsulation value.

Access ports can be configured on all the 7210 SAS platforms.

- **access-uplink port mode**

Access-uplink ports provide native Ethernet connectivity in service provider transport or in an infrastructure network. With this option, the encap-type can be configured to only QinQ. Access-uplink

-
- ¹ By default, the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C boot up in the [network](#) mode of operation. These platforms also allow the use of [access-uplink port mode](#) (without explicit BOF configuration), which provides the option to use Layer 2 uplinks instead of IP/MPLS uplinks to the network core, similar to the 7210 SAS-K 2F1C2T router.
- ² See section [7210 SAS port modes](#) for information about port mode configuration
- ³ Implicitly operates in [network](#) mode when [standalone](#) mode of operation is configured
- ⁴ Supports MPLS uplinks only and implicitly operates in network mode

SAPs, which are QinQ SAPs, can only be configured on an access-uplink port to allow the operator to differentiate multiple services being carried over a single uplink port.

This is the default port mode of a 7210 SAS node in the [access-uplink](#) mode of operation.

- **network port mode**

Network ports are configured for network-facing traffic in the service provider transport or infrastructure network, and provide IP/MPLS uplinks.

This is the default port mode of a 7210 SAS node in the [network](#) or [standalone](#) mode of operation.

- **hybrid port mode**

Hybrid ports are configured for access and network facing traffic, and allow a single port to operate in both access and network modes.

Port modes available for configuration on a 7210 SAS node are determined by the current mode of operation of the router.



Note:

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C are unique; all port modes listed in [Table 2: Supported port modes by mode of operation](#) are available for configuration on the router, regardless of the current mode of operation.

The following table lists the port mode configuration support per 7210 SAS mode of operation.

Table 2: Supported port modes by mode of operation

Mode of operation	Supported port mode			
	Access	Network	Hybrid	Access-uplink
Access-uplink	✓			✓
Network	✓	✓	✓	
Satellite ⁵				
Standalone	✓	✓	✓	
Standalone-VC	✓	✓	✓	

The following table lists the port mode configuration supported by the 7210 SAS product family. See the appropriate *Interface Configuration Guide* for more information about configuring the port modes for a specific platform.

⁵ Port modes are configured on the 7750 SR host and managed by the host.

Table 3: 7210 SAS platforms supporting port modes

Platform	Port mode			
	Access	Network	Hybrid	Access-uplink
7210 SAS-D	Yes	No	No	Yes
7210 SAS-Dxp	Yes	No	No	Yes
7210 SAS-K 2F1C2T	Yes	No	No	Yes
7210 SAS-K 2F6C4T	Yes	Yes	Yes	Yes
7210 SAS-K 3SFP+ 8C	Yes	Yes	Yes	Yes
7210 SAS-Mxp	Yes	Yes	Yes	No
7210 SAS-R6 IMM-b (IMMv2)	Yes	Yes	Yes	No
7210 SAS-R6 IMM-c 100GE (IMM-c 1QSFP28)	Yes	Yes	Yes	No
7210 SAS-R12 IMM-b	Yes	Yes	Yes	No
7210 SAS-R12 IMM-c 100GE (IMM-c 1QSFP28)	Yes	Yes	Yes	No
7210 SAS-Sx/S 1/10GE	Yes	Yes	Yes	No
7210 SAS-Sx 10/100GE	Yes	Yes	Yes	No
7210 SAS-T	Yes	Yes ⁶	Yes ⁷	Yes ⁸

1.4 7210 SAS services configuration process

The following table lists the tasks necessary to configure subscriber services and configure mirroring.

This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

⁶ Network ports are supported only if the node is operating in network mode.

⁷ Hybrid ports are supported only if the node is operating in network mode.

⁸ Access-uplink ports are supported only if the node is operating in access-uplink mode.

Table 4: Configuration process

Area	Task	Chapter
Subscribers	Subscriber services	
	Global entities	Configuring global service entities with CLI
	VLL services	VLL services
	VPLS service	Virtual Private LAN Service
	IEEE 802.1ah Provider Backbone Bridging	IEEE 802.1ah Provider Backbone Bridging (PBB)
	IES service	Internet Enhanced Service
	VPRN service	Virtual Private Routed Network service
Reference	List of IEEE, IETF, and other proprietary entities.	Standards and protocol support

1.5 Conventions

This section describes the general conventions used in this guide.

1.5.1 Precautionary and information messages

The following information symbols are used in the documentation.



DANGER: Danger warns that the described activity or situation may result in serious personal injury or death. An electric shock hazard could exist. Before you begin work on this equipment, be aware of hazards involving electrical circuitry, be familiar with networking environments, and implement accident prevention procedures.



WARNING: Warning indicates that the described activity or situation may, or will, cause equipment damage, serious performance problems, or loss of data.



Caution: Caution indicates that the described activity or situation may reduce your component or system performance.



Note: Note provides additional operational information.



Tip: Tip provides suggestions for use or best practices.

1.5.2 Options or substeps in procedures and sequential workflows

Options in a procedure or a sequential workflow are indicated by a bulleted list. In the following example, at step 1, the user must perform the described action. At step 2, the user must perform one of the listed options to complete the step.

Example: Options in a procedure

1. User must perform this step.
2. This step offers three options. User must perform one option to complete this step:
 - This is one option.
 - This is another option.
 - This is yet another option.

Substeps in a procedure or a sequential workflow are indicated by letters. In the following example, at step 1, the user must perform the described action. At step 2, the user must perform two substeps (a. and b.) to complete the step.

Example: Substeps in a procedure

1. User must perform this step.
2. User must perform all substeps to complete this action:
 - a. This is one substep.
 - b. This is another substep.

2 Services overview

This chapter provides an overview of the 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE subscriber services, service model, and service entities. Additional information about the individual subscriber services is described in subsequent chapters.

2.1 Introduction

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

In 7210 SAS-series routers, services can provide Layer 2/bridged service between a service access point (SAP) on one router and another service access point (a SAP is where traffic enters and exits the service) on the same (local) router or another router (distributed). A distributed service spans more than one router.

Distributed services use service distribution points (SDPs) to direct traffic to another 7210 SAS or SR router or other router that supports MPLS, through a service tunnel. SDPs are created on each participating router, specifying the origination address (the router participating in the service communication) and the destination address of another router. SDPs are then bound to a specific customer service. Without the binding process, the far-end router is not able to participate in the service (there is no service without associating an SDP with the service).



Note:

SDPs are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode. Only local services can be configured on 7210 SAS platforms operating in access-uplink mode.

2.1.1 Service types

The 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx 1/10GE: standalone and standalone-VC, and 7210 SAS-Sx 10/100GE offer the following types of subscriber services, which are described in more detail in the following chapters:

- **Virtual Leased Line (VLL) services:**
 - **Ethernet pipe (Epipe)**
A Layer 2 point-to-point VLL service for Ethernet frames. See [Epipe services](#).
- **Virtual Private LAN Service (VPLS)**
A Layer 2 multipoint-to-multipoint VPN. See [Virtual Private LAN Service](#).
- **Internet Enhanced Service (IES)**

A routed connectivity service used to provide IP services. This service is supported in 7210 SAS platforms operated in access-uplink mode for only inband management of the node (that is, it cannot be used for configuring customer service in access-uplink mode). See [Internet Enhanced Service](#).

- **Virtual Private Routed Network (VPRN)**

A Layer 3 IP multipoint-to-multipoint VPN service as defined in RFC 2547bis. See [Virtual Private Routed Network service](#).

2.1.2 Service policies

Common to all 7210 SAS-series connectivity services are policies that are assigned to the service. Policies are defined at a global level, then applied to a service on the router. Policies are used to define 7210 SAS-series service enhancements. The types of policies that are common to all 7210 SAS-series connectivity services, and their functions, are the following:

- 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 ingress policy applied to a SAP specifies the number of meters, meter characteristics (such as forwarding class, committed, and peak information rates, and so on) and the mapping of traffic to a forwarding class. A QoS egress policy defines the queue characteristics (such as CBS, CIR, PIR). A QoS policy must be created before it can be applied to a SAP. A single ingress and egress QoS policy can be associated with a SAP. A single access egress QoS policy can be associated with a port.
- Filter policies allow selective blocking of traffic matching criteria from ingressing or egressing a SAP. Filter policies, also referred to as access control lists (ACLs), control the traffic allowed in or out of a SAP, based on MAC or IP match criteria. Associating a filter policy with a SAP is optional. Filter policies are identified by a unique filter policy ID. A filter policy must be created before it can be applied to a SAP. A single ingress and single egress filter policy can be associated with a SAP.
- Scheduler policies define the operating parameters (such as scheduling algorithm, weights per priority). Depending on the platform, these are either associated with SAPs or physical ports.
- Accounting policies define how to count the traffic usage for a service, for billing purposes.

The routers provide a comprehensive set of service-related counters. Accounting data can be collected on a per-service, per-forwarding class basis, 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.

2.2 Nokia service model

In the Nokia service model, the service edge routers are deployed at the provider edge. Services are provisioned on the service routers and transported across an IP or IP/MPLS provider core network in encapsulation tunnels created using MPLS label switched paths (LSPs).

The 7210 SAS devices configured in access-uplink mode support QinQ/dot1q Layer 2 uplinks to transport the services to the provider edge in a hierarchical configuration, whereas 7210 SAS devices configured in network mode support MPLS uplinks to transport the services.

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


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

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

2.3 Service entities

The following figure shows an example of the basic logical entities in the service model used to construct a service for a 7210 SAS device configured in network mode.

Figure 1: Service entities for 7210 SAS devices configured in network mode



2.3.1 Customers

The terms "customer" and "subscriber" are used synonymously. The most basic required entity is the customer ID value, 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.

2.3.2 SAPs

Each subscriber service type is configured with at least one SAP. A SAP identifies the customer interface point for a service on a 7210 SAS router ([Figure 2: SAPs for 7210 SAS configured in network mode](#)). The SAP configuration requires that slot, MDA, and port information be specified. The slot, MDA, and port

parameters must be configured before provisioning a service (see the Cards, MDAs, and Ports sections of the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide*).


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

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

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


The following figure shows SAPs used for customer service delivery, with SDP used for service transport on 7210 SAS devices that support MPLS uplinks (also known as network mode platforms).

Figure 2: SAPs for 7210 SAS configured in network mode



The following figure shows multiple SAPs used for customer service delivery, with access-uplink SAPs (also known as QinQ SAPs) used for service transport on 7210 SAS devices that support only Layer 2 uplinks (also known as access-uplink mode platforms).

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



2.3.2.1 SAP encapsulation types and identifiers

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

2.3.2.2 Ethernet encapsulations

The following are the encapsulation service options on Ethernet ports:

- **null**

Supports a single service on the port. For example, where a single customer with a single service customer edge (CE) device is attached to the port. The encapsulation ID is always 0 (zero).

- **dot1q**

Supports multiple services for one customer or services for multiple customers ([Figure 4: Multiple SAPs on a single port \(7210 SAS configured in network mode\)](#)). The encapsulation ID used to distinguish an individual service is the VLAN ID in the IEEE 802.1Q header. For example, the port is connected to a Ethernet switch with multiple downstream customers.

- **QinQ**

The QinQ encapsulation type adds a IEEE 802.1Q tag to the 802.1Q tagged packets entering the network, to expand the VLAN space by tagging tagged packets, producing a double-tagged frame. The 7210 SAS supports QinQ encapsulation for access ports in network mode. In access-uplink mode, QinQ encapsulation is supported for both access ports and access-uplink ports.


The following are the encapsulation service options on Ethernet access-uplink ports:

- **QinQ**

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

The following figure shows multiple SAPs used for customer service delivery on the same port and belonging to the same service, along with SDP used for service transport on 7210 SAS devices that support MPLS uplinks (also known as network mode platforms). This is supported only in network mode.

Figure 4: Multiple SAPs on a single port (7210 SAS configured in network mode)



OSSG003-7210M

2.3.2.3 Services and SAP encapsulations

The following table lists the service and SAP encapsulation information for Ethernet ports.

Table 5: Service and SAP encapsulation

Port type	Encapsulation
Ethernet	Null
Ethernet	Dot1q
Ethernet	QinQ

2.3.2.4 Default SAP on a dot1q port

This feature provides default SAP functionality on dot1q-encapsulated ports. On a dot1q-encapsulated port where a default SAP is configured, all packets with q-tags not matching any explicitly defined SAPs will be assigned to this SAP. SAPs with default dot1q encapsulation are supported in VPLS and Epipe services. Dot1q default SAPs are not supported in VPRNs.

In this context, the character "*" indicates default, which means allow through. The default SAP also accepts untagged or priority-tagged packets. A default SAP must be configured explicitly. When a default SAP is not configured explicitly, packets not matching any explicitly defined SAPs will be dropped.

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

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

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

- This type of SAP is supported only on VPLS and Epipe services, and cannot be created in IES and VPRN services because it cannot preserve VLAN tag markings.
- For VPLS SAPs with STP enabled, STP listens to untagged and null-tagged BPDUs only. All other tagged BPDUs are forwarded like other customer packets. This is the same behavior as null-encapsulated ports.
- This type of SAP is mutually exclusive of a SAP defined by explicit null encapsulation (for example, 1/1/1:0). This avoids conflict as with which SAP untagged frames should be associated.
- IGMP snooping is not supported on a default SAP. This would require remembering VLAN tags per hosts. By not allowing IGMP snooping on this SAP, all IGMP packets will be transparently forwarded.

2.3.2.5 Default SAPs on a QinQ port (supported only on 7210 SAS devices configured in access-uplink mode)

Default QinQ SAPs (notation *.*) are used in ring ports to avoid the need to configure services on all the intermediate nodes in the ring that are transiting the service. Default QinQ SAPs match all VLAN-tagged traffic that is not classified into any other SAP configured on the same port. Only one Epipe service with default QinQ SAPs is needed for transit service traffic on access-uplink ports.

Default QinQ SAPs are only allowed on access-uplink ports and access ports. A default QinQ SAP can coexist with a 0.* SAP on an access-uplink or access port. A default QinQ SAP accepts only tagged packets. Untagged packets or priority-tagged packets are not accepted on default QinQ SAPs.

When an Epipe service with default QinQ SAPs on the ring ports is used for transit traffic in a ring deployment, no protection mechanism (for example, STP or G.8032) is supported for default QinQ SAPs. The upstream or head-end node on which the service originates must ensure that the correct path on the ring is selected using either G.8032 or STP.

When a VPLS service with default QinQ SAPs on the ring ports is used for transit traffic in a ring deployment, users can use either G.8032 or M-VPLS with xSTP for ring protection. When using G.8032, the state of the default QinQ SAPs in the VPLS service can be managed using a separate G.8032 control instance.



Note:

A G.8032 control instance cannot use default QinQ SAPs.

A default QinQ SAP is available for use only in an Epipe and a VPLS service created with the **svc-sap-type** parameter set to **null-star**. A default QinQ SAP can be configured along with other SAPs allowed in the same service (that is, service with the **svc-sap-type** parameter set to **null-star**).

The following features are available for use with default QinQ SAPs configured in Epipe and VPLS service (unless explicitly specified, the following features are applicable for both Epipe and VPLS service).

The following features are available for default QinQ SAPs on either access ports or access-uplink ports:

- MAC learning and aging is available for use in a VPLS service.
- Per-SAP MAC limit is available for use in a VPLS service.
- Mac-move detection and Mac-pinning is available for use in a VPLS service.
- Discard-unknown and discard-unknown-source is available for use in a VPLS service.
- Eth-CFM and Y.1731 is not available for use.
- STP (and all its different flavors) cannot be enabled in the service with default QinQ SAPs.

- MVPLS with xSTP can be used for loop prevention. The default QinQ SAPs inherit the state from the associated MVPLS instance.
- A G.8032 control instance cannot be configured in a service with a default QinQ SAP.
- G.8032 can be used for loop prevention in ring deployments, where the default QinQ SAPs are configured on the ring ports in a VPLS service. A separate G.8032 control instance must be configured for use on the ring ports, and the service with default QinQ ports needs to be associated with this G.8032 control instance.
- IGMP snooping is not available for use.
- L2PT and BPDU translation is not available for use.
- IP interface in a VPLS service is not supported in a service using this SAP.

The following features are available for default QinQ SAPs created on access-uplink port:

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

The following features are available for default QinQ SAPs created on access ports:

- SAP ingress QoS policy is available for use.
- Egress QoS policy applied on an access port is available for egress shaping, scheduling, and marking.
- SAP ingress ACLs are available for use.
- SAP egress ACLs are not available for use.
- SAP ingress meter counters, SAP ingress received counters, and SAP egress forwarded counters are available for use (appropriate accounting records can be used).

2.3.2.5.1 Configuration notes for use of default QinQ SAPs for transit service in a ring deployment

The following considerations apply to default QinQ SAPs for transit service in a ring deployment configurations:

- If an Epipe service is used with default QinQ SAPs on the ring ports for transit service in a ring deployment, no protection mechanism is available for the transit service (that is, Epipe service with the default QinQ SAPs on ring ports). Both Epipe and VPLS services that are originating on different nodes in the ring can use the transit service.

Protection and loop-detection mechanisms can be implemented for VPLS service configured in the ring nodes, by using MVPLS with xSTP on the nodes where the VPLS service is configured. No protection mechanisms are available for use with Epipe services on the node that originates the service.

- If a VPLS service is used with default QinQ SAPs on the ring ports for transit service in a ring deployment, either MVPLS/xSTP or G.8032 can be used to protect the transit service (that is, VPLS service with the default QinQ SAPs on ring ports). In this case, VPLS services that are originating on

different nodes in the ring and use the transit VPLS service are also protected. Epipe services that are originating on different nodes in the ring cannot use the transit VPLS service.

- When using VPLS service with default QinQ SAPs for transit service with either G.8032 or MVPLS with xSTP configured for protection, load-balancing of the traffic based on the VLAN IDs is not possible. If load-balancing is needed, it is better to use Epipe service with default QinQ SAPs as the transit service.

2.3.2.6 SAP configuration considerations for network mode and access-uplink mode

The following considerations apply to network mode and access-uplink mode SAP configurations:

- A SAP is a local entity and only locally unique to a specific device. The same SAP ID value can be used on another 7210 SAS-series router.
- By default, no SAPs are configured on the node. All SAPs in subscriber services must be created.
- At creation, the default administrative state for a SAP is set to 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 in each router.
- On a port with a dot1q encapsulation type, 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 added at SAP egress. As a result, VLAN IDs only have local significance, and configuring identical VLAN IDs for each SAP on a service is not required.
- If a port is administratively shut down, all SAPs on that port are operationally out of service.
- QinQ access SAPs of type Q1.0 are supported only for IES, VPRN, and R-VPLS services. They are not supported for Layer 2 services.
- A SAP cannot be deleted until it has been administratively disabled (shutdown).
- Each SAP can have one each of the following policies assigned:
 - Ingress filter policy
 - Egress filter policy
 - Ingress QoS policy
 - Accounting policy

2.3.2.7 SAP configuration notes when operating 7210 SAS devices in network mode only

When provisioned in network mode, the following SAP configuration guidelines are applicable.

The following table describes SAP and service (**svc-sap-type**) combinations for Layer 2 services (when in network mode).

Table 6: SAP and Service (svc-sap-type) combinations for Layer 2 services when in network mode

svc-sap-type (Layer 2 service)	Access SAPs	SDP bindings
Any (without any QinQ SAP in the same service)	Null SAP, dot1q SAP, dot1q default SAP, dot1q explicit null SAP, Q1.*	PW with vc-type set to vc-ether or vc-vlan

svc-sap-type (Layer 2 service)	Access SAPs	SDP bindings
	SAP, 0.* SAP (accepts only untagged frames)	
Any (with QinQ SAP in the same service)	Null SAP (accepts only untagged frames), Dot1q SAP (accepts only a single tag frame, Dot1q explicit null SAP (accepts only untagged frame), Q1.Q2 SAP (accepts frames with only two tags), 0.* SAP (accepts only untagged frames)	PW with vc-type set to vc-ether or vc-vlan (accepts only untagged frames on vc-ether PW and only tagged frames on vc-vlan PW)
Any (with dot1q-range SAP)	Dot1q range SAP, Q1.* SAP	PW with vc-type set to vc-ether or vc-vlan (no checks on VLAN ID for packets received on PW)
Qinq-inner-tag-preserve	Q1.Q2 SAP (inner tag Q1 is not stripped), dot1q SAP (no tags are stripped)	PW with vc-vlan (vc-vlan is not stripped)

2.3.2.8 QinQ SAP Configuration Restrictions for 7210 SAS platforms in network operating mode

The following are the QinQ access SAP configuration guidelines for 7210 SAS in network mode only.

These guidelines are not applicable when 7210 SAS devices are operating in access-uplink mode and access-uplink SAPs are in use:

- Tagged packets received on SAPs configured in a service in which a QinQ SAP is also in use are processed (not applicable when a QinQ SAP is not provisioned in a service).
- When a QinQ SAP is configured in a service, the number of VLAN tags in the packets received on null SAP, dot1q SAP, and QinQ SAP configured in the same service should match the number of VLAN tags implied by the port encapsulation mode. Packets that do not match are dropped by the hardware. That is, packets received with more than two VLAN tags on a QinQ SAP are dropped, packets received with more than one VLAN tag on a Dot1q SAP are dropped, and packets received with tags (even packets with a priority tag) on a null SAP are dropped. In this document, such packets are referred to as extra-tag packets.
- When a QinQ SAP is configured in a service, the number of VLAN tags in the packets received on the VC/pseudowire of type vc-vlan should be exactly one and packets received on the VC/pseudowire of type vc-ether should contain no tags (not even priority tags). If either case, packets that contain more VLAN tags than the number specified previously are dropped. In this document, such packets are referred to as extra-tag packets.
- The system will provide a limited number of counters for the extra-tag packets dropped on SAP ingress. These counters are intended for diagnostic use.

The following table shows the SAP types allowed in a service when QinQ SAP is in use.

Table 7: SAP types in a service when QinQ SAP is in use (network mode operation)

SAP configured in the service	SAPs not allowed for configuration in the same service
QinQ	Q.* SAP, dot1q default SAP
Q.*	Q1.Q2
Dotq1 default SAP	Q1.Q2

A 0.* QinQ SAP configured in the service will only accept untagged or priority-tagged packets, regardless of whether a QinQ SAP is configured in the service.

**Note:**

The 7210 SAS platforms support a mechanism to transport QinQ packets in an Epipe with two or more tags, with some restrictions. For more information, see [Epipe services](#).

2.3.2.9 SAP configuration notes for 7210 SAS platforms in access-uplink operating mode

When provisioned in access-uplink mode, the following SAP configuration guidelines are applicable.

The following table describes SAP and service combinations allowed in access-uplink mode.

Table 8: SAP and service combinations for 7210 SAS-T in access-uplink mode

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

The following guidelines apply to SAPs:

- The **svc-sap-type** parameter value determines the type of SAPs that are allowed to be provisioned in a service.
- A physical port can only have one SAP to be part of one service. Multiple SAPs can be defined over a physical port but each of these SAPs must belong to a different service.

- If a service sap-type is specified as **dot1q-preserve**, all the SAPs configured in the service must have the same VLAN ID. The outermost VLAN tag of the packets received on the access port is not stripped, when **svc-sap-type** is set to dot1q-preserve.
- A dot1q default SAP cannot be configured when svc-sap-type is set to **any**.
- When **svc-sap-type** is set to **any** for a null SAP, the system processes and forwards only packets with no VLAN tag (that is, untagged). All other packets with one or more VLAN tags (even those with priority tag only) are not processed and are dropped. Users can use the service with **svc-sap-type** set to **null-star**, to process and forward packets with one or more tags (including priority tag) on a null SAP.
- An ingress QoS policy and accounting policy is assigned per access-uplink port and cannot be assigned per access-uplink SAP.
- The default QinQ SAP processes only tagged packets received on a QinQ port. All tagged packets that do not match the specific SAP tags configured on the same port are processed by this SAP. The default QinQ SAP cannot process un-tagged packets, even if 0.* SAP is not configured for use on that port.
- The default QinQ SAP is available for use with 0.* SAPs configured on the same port or in the same service. It is available for use with another default QinQ SAP configured in the same service (on a different port). In a VPLS service, the default QinQ SAP is available for use with any other SAP type configured in a service configured with **svc-sap-type** parameter set to **null-star**.
- SAPs using connection-profile (to specify dot1q VLAN ranges or individual VLAN IDs) can be configured in a service only when **svc-sap-type** is set to **dot1q-range**.
- When a service is configured to use **svc-sap-type dot1q-range**, the outermost V-LAN tag of the packets is not stripped when the packet is received on access port ingress. See [Epipe services](#) for more information about the processing behavior for this type of service.

2.3.3 SDPs



Note:

SDPs are supported by all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

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

An SDP has the following characteristics:


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

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

2.3.3.1 SDP binding

The following figure shows an example SDP binding. To configure a distributed service from ALA-A to ALA-B, the SDP ID must be specified in the service creation process to bind the service to the tunnel (the SDP). Otherwise, service traffic is not directed to a far-end point and the far-end devices cannot participate in the service (there is no service). To configure a distributed service from ALA-B to ALA-A, the SDP ID must be specified.

Figure 5: MPLS SDP pointing from ALA-A to ALA-B



2.3.3.2 Spoke and mesh SDPs



Note:

SDPs are supported by all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

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

A spoke-SDP is treated like the equivalent of a traditional bridge port where flooded traffic received on the spoke-SDP is replicated on all other ports and not transmitted on the port it was received.

All mesh SDPs bound to a service are logically treated like a single bridge port for flooded traffic where flooded traffic received on any mesh SDP on the service is replicated to other ports (spoke-SDPs and SAPs) and not transmitted on any mesh SDPs.

2.3.3.3 SDP using BGP route tunnel

SDPs are enhanced to use BGP route tunnels to extend inter-AS support for Layer 2 and Layer 3 VPN services. An SDP can be configured to use the MPLS transport method. MPLS SDP support is enhanced to allow a BGP route tunnel to reach the far-end PE. A single method of tunneling is allowed per SDP (for example, LDP, RSVP-TE LSP, or BGP route tunnel). The BGP route tunnel method is excluded if multimode transport is enabled for an SDP.

For 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 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 entered to select an LSP to reach the PE node from the ASBR node. On the last BGP segment, both BGP+LDP and LDP routes may be available to reach the far-end PE from the ASBR node. An LDP LSP must be preferred because of higher protocol priority. This leads to just one label, besides other labels in the stack to identify the VC/VPN at far-end PE nodes.

2.3.3.4 SDP keepalives

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

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

- admin up/admin down state
- hello time
- message length
- max drop count
- hold down time

SDP keepalive echo request messages are only sent when the SDP is completely configured and administratively up and SDP keepalives 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. If max drop count echo request messages do not receive an echo reply, the SDP will immediately be brought operationally down.

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

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

See [Configuring an SDP](#) for information about configuring keepalive parameters.

2.3.3.5 Mixed-LSP mode of operation

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

An LDP LSP can be configured as a primary LSP type, which can then be backed up by a BGP LSP type.

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

1. RSVP LSP type

One RSVP LSP can be configured per SDP. This is the highest priority LSP type.

2. LDP LSP type

One LDP FEC will be used per SDP. The 7210 SAS does not support LDP ECMP.

3. BGP LSP type

One RFC 3107-labeled BGP prefix programmed by the service manager.

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

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



Note:

LDP uses a tunnel down damp timer that is set to three seconds by default. If the LDP LSP fails, the SDP reverts to the RSVP LSP type after the timer expires. For an immediate switchover, the timer should be set to zero using the **configure>router>ldp>tunnel-down-damp-time** command. See the *7210 SAS-Mxp, R6, R12, S, Sx, T MPLS Guide* for more information.

If the value of the **revert-time** timer is changed, it comes into effect only at the next use of the timer. Any timer that is outstanding at the time of the change is restarted with the new value.

In the case of the LDP/BGP SDP, the service manager prefers the LDP LSP type over the BGP LSP type. The service manager reprograms the line card with the BGP LSP, if available; otherwise, the SDP is placed in the operationally down state.



Note:

An LDP/BGP SDP has differences in behavior compared to an RSVP/LDP SDP. For a specific /32 prefix, only a single route will exist in the routing table: the IGP route or the BGP route. Therefore, either the LDP FEC or the BGP label route is active at any time. As a result, the tunnel table needs to be reprogrammed each time a route is deactivated and another is activated. Also, the SDP revert-time cannot be used, because there is no situation where both LSP types are active for the same /32 prefix.

2.3.4 SAP and service scaling with high SAP scale mode



Note:

This feature is supported only on the 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE (standalone).

In Layer 2 access networks used to backhaul service traffic from business services, mobile backhaul, and residential services, the 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE platforms act as Layer 2 carrier Ethernet switching platforms with VLAN-based Layer 2 uplinks. To perform this role, these platforms must support higher SAP and service scaling. To do so, these platforms use the SAP scale mode and port-based access ingress policies.




Note:

On the 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE, **sap-scale-mode** can be configured in standalone mode and does not require a BOF configuration.

The following figure shows the use of Layer 2 uplinks in a Layer 2 access network with a 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, or 7210 SAS-Sx 10/100GE.

Figure 6: 7210 SAS using Layer 2 uplinks in an Layer 2 network



The SAP scale mode is configured using the **configure>system>resource-profile>sap-scale-mode {high | low}** command. By default, the **low** option is configured for low SAP scale mode, which provides backward compatibility. The user can configure the **high** option to use high SAP scale mode, which allows the configuration of a higher number of services and SAPs. Before changing the **sap-scale-mode** value, the user must perform the following:

- Remove all service and SAP configurations.
- Change the value of **sap-scale-mode** and enable per-port egress queuing using the **configure>system>resource-profile>qos>port-scheduler-mode** command.
- Reboot the node.
- Reconfigure all SAPs and services as required.

In high SAP scale mode, the system supports higher SAP and service scaling for Epipe/VLL and VPLS services only. SAP and service scaling for IES, VPRN, and RVPLS services remain unchanged.

QoS policies support port-based access ingress policies on access ports to facilitate the use of access ports as Layer 2 uplinks. With the use of ports as Layer 2 uplinks, the user can apply a single port-based access ingress policy at ingress of an access port, instead of using per-SAP ingress policies. This allows a single policy definition to be used to classify and rate-limit all traffic received over access ports used as Layer 2 uplinks (similar to a network port-based policy applied to network ports used as uplinks) instead of using per SAP ingress policies. Resources must be allocated using the **configure>system>resource-profile** command to use access ingress QoS policies on an access port.

In addition, only the following QoS policies can be used in the high SAP scale mode to achieve a higher scale:

- access port-based egress queuing and shaping on all ports, including service delivery ports and uplinks
- access port-based ingress classification and policing on uplinks
- Epipe and VPLS SAPs using ingress table-based classification and policing on service delivery ports for higher SAP scale
- IES and VPRN SAPs using table-based classification or CAM-based classification
- RVPLS SAPs using CAM-based classification and policing
- default SAP ingress QoS policy of 65536 for the 7210 SAS-Mxp and 1 for the 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE.

The following SAP configuration restrictions apply to the high SAP scale mode; see [SAPs](#) for additional SAP configuration guidelines:

- If an RVPLS Q1.* SAP is configured, SAPs (Q1.Q2 SAP) with a matching Q1 tag cannot be configured in other VPLS, Epipe, IES, and VPRN services on the same port.
- If a VPLS Q1.* SAP enabled with DHCP snooping is configured, SAPs (Q1.Q2 SAP) with a matching Q1 tag cannot be configured in other VPLS, Epipe, IES, and VPRN services on the same port. They can use other values for the Q1 tag. The reverse is also true.
- The dot1p default SAP cannot be configured in RVPLS services; it is only supported in Epipe, VPLS, IES, and VPRN services.
- If the following criteria are met on the 7210 SAS-Sx 10/100GE, IES and VPRN SAPs for IES or VPRN services are not supported:
 - the SAP has a QinQ encapsulation
 - **access-ingress-qos-mode** is set to **port-mode**
 - **sap-scale-mode** is set to **high**

2.3.4.1 Guidelines for configuring high SAP scale mode

About this task

Perform the following procedure to change the **sap-scale-mode low** to the **sap-scale-mode high** configuration:

Procedure

- Step 1.** Delete all SAPs.
- Step 2.** Configure the **config>system>resource-profile>qos>port-scheduler-mode** command.
- Step 3.** Configure the **sap-scale-mode** command to use the **high** option.

Step 4. Save the configuration and reboot the node.

2.3.4.2 Guidelines for configuring low SAP scale mode

About this task

Perform the following procedure to change the **sap-scale-mode high** to the **sap-scale-mode low** configuration:

Procedure

- Step 1.** Delete all SAPs.
- Step 2.** If the access ingress QoS policy has attachments, reset the policy.
- Step 3.** If the **access-ingress-qos-mode** command is set to **port-mode**, configure the command to use the **sap-mode** option.
- Step 4.** Configure the **sap-scale-mode** command to use the **low** option.
- Step 5.** Save the configuration and reboot the node.

2.3.5 G.8032 Ethernet ring protection switching

Ethernet ring (Eth-ring) protection switching provides ITU-T G.8032 specification compliance to achieve resiliency for Ethernet Layer 2 networks. Similar to G.8031 linear protection (also called Automatic Protection Switching (APS)), G.8032 Eth-ring is implemented on Ethernet OAM and often referred to as Ring Automatic Protection Switching (R-APS).

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

Rings are preferred in data networks where the native connectivity is laid out in a ring or there is a requirement for simple resilient LAN services. Because of the symmetry and the simple topology, rings are viewed a good solution for access and core networks where resilient LANS are required. The Nokia implementation of G.8032 Eth-rings can be used for interconnecting access rings and to provide traffic engineered backbone rings. The 7210 SAS implementation of G.8032 Eth-rings supports dual interconnected rings with sub-rings.

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

2.3.5.1 Overview of G.8032 operation


R-APS messages that carry the G.8032 protocol are sent on a dedicated protocol VLAN called ERP VLAN (or ring control instance). In a revertive case, the G.8032 protocol ensures that one Ring Protection Link (RPL) owner blocks the RPL link. R-APS messages are periodically sent around in both directions to inform other nodes in the ring about the blocked port in the RPL owner node. In non-revertive mode, any link may be the RPL link.

Y.1731 Ethernet OAM CC is the basis of the R-APS messages. Y.1731 CC messages are typically used by nodes in the ring to monitor the health of each link in the ring in both directions. However, CC messages are not mandatory. Other link layer mechanisms could be considered; for example, loss of signal (LoS) 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. When a ring node in the ring learns that another link is blocked, the node unblocks its blocked link, possibly causing an FDB flush in all links of the ring for the affected service VLANs, controlled by the ring control instance. This procedure results in unblocking all links except the one link and the ring normal (or idle) state is reached.


In revertive mode, the RPL link will be the link that is blocked when all links are operable after the revert time. In non-revertive mode, the RPL link is no different from other ring links. Revertive mode provides predictability, particularly when there are multiple ring instances, and the operator can control which links are blocked on the different instances. Each time that there is a topology change that affects reachability, the nodes may flush the FDB and MAC learning takes place for the affected service VLANs, allowing forwarding of packets to continue. The following figure shows this initial operational state.

Figure 7: G.8032 Ring in the Initial State



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

Figure 8: 0 to 1 G.8032 ring in the protecting state



When the failed link recovers, the nodes that blocked the link again send the R-APS messages indicating no failure this time. This causes the RPL owner to block the RPL link and indicate the blocked RPL link to the ring in an R-APS message. When the message is received by the nodes at the recovered link, they unblock that link and restore connectivity (again all nodes in the ring perform an FDB flush and MAC learning takes place). The ring is back in the normal (or idle) state.

Within each path, Y.1731 Maintenance Entity Group (MEG) Endpoints (MEPs) are used to exchange R-APS specific information (specifically to coordinate switchovers). As well, MEPs optionally exchange fast Continuity Check Messages (CCMs) providing an inherent failure detection mechanism as part of the protocol. Failure detection of a ring path by one of the mechanisms activates the protection links. Upon failure, reconvergence times are dependent on the failure detection mechanisms.

In the case of Y.1731, the CCM transmit interval determines the response time. The 7210 SAS device supports 100 ms message timers that allow for quicker restoration times. Alternatively, 802.3ah (Ethernet in the First Mile) or LoS can trigger a protection switch where appropriate. In the case of direct connectivity between the nodes, there is no need to use Ethernet CC messaging for liveness detection.

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

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

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

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

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


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

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

The following figure shows a resilient ring service example, in which a PBB ring (solid line) using VID 500 carries two service VLANs on I-SID 1000 and 1001 for service VIDs (dot1q 100 and QinQ 400.1, respectively). The RPL for the PBB ring is between A and B, where B is the RPL owner. The following figure also shows a QinQ service on the (dotted line) ring that uses Dot1q VID 600 for the ring to connect service VLAN 100.50.

The two rings have RPLs on different nodes, which allows a form of load balancing. The example demonstrates that service and ring encapsulation can be mixed in various combinations. Also, note that neither of the rings is a closed loop. A ring can restore connectivity when any one node or link fails to all remaining nodes within the 50ms transfer time (signaling time after detection).

Figure 9: 0 to 3 resilient ring service example



Example: Ethernet ring configuration

The following is a sample configuration that shows an Ethernet ring configuration.

```
configure eth-ring 1
```

```

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

service
    vpls 10 customer 1 create // Ring APS SAPs
    description "Ring Control VID 100"
    sap 1/1/1:100 eth-ring 1 create // TAG for the Control Path a
    exit
    sap 6/6/2:100 eth-ring 1 create // TAG for the Control Path b
    exit
    no shutdown
    exit
service
    vpls 40 customer 1 b-vpls create //Data Channel on Ring
    description "Ethernet Ring 1 VID 500"
    sap 1/1/1:500 eth-ring 1 create // TAG for the Data Channel Path a
    exit
    sap 6/6/2:500 eth-ring 1 create // TAG for the Data Channel Path b
    exit
    exit
service
    epipe 100 pbb-epipe // CPE traffic
    description "PBB epipe service for CPE"
    pbb-tunnel 40 backbone-dest-mac 00:bb:bb:bb:bb:bb isid 100
    sap 3/1/1:100 create
        description "Default sap description for service id 100"
    exit
    no shutdown
exit


```

2.3.5.2 Ethernet ring sub-rings

Ethernet sub-rings offer a dual redundant way to interconnect rings. The 7210 SAS supports sub-rings connected to major rings, and a sub-ring connected to a VPLS (LDP based) for access ring support in VPLS networks.

The following figure shows 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. Also, the sub-ring (ERP2) relies on the major ring (ERP1) as part of its protection for the traffic from C and D. The nodes C and D are configured as interconnection nodes.

Figure 10: 0 to 4 G.8032 Sub-ring



Sub-rings and major rings run similar state machines for the ring logic; however, there are some differences. When sub-rings protect a link, the flush messages are propagated to the major ring. (A special configuration allows control of this option on the 7210 SAS.) When major rings change topology, the flush is propagated around the major ring and does not continue to any sub-rings. The reason for this is that major rings are completely connected but sub-rings are dependent on another ring or network for full

connectivity. The topology changes need to be propagated to the other ring or network usually. Sub-rings offer the same capabilities as major rings in terms of control and data so that all link resources may be used.

2.3.5.2.1 Virtual and non-virtual channel

Example

The following is a sample sub-ring using virtual-link configuration output on Node C, interconnecting node.

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

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

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

# Data configuration for the Sub-Ring

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

```

sap 3/2/1:1 create
description "Local Data SAP"
stp shutdown
no shutdown
exit

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


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

```

2.3.5.2.2 Ethernet ring sub-ring using non-virtual link

The following figure shows an Ethernet sub-ring that uses a non-virtual link.

Figure 11: 0 to 6 sub-ring homed to VPLS



Example

The following is a sample sub-ring using non-virtual link configuration output on PE1, the interconnecting node.

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

Example

All the sub-ring nodes that are part of a sub-ring with non-virtual link should be configured with the **sub-ring non-virtual-link** option, as shown in the following sample configuration.

```
eth-ring 1
  sub-ring non-virtual-link
  exit
  path a 1/1/1 raps-tag 1.1
    eth-cfm
    mep 5 domain 1 association 4
      ccm-enable
      control-mep
      no shutdown
    exit
  exit
  no shutdown
exit
path b 1/1/2 raps-tag 1.1
  eth-cfm
  mep 6 domain 1 association 3
    ccm-enable
    control-mep
    no shutdown
  exit
  exit
  no shutdown
exit
no shutdown
exit

# Control Channel for Sub-Ring using non-virtual-link on interconnecting node:
```

```

vpls 1 customer 1 create
  description "Ring 1 Control termination"
  stp shutdown
  sap 1/1/3:1.1 eth-ring 1 create //path a control
    stp shutdown
  exit
  no shutdown
exit

# Configuration for the ring data into the VPLS Service

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

  no shutdown
exit

# Control Channel for Sub-Ring using non-virtual-link on sub-Ring nodes:

vpls 1 customer 1 create
  stp
    shutdown
  exit
  sap 1/1/1:1.1 eth-ring 1 create
    stp
      shutdown
  exit
  exit
  sap 1/1/2:1.1 eth-ring 1 create
    stp
      shutdown
  exit
  exit
  no shutdown
exit

```

Example

The following is a sample sub-ring using non-virtual link configuration output homed to a major ring.

```

eth-ring 1
  description "Ethernet Ring 1"
  guard-time 20
  no revert-time
  rpl-node nbr
  sub-ring non-virtual-link
    interconnect ring-id <major ring index>
    propagate-topology-change

```

```

        exit
    exit
    path a 1/1/3 raps-tag 1.1
        description "Ethernet Ring : 1 Path on LAG"
        eth-cfm
        mep 8 domain 1 association 8
            ccm-enable
            control-mep
            no shutdown
        exit
    exit
    no shutdown
    exit
    no shutdown
    exit

```

2.3.5.2.3 Lag support

The 7210 SAS does not support G.8032 Ethernet rings on LAGs.

2.3.5.3 OAM considerations

Ethernet CFM can be enabled on each individual path under an Ethernet ring. Only Down MEPs can be configured on each path, and CCM sessions can be enabled to monitor path liveliness using one of the following intervals: 3.3ms/300Hz, 10 ms, 100 ms, 1 s, 10 s, 60 s, 600 s. The 3.3ms/300Hz interval is only supported on the 7210 SAS-Sx/S 1/10GE in standalone and standalone-VC mode. Different CCM intervals can be supported on path A and path B in an Ethernet ring. CFM is optional if the hardware supports LOS, for example.

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

2.3.5.4 QoS considerations

When Ethernet ring is configured on two ports located on different IOMs, the SAP queues and virtual schedulers will be created with the actual parameters on each IOM.

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

See the *7210 SAS-Mxp, S, Sx, T Services Guide* for more information about Ethernet ring applicability in the services solution.

2.3.5.5 Support service and solution combinations

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

- Only ports in access mode can be configured as Ethernet-ring paths.

- Dot1q and QinQ ports are supported as Ethernet-ring path members.
- A mix of regular and multiple Ethernet-ring SAPs and PWs can be configured in the same services.

2.3.5.6 Configuration guidelines for G.8032


The configuration guidelines for G.8032 are the following:

- For 7210 SAS-T devices in network mode, users must enable the fast-flood feature and allocate resources from the ingress internal TCAM pool using the **config>system>resource-profile>g8032-fast-flood-enable** command.
- For 7210 SAS-T devices in access-uplink mode, users do not need to enable the fast-flood feature or allocate resources from the ingress internal TCAM pool using the **config>system>resource-profile>g8032-fast-flood-enable** command, because the resources are automatically allocated by the software.
- For 7210 SAS-Mxp devices in network mode, the fast-flood feature is enabled by default to improve the service failover time caused by failures in the ring path. If a failure is detected on one of the Ethernet ring paths, along with MAC flush, the system also floods the traffic to the available path. No explicit user configuration is required to facilitate this, and resource allocation from the ingress internal TCAM pool is not needed.
- For 7210 SAS-Sx/S 1/10GE and 7210 SAS-Sx 10/100GE, no explicit user configuration is needed to enable G.8032 fast-flood, and resources do not need to be allocated from the ingress internal TCAM pool.
- Service-level MEPs are not available on SAPs tied to an Ethernet-ring instance on a port.
- G.8032 instances cannot be configured over a LAG.
- G.8032 version 2 is the supported version, by default. Use the **config>eth-ring>compatible-version** command to change the G.8032 version to 1, if required. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide* for more information.
- Configuring the same Maintenance Association Identifier (MAID) on two different MEPs on the same node and service is allowed. This is only applicable to the 7210 SAS-Sx/S 1/10GE in standalone and standalone-VC mode and 7210 SAS-Sx 10/100GE in standalone mode.

2.4 Service creation process overview

The following figure shows the overall process to provision core and subscriber services.

Figure 12: Service creation and implementation flow



2.5 Deploying and provisioning services

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

- core network construction
- service administration
- service provisioning

2.5.1 Phase 1: core network construction

Before the services are provisioned, perform the following tasks:

- Build the IP or IP/MPLS core network.
- Configure routing protocols.
- Configure MPLS LSPs (if MPLS is used).

2.5.2 Phase 2: service administration

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

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

2.5.3 Phase 3: service provisioning

Perform the following tasks to complete service provisioning:

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

2.6 Configuration notes

This section describes service configuration restrictions.

2.6.1 General

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

Core tasks include the following:

- Create customer accounts.
- Create template QoS, filter, scheduler, and accounting policies.
- Create SDPs (not applicable for devices configured in access-uplink mode).

Subscriber services tasks include the following:

- Create Epipe and VPLS services.
- Create a VPRN service (supported only when operating in network mode).
- Bind SDPs (not applicable for 7210 SAS devices configured in access-uplink mode).
- Configure interfaces (where required) and SAPs.
- Create exclusive QoS and filter policies.

To send and receive inband management traffic (for 7210 SAS configured in access-uplink mode), create an IES service.

2.7 Configuring global service entities with CLI

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

2.8 Service model entities

The Nokia service model uses logical entities to construct a service. The service model contains four main entities to configure a service.

2.9 Basic configuration

The most basic service configuration must have the following:

- a customer ID
- a service type
- a service ID
- a SAP identifying a port and encapsulation value
- an associated SDP for distributed services in the network mode

The SDPs are not supported in the access-uplink mode.

Example

The following is a sample Epipe service configuration output showing the SDP and Epipe service entities. SDP ID 1 was created with the far-end node 10.20.1.2. Epipe ID 101 was created for customer ID 1, which uses the SDP ID 1.

```
A:ALA-7210M>config>service#
-----
...
    sdp 1 mpls create
        description "Default sdp description"
        far-end 10.20.1.2
        lsp "lsp_1_to_B"
        signaling tldp
        no vlan-vc-etype
        path-mtu 9194
        no adv-mtu-override
        keep-alive
            shutdown
            hello-time 10
            hold-down-time 10
            max-drop-count 3
            timeout 5
            no message-length
        exit
        no collect-stats
        no accounting-policy
        no shutdown
    exit
```

```

...
    pipe 101 customer 1 vpn 101 create
        description "Default epipe description for service id 101"
        service-mtu 9194
        sap lag-2:101 create
            description "Default sap description for service id 101"
            no tod-suite
            dotlag
            exit
            ingress
                qos 1
                no filter
            exit
        spoke-sdp 101:101 vc-type ether create
            no vlan-vc-tag
            ingress
                no vc-label
            exit
            egress
                no vc-label
            exit
            no control-word
            no
            dotlag
                mep 1 domain 5 association 101 direction down
                    ccm-enable
                    no ccm-ltm-priority
                    low-priority-defect remErrXcon
                    no mac-address
                    no shutdown
                exit
                mep 1 domain 6 association 101 direction down
                    ccm-enable
                    no ccm-ltm-priority
                    low-priority-defect remErrXcon
                    no mac-address
                    no shutdown
                exit
            exit
            no collect-stats
            no accounting-policy
            no precedence
            no shutdown
        exit
        no shutdown
...
-----
A:ALA-7210M>config>service#

```

2.10 Common configuration tasks

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

2.10.1 Configuring customers accounts

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

- description
- contact name
- telephone number

2.10.1.1 Customer information

Use the following syntax to create and input customer information.

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

Example

The following is a sample basic customer account configuration output.

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

2.10.2 Configuring an SDP



Note:

SDPs are supported by all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

The most basic SDP must have the following:

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

2.10.2.1 SDP configuration tasks

About this task

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

Consider the following SDP characteristics:

- SDPs can be created as MPLS.

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

If signaling is disabled for an SDP, services using that SDP must configure ingress and egress VC labels manually.

Procedure

To configure a basic SDP, perform the following steps:

- Step 1.** Specify an originating node.
- Step 2.** Create an SDP ID.
- Step 3.** Specify an encapsulation type.
- Step 4.** Specify a far-end node.

2.10.2.2 Configuring an SDP

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



Note:

When you specify the far-end IP address, you are creating the tunnel; in essence, you are creating the path from point A to point B. When you configure a distributed service, you must identify an SDP ID. Use the **show service sdp** command to display the qualifying SDPs.

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

LSPs are configured in the **config>router>mpls** context. See the *7210 SAS-Mxp, R6, R12, S, Sx, T MPLS Guide* for configuration and command information.

Use the following syntax to create an MPLS SDP.

```
config>service>sdp sdp-id [mpls] create
  adv-mtu-override
  description description-string
  far-end ip-address
  keep-alive
  hello-time seconds
  hold-down-time seconds
  max-drop-count count
```

```

message-length octets
timeout timeout
no shutdown
    lsp lsp-name [lsp-name](only for MPLS SDPs)
path-mtu octets
signaling {off | tldp}
no shutdown

```

Example

The following is a sample LSP-signaled MPLS SDP configuration output.

```

A:ALA-12>config>service# info
-----
...
    sdp 8 mpls create
        description "MPLS-10.10.10.104"
        far-end 10.10.10.104
        lsp "to-104"
        keep-alive
        mixed-lsp-mode
            revert-time 1
        shutdown
    exit
    no shutdown
exit
...
-----
A:ALA-12>config>service#

```

2.10.2.3 Configuring a mixed-LSP SDP

The following is the command usage to configure an SDP with mixed-LSP mode of operation:

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

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

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

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

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

An *infinite* value for the timer dictates that the SDP must never revert to another higher priority LSP type unless the currently active LSP type is down:

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

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

2.11 Ethernet connectivity fault management

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

ETH-CFM configuration is split into multiple CLI contexts. The ETH-CFM configuration, which defines the different management constructs and administrative elements, is performed in the **eth-cfm** context. The individual management points are configured within the specific service contexts in which they are applied (port, SAP, and so on).

See the *7210 SAS-Mxp, S, Sx, T Services Guide* for detailed information about the basic service-applicable material to build the service-specific management points, MEPs, and MIPs. The different service types support a subset of the features from the complete ETH-CFM suite.

Ethernet continuity check (ETH-CC) used for continuity is available to all MEPs configured within a service. 7210 SAS devices support Down MEPs and UP MEPs, though the support is not available on all platforms. See the *7210 SAS-Mxp, R6, R12, S, Sx, T OAM and Diagnostics Guide* for more information.

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

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

See the *7210 SAS-Mxp, R6, R12, S, Sx, T OAM and Diagnostics Guide* for a description of the individual features and functions that are supported and configuration guidelines applicable to CFM entities on the 7210 SAS.

The following table lists ETH-CFM acronym expansions.

Table 9: ETH-CFM acronym expansions

Acronym	Expansions
1DM	One-way Delay Measurement (Y.1731)
AIS	Alarm Indication Signal
BNM	Bandwidth Notification Message (Y.1731 sub OpCode of GMN)
CCM	Continuity Check Message
CFM	Connectivity Fault Management
DMM	Delay Measurement Message (Y.1731)
DMR	Delay Measurement Reply (Y.1731)
GMN	Generic Message Notification
LBM	Loopback Message

Acronym	Expansions
LBR	Loopback Reply
LTM	Linktrace Message
LTR	Linktrace Reply
ME	Maintenance Entity
MA	Maintenance Association
MA-ID	Maintenance Association Identifier
MD	Maintenance Domain
MEP	Maintenance Association Endpoint
MEP-ID	Maintenance Association Endpoint Identifier
MHF	MIP Half Function
MIP	Maintenance Domain Intermediate Point
OpCode	Operational Code
RDI	Remote Defect Indication
TST	Ethernet Test (Y.1731)
SLM	Synthetic Loss Message (Y.1731)
SLR	Synthetic Loss Reply (Y.1731)

2.11.1 MA, MEP, MIP, and MD levels

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

Figure 13: Ethernet OAM model for Ethernet access – business





Figure 14: Ethernet OAM model for Ethernet access – wholesale



The following functions are supported:

- CFM can be enabled or disabled on a SAP or SDP bindings basis.
- The eight ETH-CFM levels are suggested to be broken up numerically between customer 7-5, service provider 4-3 and operator 2-1. Level 0 typically is meant to monitor direct connections without any MIPs and should be reserved for port-based G8032 MEPs.
- Down MEP and UP MEP with an MEP-ID on a SAP/SDP binding for each MD level can be configured, modified, or deleted. Each MEP is uniquely identified by the MA-ID, MEP-ID tuple:
 - MEP creation on a SAP is allowed only for Ethernet ports (with null, q-tags, QinQ encapsulations).
 - MEP support in different services and the endpoints configured in the services (SAPs, SDPs, IP interfaces, and so on) varies across services and 7210 SAS platforms.
- MIP creation on a SAP for each MD level can be enabled and disabled. MIP creation is automatic or manual when it is enabled. When MIP creation is disabled for an MD level, the existing MIP is removed. 7210 SAS platforms have the notion of ingress and egress MIPs. Ingress MIP responds to OAM

messages that are received. Egress MIP responds to OAM messages that are sent. Ingress and egress MIP support for SAP, SDP bindings, and services varies and is listed in [Table 10: Defect conditions and priority settings](#). See [MEP and MIP support](#) for more information about MEP and MIP support.

2.11.1.1 Common actionable failures

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

Table 10: Defect conditions and priority settings

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

Defect	Low priority defect	Description	Causes	Priority
			service, incorrect association identification.	

2.11.1.2 MEP and MIP support

See the *7210 SAS-Mxp, R6, R12, S, Sx, T OAM and Diagnostics Guide* for more information about ETH-CFM support for different services and endpoints.

2.11.2 Configuring ETH-CFM parameters



Note:

See the *7210 SAS-Mxp, R6, R12, S, Sx, T OAM and Diagnostics Guide* for more information about ETH-CFM configuration guidelines for 7210 SAS platforms.

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

This section provides a sample of the global ETH-CFM configuration, which defines the domains, associations, linkage of the service ID or function, and the globally applicable CCM parameters, including the interval and building of the remote MEPs database.

Example

The following is a sample configuration output.

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

Example

Defining the MEP and configuring service-specific ETH-CFM parameters is performed within the service on the specific SAP or SDP binding. The following is sample output using the service VPLS 100 on the SAP.

```
##A:ALU-7_A>config>service# info
-----
    vpls 100 customer 1 create
      description "VPLS service 100 - Used for MEP configuration example"
        sap 2/2/1:20 create
          description "2/2/1:20"
            eth-cfm
              mep 1 domain 1 association 1 direction down
```

```

                                no shutdown
                                exit
                                exit
                                exit
                                no shutdown
                                exit
                                customer 1 create
                                    description "Default customer"
                                exit
                                exit
-----
*A:ALU-7_A>config>service#

```

The preceding samples were based on IEEE 802.1ag. They are not capable of running Y.1731 functions. To build a Y.1731 context, the domain format must be none.

Example

The following are sample global ETH-CFM configuration outputs and the advanced Y.1731 functions that can be configured. The configuration will reject the configuration of Y.1731 functions within an IEEE 802.1ag context.

```

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

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

```

**Note:**

- To transmit and receive AIS PDUs, Y.1731 MEPs must have **ais-enable** configured.
- To transmit and receive ETH-Test PDUs, Y.1731 MEPs must have **eth-test-enable** configured.

2.11.3 Applying ETH-CFM parameters

Use the following syntax to apply ETH-CFM parameters to the following entities.

```
config>service>epipe>sap
eth-cfm
mep mep-id domain md-index association ma-index [direction
{up | down}]
    ais-enable
        client-meg-level [level [level ...]]
        interval {1 | 60}
        priority priority-value
    ccm-enable
    ccm-ltm-priority priority
    eth-test-enable
        test-pattern {all-zeros | all-ones} [crc-enable]
    low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon | xcon | noXcon}
    [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
    eth-test-enable
        test-pattern {all-zeros | all-ones} [crc-enable]
    low-priority-defect {allDef|macRemErrXcon|remErrXcon|
errXcon|xcon|noXcon}
    [no] shutdown
```

```
config>service>vppls>sap
eth-cfm
mip
mep mep-id domain md-index association ma-index [direction {up | down}]
no mep mep-id domain md-index association ma-index
    ccm-enable
    ccm-ltm-priority priority
    eth-test-enable
        test-pattern {all-zeros | all-ones} [crc-enable]
    low-priority-defect {allDef|macRemErrXcon|remErrXcon|errXcon|xcon|noXcon}
    mac-address mac-address
    [no] shutdown
```

```
config>service>vppls>mesh-sdp sdp-id[:vc-id] [vc-type {ether|vlan}]
eth-cfm
mep mep-id domain md-index association ma-index [direction
{up | down}]
    ccm-enable
    ccm-ltm-priority priority
    eth-test-enable
        test-pattern {all-zeros | all-ones} [crc-enable]
    low-priority-defect {allDef|macRemErrXcon|remErrXcon|
```

```

errXcon|xcon|noXcon}
mac-address mac-address
no] shutdown

config>service>vpls
  spoke-sdp sdp-id:vc-id [vc-type {ether | vlan}] [split-horizon-group group-name] [no-
endpoint]
  spoke-sdp sdp-id:vc-id [vc-type {ether | vlan}] [split-horizon-group group-name]
endpoint endpoint
  eth-cfm
  map mep-id domain md-index association ma-index [direction
  {up | down}]
  ccm-enable
  ccm-ltm-priority priority
  eth-test-enable
    test-pattern {all-zeros | all-ones} [crc-enable]
  low-priority-defect {allDef | macRemErrXcon|remErrXcon|errXcon|xcon|noXcon}
  mac-address mac-address
  no] shutdown

oam
  eth-cfm linktrace mac-address mep mep-id domain md-index association ma-index [ttl ttl-
value]
  eth-cfm loopback mac-address mep mep-id domain md-index association ma-index [send-
count send-count] [size data-size] [priority priority]
  eth-cfm eth-test mac-address mep mep-id domain md-index association ma-index
[priority priority] [data-length data-length]
  eth-cfm one-way-delay-test mac-address mep mep-id domain md-index association ma-index
[priority priority]
  eth-cfm two-way-delay-test mac-address mep mep-id domain md-index association ma-index
[priority priority]
  eth-cfm two-way-slm-test mac-address mep mep-id domain md-index association ma-index
[priority priority]

```

2.12 Layer 2 control processing

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

- an option to tunnel, discard, or peer for EFM OAM, LLDP, dot1x, and LACP
- BPDU translation and Layer 2 protocol tunneling support for xSTP and Cisco control protocols. This is supported only in a VPLS service.

See [L2PT and BPDU translation](#) for more information.



Note:

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

By default, LACP, LLDP, EFM OAM, and dot1x Layer 2 control protocol untagged packets are discarded if the protocol is not enabled on the port where these frames are received. The user has an option to enable peering by enabling the protocol on the port and configuring the appropriate parameters for the protocol. The user also has an option to tunnel these packets using an Epipe or VPLS service.

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

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

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

In addition to the protocols listed previously, protocols that are not supported on the 7210 SAS (for example, GARP, GVRP, ELMI, and others) are transparently forwarded in case of a VPLS service. These protocols are transparently forwarded if a null SAP, dot1q default SAP, dot1q explicit null SAP or 0.* SAP is configured on the port and the received packet is untagged. If the received packet is tagged and matches the tag of any of the SAPs configured on the port, it is forwarded in the context of the SAP and the service. Otherwise, if the received packet is untagged and none of the null or dot1q default or dot1q explicit null or 0.* SAP is configured, it is discarded.

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

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

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

The following table describes L2CP support for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Mxp access-uplink and network media platforms.

Table 11: L2CP support for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp access-uplink and network mode platforms

Packet type	7210 SAS-T	7210 SAS-Mxp	7210 SAS-Sx/S 1/10GE	7210 SAS-Sx 10/100GE
LACP	Option to tunnel or discard or peer	Option to tunnel or discard or peer	Option to tunnel or discard or peer	Option to tunnel or discard or peer

Packet type	7210 SAS-T	7210 SAS-Mxp	7210 SAS-Sx/S 1/10GE	7210 SAS-Sx 10/100GE
Dot1x	Option to tunnel or discard or peer	Option to tunnel or discard or peer	Option to tunnel or discard or peer	Option to tunnel or discard or peer
LLDP	Option to tunnel or discard or peer ⁹	Option to tunnel or discard or peer ⁹	Option to tunnel or discard or peer ⁹	Option to tunnel or discard or peer ⁹
EFM	Option to tunnel or discard or peer	Option to tunnel or discard or peer	Option to tunnel or discard or peer	Option to tunnel or discard or peer
L2PT	Supported ¹⁰	Supported ¹⁰	Supported ¹⁰	Supported ¹⁰
BPDU Tunneling	Supported	Supported	Supported	Supported
xSTP	Option to peer or tunnel	Option to peer or tunnel	Option to peer or tunnel	Option to peer or tunnel
ESMC protocol	Option to tunnel or discard or peer	Option to tunnel or discard or peer	Option to tunnel or discard or peer ¹¹	Option to tunnel or discard or peer

2.13 Dynamic VLAN assignment using dot1x RADIUS authentication with EHS



Note:

Dynamic VLAN assignment using dot1x RADIUS authentication with EHS is only supported on the 7210 SAS-Sx/S 1/10GE (standalone).

On the 7210 SAS, users can assign a VLAN using the RADIUS tunnel attribute. Only the VLAN is returned by RADIUS, while other policies (such as QoS, ACLs, accounting) are not. The locally configured policies can be applied when the VLAN ID is used to configure the SAP after a successful authentication of the host using dot1x (including MAC authentication).

⁹ See the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide* for more information about options available for LLDP tunneling.

¹⁰ L2TP support on 7210 SAS platforms varies among the platforms. Not all platforms support tunneling of all CISCO protocols. See [L2PT and BPDU translation](#) for more information.

¹¹ ESMC tunneling is not supported in the standalone-VC mode of operation.

2.13.1 Assignment of a VLAN after a host is authenticated using dot1x

The 7210 SAS provides support for RADIUS-based dot1x or MAC authentication to authenticate network access in enterprise networks. If RADIUS authentication is successful on the port on which the dot1x frames are received or the packet that triggered MAC authentication, access is unblocked and the user is granted access to the network. In addition, the RADIUS message contains the VLAN ID that is assigned to the traffic received on the port. The SAP for the user along with the QoS policy, ACLs, or accounting policy can be configured using EHS only after the VLAN tag information is received upon a successful authentication.

During an unsuccessful authentication, the port is blocked, preventing network access.

It is expected that the service is preconfigured with uplinks and service parameters (if any), including preconfiguration of the access port (for example, port mode and encap, MTU, and so on) to which the device is connected. During a successful authentication, only the access SAP with the appropriate VLAN along with any policies associated with SAP can be configured using EHS scripts.

To achieve this functionality, the event handling system (EHS) script must be configured so that the script is invoked with the appropriate parameters received from the RADIUS server to allow the script to configure the SAP and the policies associated with the SAP for the user.


The following figure shows how devices are mapped to a VLAN (service and SAP in the SR OS context) in an enterprise network. The devices can send tagged or untagged traffic, and use port, MAC, or VLAN authentication.



Note:

When using VLAN authentication, dynamic VLAN assignment procedures are triggered. The device must know the VLAN to use for dot1x authentication which can be assigned using LLDP-MED. It is expected that the device continues to use the same VLAN for service traffic after authentication is completed successfully, though there are no checks enforced by the node to ensure the same VLAN is used.

Figure 15: Dynamic VLAN assignment using dot1x in enterprise



In the preceding figure, the user has connected a camera that sends untagged traffic using MAC authentication, a Wi-Fi AP that sends tagged traffic using port authentication, and a laptop that sends

untagged traffic using MAC authentication. Based on the VLAN returned by the RADIUS server, the traffic from each of these devices is mapped to a different VPLS with the appropriate SAP configured to accept tagged and untagged traffic from the devices connected to the 7210 SAS.

**Note:**

VPLS with SAPs to accept tagged and untagged traffic is used in SR OS to implement a VLAN bridge service.

2.13.2 Assigning a VLAN to a device based on dot1x authentication

About this task

Perform this procedure to achieve the scenario described in [Assignment of a VLAN after a host is authenticated using dot1x](#).

Procedure

- Step 1.** Configure dot1x, MAC, or VLAN authentication on the required ports.
- Step 2.** If an EAPOL request is received on the port, use the EAPOL message to authenticate the connected device; or if the device connected to the port is not capable of dot1x, then use the source MAC address from the first packet received on the port to authenticate the device by sending the dot1x request to the RADIUS server for authentication.
- Step 3.** If the RADIUS server has returned an authentication "SUCCESS", it is expected to also send the VLAN to be configured for the port. The software triggers the EHS scripts associated with the dot1x "tmnxPortDot1xAuthSuccess" event and it is expected that the EHS script performs the following:
 - a.** Configure the access SAP using the VLAN and authenticated port, and associate it with the appropriate service.

**Note:**

The VPLS service is preconfigured with only the uplinks (access SAPs are added dynamically). In addition, the access port is preconfigured with appropriate port mode and encap type.

- b.** Configure the other SAP policies (QoS, ACL, accounting, and so on) and other parameters.

**Note:**

See the *7210 SAS-Mxp, R6, R12, S, Sx, T System Management Guide* for more information about EHS and script configuration examples.


2.13.3 Guest and restricted VLAN service support

The 7210 SAS allows a host or device configured with a corporate VLAN to connect to the company network (for example, devices used by employees), and a guest or restricted VLAN for devices used by visitors to allow internet access. To access the corporate VLAN, the employee device must successfully authenticate using 802.1x through a RADIUS server that acts as the AAA server.

The following figure shows a simplified architecture in which a host connected to the switch needs to connect to the network. There are three possibilities:

- VPLS 1 – corporate VLAN access if 802.1x authentication is successful
- VPLS 2 – restricted VLAN access if 802.1x authentication fails
- VPLS 3 – guest VLAN if the device does not send EAPOL packets or if there is no response from the RADIUS server

Figure 16: Corporate or guest VPLS assignment based on 802.1x authentication outcome



When the RADIUS server returns an authentication failure and there is no response from the RADIUS server, or if dot1x authentication is enabled on a port but no EAPOL packets are received from the connected device, the connected user or device must assign a guest or restricted VLAN.

Assign a guest VLAN when there is no response from the RADIUS server or when dot1x authentication is enabled on a port but no EAPOL packets are received from the connected device.

Assign a restricted VLAN when the RADIUS server returns authentication failure.

Use the following commands to configure a guest or restricted VLAN:

- **config>port>ethernet>dot1x>guest-service** *service-id* [*vlan-id* *vlan-id*]
- **config>port>ethernet>dot1x>restricted-service** *service-id* [*vlan-id* *vlan-id*]

Use the EHS script to configure a *service-id* or *svc-name* to specify the service under which the SAP must be configured. The EHS script must use the NULL (":0") tag to configure a NULL SAP if the optional *vlan-id* is not configured. If the *vlan-id* is configured, it is used as the SAP tag.



Note: The configured *vlan-id* cannot be modified if the **config>port>ethernet>dot1x>port-control** command is set to **auto**. To modify the service and VLAN ID information, users must set the **config>port>ethernet>dot1x>port-control** command to either **force-auth** or **force-unauth**.

Similar to the successful authentication scenario in which the VLAN provided by the RADIUS server is passed to the EHS script, information about the configured guest VLAN and restricted VLAN is passed to the EHS script for configuration of the SAP in the appropriate service. Specifically, the software passes the *service-id* or *svc-name* and the configured *vlan-id* to EHS script. The configured *service-id* can reference an Epipe or VPLS service (or a Layer 3 service).

The following log events have been added to allow the user to create and delete the guest and restricted VLAN service:

- **tmnxPortDot1xAuthLostGRvC**

This log event invokes the EHS script that must be used to configure the SAP in the service used for guest and restricted VLAN services.

- **tmnxPortDot1xAuthLostGRvD**

This log event invokes the EHS script that must be used to delete the SAP in the service used for guest and restricted VLAN services.

2.13.4 Configuration guidelines

The following configuration guidelines apply to dynamic VLAN assignment using dot1x RADIUS authentication with EHS:

- To assign the SAP and a service based on the VLAN ID obtained from the RADIUS server, an EHS script must be provided. The EHS script is expected to use the parameters passed to the script and the appropriate CLI commands supported by the node to create the SAP in the required service. For an example see the *7210 SAS-Mxp, R6, R12, S, Sx, T System Management Guide*.
- The SAPs created through EHS scripts are treated as dynamic SAPs. Dynamic SAPs are not saved in the configuration file and do not survive reboots. Upon a reboot, the device or user is authenticated afresh, and as a result, the EHS script is executed again to create the SAP and associate it with a service.
- Dynamic SAPs are identified as "Dynamic SAPs" in the **show** command output.
- It is mandatory to use the VLAN ID returned by the RADIUS server as the service-delimiting VLAN tag of the SAP created through the EHS script. The software uses the VLAN ID to identify dynamic SAPs created through EHS scripts.
- An EHS script must be provided for the dot1x authentication lost event (for example, `tmnxPortDot1xAuthLost`). The script is expected to remove the SAP configuration.
- For the RADIUS server, the "Tunnel-Type" value is set to 13, the "Tunnel-Medium-Type" value is 6, and the "Tunnel-Private-Group-Id" value refers to the VLAN ID.
- The value configured with the **config>port>ethernet>dot1x>server>timeout** command must be greater than the timeout value configured with the **config>system>security>radius>timeout** and **config>system>security>radius>retry** commands. This allows for sufficient time before the software can timeout the authentication attempt with the RADIUS server before attempting dot1x authentication again.

2.14 Service management tasks

This section describes the service management tasks.

2.14.1 Modifying customer accounts

To access a specific customer account, you must specify the customer ID.

Use the following syntax to display a list of customer IDs.

Enter the parameter (description, contact, phone) and then enter the new information.

```
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"
```

2.14.2 Deleting customers

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

```
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
```

2.14.3 Modifying SDPs



Note:

SDPs are supported by all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

To access a specific SDP, you must specify the SDP ID. To display a list of SDPs, use the **show service sdp** command. Enter the parameter, such as description, far-end, and lsp, and then enter the new information.



Note:

When an SDP is created, the SDP encapsulation type cannot be modified.

```
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
```

2.14.4 Deleting SDPs

The **no** form of the **sdp** command removes an SDP ID and all associated information. Before an SDP can be deleted, the SDP must be shutdown and removed (unbound) from all customer services where it is applied.

```
config>service# no sdp 79
```

Example:

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

2.15 Global services command reference

2.15.1 Command hierarchies

- [Global service configuration commands](#)
 - [Customer commands](#)
 - [Pseudowire \(PW\) commands](#)
 - [SDP commands](#)
 - [SAP commands for 7210 SAS platforms operating in network mode](#)
 - [SAP commands for 7210 SAS platforms operating in access-uplink mode](#)
 - [ETH-CFM configuration commands](#)
- [Show commands](#)
- [Tools commands](#)

2.15.1.1 Global service configuration commands

2.15.1.1.1 Customer commands

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

2.15.1.1.2 Pseudowire (PW) commands



Note:

PW commands are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

```

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

config
- service
- [no] pw-template policy-id [use-provisioned-sdp] [create]
- accounting-policy acct-policy-id
- no accounting-policy
- [no] collect-stats
- [no] control-word
- [no] disable-learning
- [no] disable-aging
- [no] discard-unknown-source
- [no] force-vlan-vc-forwarding
- hash-label [signal-capability]
- no hash-label
- igmp-snooping
- [no] disable-router-alert-check
- import policy-name
- no import
- last-member-query-interval 1/10 seconds
- no last-member-query-interval
- max-num-groups max-num-groups
- no max-num-groups
- query-interval seconds
- no query-interval
- query-response-interval seconds
- no query-response-interval
- robust-count robust-count
- no robust-count
- [no] send-queries
- version version
- no version
- limit-mac-move {blockable | non-blockable}
- no limit-mac-move
- [no] mac-pinning
- max-nbr-mac-addr table-size

```

```

- no max-nbr-mac-addr
- split-horizon-group group-name
- no split-horizon-group
  - description description-string
  - no description
- vc-type {ether | vlan}
- vlan-vc-tag 0..4094
- no vlan-vc-tag

```

2.15.1.1.3 SDP commands



Note:

SDP commands are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

```

config
- service
- sdp sdp-id [mpls] [create]
- no sdp sdp-id
  - accounting-policy acct-policy-id
  - no accounting-policy
  - collect-stats acct-policy-id
  - no collect-stats
  - [no] adv-mtu-override
  - [no] bgp-tunnel
  - [no] collect-stats
  - description description-string
  - no description
  - far-end ip-address | ipv6-address
  - no far-end [ip-address | ipv6-address]
  - keep-alive
    - hello-time seconds
    - no hello-time
    - hold-down-time seconds
    - no hold-down-time
    - max-drop-count count
    - no max-drop-count
    - message-length octets
    - no message-length
    - [no] shutdown
    - timeout timeout
    - no timeout
  - [no] ldp
  - metric metric
  - no metric
  - no mixed-lsp-mode
  - mixed-lsp-mode
    - no revert-time
    - revert-time {revert-time | infinite}
  - [no] lsp lsp-name
  - path-mtu octets
  - no path-mtu
  - [no] shutdown
  - signaling [off | tldp | bgp]
  - [no] sr-isis
  - [no] sr-ospf

```

2.15.1.1.4 SAP commands for 7210 SAS platforms operating in network mode

```

config
- service
- epipe
  - sap sap-id [create]
  - no sap sap-id
- ies
  - sap sap-id [create]
  - no sap sap-id
- vpls
  - sap sap-id [split-horizon-group group-name] [eth-ring ring-index] [create]
  - no sap sap-id
- vprn
  - interface ip-int-name [create]
  - no interface ip-int-name
    - sap sap-id [create]
    - no sap sap-id

```

2.15.1.1.5 SAP commands for 7210 SAS platforms operating in access-uplink mode

```

config
- service
- epipe service-id [customer customer-id] [create] [svc-sap-type {null-star | dot1q-preserve|any|dot1q-range}] [customer-vid vlan-id]
  - no epipe service-id
    - sap sap-id [create]
    - no sap sap-id
- ies service-id [customer customer-id] [create]
  - no ies service-id
    - sap sap-id [create]
    - no sap sap-id
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type {nullstar | any | dot1q-preserve}] [customer-vid vlan-id]
  - no vpls service-id
    - sap sap-id [create]
    - no sap sap-id

```

2.15.1.1.6 ETH-CFM configuration commands



Note:

For command descriptions, refer to the *7210 SAS-Mxp, R6, R12, S, Sx, T OAM and Diagnostics Guide*.

```

config
- eth-cfm
- domain md-index [format md-name-format] [name md-name] level level
- domain md-index
- no domain md-index
  - association ma-index [format ma-name-format] name ma-name
  - association ma-index
  - no association ma-index
    - [no] bridge-identifier bridge-id
    - id-permission {chassis}

```



```

- no id-permission
- mhf-creation {none | explicit | default | static}
- no mhf-creation
- mip-ltr-priority priority
- vlan vlan-id
- no vlan
- ccm-interval {300Hz | 10ms | 100ms | 1 | 10 | 60 | 600}
- no ccm-interval
- [no] remote-mepid mep-id
- slm
- inactivity-timer timer
- no inactivity-timer
- system
- sender-id local local-name
- sender-id system
- no sender-id

```

2.15.1.2 Show commands

```

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

```

2.15.1.3 Tools commands

```

tools
- perform
- service

```

```

- eval-pw-template policy-id [allow-service-impact]
- id service-id
  - endpoint endpoint-name
    - force-switchover sdp-id:vc-id
    - no force-switchover
    - force-switchover spoke-sdp-fec [1..4294967295]
  - eval-pw-template policy-id [allow-service-impact]
- eval-expired-fec
  - eval-expired-fec spoke-sdp-fec-id
  - eval-expired-fec all
- spoke-sdp-fec-release global-id[:prefix[:ac-id]]

```

2.15.2 Command descriptions

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

2.15.2.1 Global service configuration commands

- [Generic commands](#)
- [Customer management commands](#)
- [Pseudowire Commands](#)
- [SDP commands](#)
- [SDP keepalive commands](#)

2.15.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>service>customer

config>service>sdp (not supported in access-uplink mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

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 configuration.

Default

No description associated with the configuration context.

Parameters

string

Specifies the description character string. Allowed values are any string up to 80 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax

[no] shutdown

Context

config>dot1ag>mep

config>service>sdp (not supported in access-uplink mode)

config>service>sdp>keep-alive (not supported in access-uplink mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

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

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

Services are created in the administratively down (**shutdown**) state. Default administrative states for services and service entities is described as follows in Special Cases.

The **no** form of this command places the entity into an administratively enabled state and then tries to enter the operationally up state.

Special Cases

Service Admin State

Bindings to an SDP within the service are put into the out-of-service state when the service is shutdown. While the service is shutdown, all customer packets are dropped and counted as discards for billing and debugging purposes.

SDP (global)

When an SDP is shutdown at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would be transmitted using this SDP binding are 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.

2.15.2.1.2 Customer management commands**customer****Syntax**

customer *customer-id* [**create**]

no customer *customer-id*

Context

config>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

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

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

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

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

Default

customer 1

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 *contact-information*

Context

config>service>customer

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies contact information for a customer, such as a technician name or account contract name.

Default

no contact

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

Parameters

contact-information

Specifies the customer contact information entered as an ASCII character string up to 80 characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Any printable, seven bit ASCII characters may be used within the string.

phone

Syntax

[no] **phone** *string*

Context

config>service>customer

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies telephone number information for a customer ID.

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

Parameters

string

Specifies the customer phone number entered as an ASCII string, up to 80 characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Any printable, seven bit ASCII characters may be used within the string.

2.15.2.1.3 Pseudowire Commands

pw-routing

Syntax

pw-routing

Context

config>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure dynamic multi-segment pseudowire (MS-PW) routing. Pseudowire routing must be configured on each node that will be a T-PE or an S-PE.

Default

disabled

boot-timer

Syntax

boot-timer secs

no boot-timer

Context

config>service>pw-routing

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

Default

10

Parameters

timer-value

The value of the boot timer in seconds.

Values 0 to 600

local-prefix

Syntax

local-prefix *local-prefix* [**create**]

no local-prefix *local-prefix*

Context

config>service>pw-routing

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

Default

no local-prefix

Parameters

local-prefix

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

Values *global-id:ip-addr | raw-prefix*
 ip-addr: a.b.c.d
 raw-prefix: 1 to 4294967295
 global-id: 1 to 4294967295

advertise-bgp

Syntax

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

Context

config>service>pw-routing

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables a specific prefix to be advertised in MP-BGP for dynamic MS-PW routing. The **no** form of this command explicitly withdraws a route if it has been previously advertised.

Default

no advertise-bgp.

Parameters

rd

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

Values (6 bytes, other 2 Bytes are automatically generated) asn:number1
 (RD Type 0): 2bytes ASN and 4 bytes locally administered number

ip-address:number2 (RD Type 1): 4bytes IPv4 and 2 bytes locally administered number;

community *community*

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

Values *community*: {2-byte-as-number:comm-val}
 2-byte-asnumber: 1 to 65535
 comm.-val: 0 to 65535

path

Syntax

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

Context

config>service>pw-routing

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

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

Default

no path

Parameters

name

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

hop

Syntax

hop *hop-index ip-address*

no hop *hop-index*

Context

config>service>pw-routing>hop

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

The **no** form of this command deletes hop list entries for the path. All the MS-PWs currently using this path are unaffected. Additionally, all services actively using these MS-PWs are unaffected. The path must be shutdown first to delete the hop from the hop list. The '**no hop hop-index**' command does result in any action, except for a warning message on the console indicating that the path is administratively up.

Default

no hop

Parameters

hop-index

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

Values 1 to 16

ip-address

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

retry-count

Syntax

retry-count [10..10000]

no retry-count

Context

config>service>pw-routing

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This optional command specifies the number of attempts software should make to reestablish the spoke-SDP after it has failed. After each successful attempt, the counter is reset to zero.

When the specified number is reached, no more attempts are made and the spoke-SDP is put into the shutdown state. Use the no shutdown command to bring up the path after the retry limit is exceeded.

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

Default

30

Parameters

retry-count

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

Values 10 to 10000

retry-timer

Syntax

retry-timer secs

no retry-timer

Context

config>service>pw-routing

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies a retry-timer for the spoke-SDP. This is a configurable exponential back-off timer that determines the interval between retries to reestablish a spoke-SDP if it fails and a label withdraw message is received with the status code "All unreachable".

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

Default

30

Parameters

retry-count

Specifies the initial retry-timer value in seconds.

10 – 480

spe-address

Syntax

spe-address *global-id:prefix*

no spe-address

Context

config>service>pw-routing

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

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

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

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

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

Default

no spe-address

Parameters

global-id

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

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

global-id 1 to 4294967295

prefix 1 to 4294967295

ipaddress a.b.c.d

static-route

Syntax

[no] static-route *route-name*

Context

config>service>pw-routing

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

A default static route is entered as follows:

static-route 0:0:*next_hop_ip_address*s

or

static-route 0:0.0.0.0:*next_hop_ip_address*

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

Default

no static-route

Parameters

route-name

Specifies the static pseudowire route.

Values

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

<global-id>: 0 to 4294967295

prefix a.b.c.d | 0 to 4294967295

ip_addr a.b.c.d

pw-template

Syntax

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

Context

config>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures an SDP template.

Parameters

use-provisioned-sdp

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

create

Required keyword when first creating the configuration context. When the context is created, it is possible to navigate into the context without the create keyword.

control-word

Syntax

[no] control-word

Context

config>service>pw-template

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the use of the control word on pseudowire packets in VPLS and enables the use of the control word individually on each mesh SDP or spoke-SDP. By default, the control word is disabled. When the control word is enabled, all VPLS packets, including the BPDU frames, are encapsulated with the control word when sent over the pseudowire. The T-LDP control plane behavior is the same as in the implementation of control word for VLL services. The configuration for the two directions of the Ethernet pseudowire should match.

The **no** form of this command reverts the mesh SDP or spoke-SDP to the default behavior of not using the control word.

Default

no control-word

2.15.2.1.4 SDP commands



Note:

SDP commands are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

sdp

Syntax

sdp *sdp-id* [**mpls**] [**create**]

no sdp *sdp-id*

Context

config>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

The 7210 SAS supported only MPLS encapsulation as the method to reach the far-end router. It does not support GRE or other encapsulation methods. A 7210 SAS supports both signaled and non-signaled Label Switched Paths (LSPs) through the network. Non-signaled paths are defined at each hop through the network. Signaled paths are communicated by protocol from end to end using Resource ReserVation Protocol (RSVP). Paths may be manually defined or a constraint-based routing protocol (such as OSPF-TE or CSPF) can be used to determine the best path with specific constraints. An LDP LSP can also be used for an SDP when the encapsulation is MPLS. The use of an LDP LSP type or an RSVP/Static LSP type are mutually exclusive except when the mixed-lsp option is enabled on the SDP.

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

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

If *sdp-id* exists, the current CLI context is changed to that SDP for editing and modification. For editing an existing SDP, the **mpls** keyword is specified. If a keyword is specified for an existing *sdp-id*, an error is generated and the context of the CLI 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 and generates 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.

Parameters

sdp-id

Specifies the SDP identifier.

Values 1 to 17407

accounting-policy

Syntax

accounting-policy acct-policy-id

no accounting-policy

Context

config>service>sdp

config>service>pw-template

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

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

Default

Default accounting policy.

Parameters

acct-policy-id

Specifies the accounting policy-id, which must be entered as configured in the **config>log>accounting-policy** context.

Values 1 to 99

collect-stats

Syntax

[no] collect-stats

Context

config>service>sdp

config>service>pw-template

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables accounting and statistical data collection for an SDP. When applying accounting policies, the data, by default, is collected in the appropriate records and written to the designated billing file.

The **no** form of this command allows the IOM cards to continue to accumulate statistics. However, the CPU does not obtain the results and write them to the billing file. When 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

no collect-stats

discard-unknown-source

Syntax

[no] discard-unknown-source

Context

config>service>pw-template

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables packets received with an unknown source MAC address to be dropped only if the maximum number of MAC addresses have been reached. When disabled, the packets are forwarded based on the destination MAC addresses.

The **no** form of this command causes packets with an unknown source MAC addresses to be forwarded by destination MAC addresses.

Default

no discard-unknown

hash-label

Syntax

hash-label [signal-capability]

no hash-label

Context

config>service>pw-template

Platforms

7210 SAS-Mxp

Description

This command enables the use of hash label on a VLL or VPLS service bound to LDP or RSVP SDP, using the autobind mode with the ldp, rsvp-te, or mpls options. When the hash-label command is enabled, the ingress datapath is modified such that the result of the hash on the packet header is communicated to the egress datapath for use, as the value of the label field of the hash label. Only the hash-2 parameters are used to compute hash-label, even if sdp is over a lag (with load-balancing set as hash-1 or hash-2) or a port. The egress datapath adds the hash label at the bottom of the stack (BoS) and sets the S-bit to one (1).



Note:

On 7210 SAS, the hash label is not used on the local node for purpose of ECMP hashing and LAG hashing. It is available for use by LSR nodes, through which the traffic flows and are capable of using the labels for hashing.

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

The TTL of the hash label is set to a value of 0.

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

- The 7210 local PE inserts the Flow Label Interface Parameters sub-TLV with T=1 and R=1 in the PW ID FEC element in the label mapping message for that spoke-SDP or mesh SDP.
- If remote PE does not send the Flow Label sub-TLV in the PW ID FEC element, or sends a Flow Label sub-TLV in the PW ID FEC element with T=FALSE and R=FALSE, then the local node disables the hash label capability. Therefore the local PE node does not insert a hash label in user and control plane packets it forwards on the spoke-SDP or mesh SDP. It also drops user and control plane packets received from remote PE if they include a hash label. Note that the latter may be caused by a remote 7210 PE which does not support the hash-label option, or which has the hash-label option enabled but does not support the signal-capability option, or does support both options but the user did not enable them because of a mis-configuration.
- If remote PE sends Flow Label sub-TLV in the PW ID FEC element with T=TRUE and R=TRUE, then the local PE enables the hash label capability. Therefore local PE inserts a hash label in user and control plane packets it forwards on the spoke-SDP or mesh SDP. It also accepts user and control plane packets remote PE, with or without hash label
 - If the hash-label option was enabled on the local configuration of the spoke-SDP or mesh SDP at the remote PE, the pseudowire packets received by the local PE have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which results in the insertion of the hash label by both PE nodes.
 - If the hash-label option is not supported or was not enabled on the local configuration of the spoke-SDP or mesh SDP at the remote PE, the pseudowire received by the local PE does not have the hash label included.

The user can enable or disable the signal-capability option in CLI as needed. When doing so, the router must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the PW ID FEC element.



Note:

- This feature is supported only for VLL and VPLS services. It is not supported for VPRN services. It is also not supported on multicast packets forwarded using RSVP P2MP LPS or mLDLP LSP in both the base router instance and in the multicast VPN (mVPN) instance.
- In 7x50 and possibly other vendor implementations, to allow applications where the egress LER infers the presence of the hash label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label is always in the range [524,288 - 1,048,575] and does not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label does not match a value in the reserved label range. This is not supported on 7210 for service traffic (for MPLS OAM traffic the MSB bit is set). That is, 7210 SAS devices do not set the MSB bit in the hash label value for service traffic. Therefore, user must ensure that both the ends are correctly configured to either process hash labels or disable it.

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

Default

no hash-label

Parameters

signal-capability

Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes.

limit-mac-move

Syntax

limit-mac-move [**blockable** | **non-blockable**]

no limit-mac-move

Context

config>service>pw-template

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command indicates whether the mac-move agent limits the MAC relearn (move) rate.

Default

blockable

Parameters

blockable

Specifies that the agent monitors the MAC relearn rate, and block it when the relearn rate is exceeded.

non-blockable

When specified, a SAP is not blocked, and another blockable SAP is blocked instead.

vc-type

Syntax

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

Context

config>service>pw-template

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command overrides the default VC type signaled for the binding to the far end SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment.

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

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

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

Parameters

ether

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

vlan

Specifies the VC type as VLAN. The **ether** and **vlan** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke-SDP bindings.

vlan-vc-tag

Syntax

vlan-vc-tag 0..4094

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

Context

config>service>pw-template

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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.

adv-mtu-override**Syntax**

[no] adv-mtu-override

Context

config>service>sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

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

The **no** form of the command disables the VC-type MTU override and reverts to the default value.

Default

no adv-mtu-override

bgp-tunnel**Syntax**

[no] bgp-tunnel

Context

config>service>sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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



Note:

The 7210 SAS provides an option to install labels for only those BGP 3107 labeled routes which are in use by services. For more information about this option, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Routing Protocols Guide*.

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

Default

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

far-end

Syntax

far-end *ip-address* | *ipv6-address* **node-id** *node-id* [**global-id** *global-id*]

no far-end

Context

config>service>sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the system IP address of the far-end destination 7210 SAS router for the Service Distribution Point (SDP) that is the termination point for a service.

The far-end IP address must be explicitly configured. The destination IP address must be a 7210 SAS system IP address.

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

If the SDP uses MPLS encapsulation, the far-end ip-address is used to check LSP names when added to the SDP. If the "to IP address" defined within the LSP configuration does not exactly match the SDP far-

end ip-address, the LSP is not added to the SDP and an error is generated. Alternatively, an SDP that uses MPLS can have an MPLS-TP node with an MPLS-TP node-id and (optionally) global-id. In this case, the SDP must use an MPLS-TP LSP and the SDP signaling parameter must be set to off.

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

The **no** form of this command removes the currently configured destination IP address for the SDP. The *ip-address* parameter is not specified and generates an error if used in the **no far-end** command. The SDP must be administratively disabled using the **config service sdp shutdown** command before the **no far-end** command can be executed. Removing the far end IP address causes all *lsp-name* associations with the SDP to be removed.

Parameters

ip-address

Specifies the system address of the far-end 7210 SAS devices for the SDP in dotted-decimal notation.

ipv6-address

Specifies the system address of the far-end 7210 SAS devices for the SDP in dotted-decimal notation.

node-id *node-id*

Specifies the MPLS-TP Node ID of the far-end system for the SDP, either in dotted-decimal notation (a.b.c.d) or an unsigned 32-bit integer (1 – 4294967295). This parameter is mandatory for an SDP using an MPLS-TP LSP.

global-id *global-id*

Specifies the MPLS-TP Global ID of the far-end system for the SDP, in an unsigned 32-bit integer (0 – 4294967295). This parameter is optional for an SDP using an MPLS-TP LSP. If not entered, a default value for the Global ID of '0' is used. A global ID of '0' indicates that the far end node is in the same domain as the local node. The user must explicitly configure a Global ID if its value is non-zero.

metric

Syntax

metric *metric*

no metric

Context

config>service>sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

Parameters

metric

Specifies the SDP metric.

Values 0 to 65535

mixed-lsp-mode

Syntax

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

Context

config>service>sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

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

At any time, the service manager programs only one type of LSP in the line card that activates it to forward service packets according to the following priority order:

1. RSVP LSP type. One RSVP LSP can be configured per SDP. This is the highest priority LSP type.
2. LDP LSP type. One LDP FEC is used per SDP. 7210 SAS does not support LDP ECMP.
3. BGP LSP type. One RFC 3107-labeled BGP prefix programmed by the service manager.

In the case of the RSVP/LDP SDP, the service manager programs the NHLFEs for the active LSP type preferring the RSVP LSP type over the LDP LSP type. If no RSVP LSP is configured, or all configured RSVP LSPs go down, the service manager reprograms the line-card with the LDP LSP if available. If not, the SDP goes operationally down.

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

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



Note:

For more information about the **`configure>router>ldp>tunnel-down-damp-time`** command, see the *7210 SAS-Mxp, R6, R12, S, Sx, T MPLS Guide*.

If the user changes the value of the `sdp-revert-time` timer, it takes effect only at the next use of the timer. Any timer that is outstanding at the time of the change is restarted with the new value.

In the case of the LDP/BGP SDP, the service manager gives preference to the LDP LSP type over the BGP LSP type. The service manager reprograms the line card with the BGP LSP if available, otherwise it brings down the SDP operationally.

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

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

Default

`no mixed-lsp-mode`

revert-time

Syntax

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

`no revert-time`

Context

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

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

Default

0

Parameters

revert-time

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

Values 0 to 600

infinite

Sets the SDP to never revert to another higher priority LSP type unless the currently active LSP type is down.

ldp

Syntax

[no] ldp

Context

config>service>sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

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

Default

no ldp (disabled)

lsp

Syntax

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

Context

config>service>sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates associations between one label switched paths (LSPs) and an Multi-Protocol Label Switching (MPLS) Service Distribution Point (SDP). This command is implemented only on MPLS-type encapsulated SDPs.

In MPLS SDP configurations either one LSP can be specified.

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

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

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

Parameters

lsp-name

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

signaling

Syntax

signaling {**off** | **tldp** | **bgp**}

Context

config>service>sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

To modify the signaling configuration, the SDP must be administratively shut down so the signaling parameter can be modified and re-enabled.

Default

tldp

Parameters

off

Specifies that ingress and egress signal auto-labeling is not enabled. If this parameter is selected, each service using the specified SDP must manually configure VPN labels. This configuration is independent of the SDP transport type, MPLS (RSVP or LDP).

tldp

Specifies that ingress and egress pseudowire signaling using T-LDP is enabled.

bgp

Specifies that ingress and egress pseudowire signaling using BGP is enabled. This option is the default value used when BGP VPLS automatically instantiates the SDP.

sr-isis

Syntax

[no] sr-isis

Context

config>service>sdp

Platforms

7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

Description

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

The **no** form of this command disables the use of the IS-IS segment routing LSP type for an MPLS SDP.

Default

no sr-isis

sr-ospf

Syntax

[no] sr-ospf

Context

config>service>sdp

Platforms

7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

Description

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

The **no** form of this command disables the use of the OSPF segment routing LSP type for an MPLS SDP.

Default

no sr-ospf

path-mtu

Syntax

path-mtu *bytes*

no path-mtu

Context

config>service>sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

If the physical **mtu** on an egress interface indicates the next hop on an SDP path cannot support the current **path-mtu**, the operational **path-mtu** on that SDP is 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 uses the system default for the SDP type.

Default

path-mtu defined on the system for the type of SDP is used

2.15.2.1.5 SDP keepalive commands

keep-alive

Syntax

keepalive

Context

config>service>sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables 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 is immediately brought operationally down. When a response is received that indicates the error has cleared and the **hold-down-time** interval has expired, the SDP ID is eligible to be put into the operationally up state. If no other condition prevents the operational change, the SDP ID enters 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 keepalive interpretation of SDP echo reply response conditions and the effect on the SDP ID operational status.

Table 12: Keepalive request results

Result of request	Stored response state	Operational state
keepalive request timeout without reply	Request Timeout	Down
keepalive request not sent because of non-existent <i>orig-sdp-id</i> ¹²	Orig-SDP Non-Existent	Down
keepalive request not sent because of 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)

hello-time

Syntax

hello-time *seconds*

no hello-time

Context

config>service>sdp>keep-alive

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

¹² This condition should not occur.

Description

This command configures the time period between SDP keepalive messages on the SDP-ID for the SDP connectivity monitoring messages.

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

Default

hello-time 10 — 10 seconds between keepalive messages

Parameters

seconds

Specifies the time period in seconds between SDP keepalive messages, expressed as a decimal integer.

Values 1 to 3600

hold-down-time

Syntax

hold-down-time seconds

no hold-down-time

Context

config>service>sdp>keep-alive

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the minimum time period the SDP remains in the operationally down state in response to SDP keepalive monitoring. It 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* is immediately brought operationally down. If a keepalive response is received that indicates the error has cleared, the *sdp-id* is eligible to be put into the operationally up state only after the **hold-down-time** interval has expired.

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

Default

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

Parameters

seconds

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

Values 0 to 3600

max-drop-count

Syntax

max-drop-count *count*

no max-drop-count

Context

config>service>sdp>keep-alive

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 is brought operationally down by the keepalive SDP monitoring.

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

Default

max-drop-count 3

Parameters

count

Specifies the number of consecutive SDP keepalive requests that are failed to be sent or replies missed, expressed as a decimal integer.

Values 1 to 5

message-length

Syntax

message-length *octets*

no message-length

Context

config>service>sdp>keep-alive

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the SDP monitoring keepalive request message length transmitted.

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

Default

0 — The message length should be equal to the SDP operating path MTU as configured in the [path-mtu](#) command. If the default size is overridden, the actual size used is the smaller of the operational SDP-ID Path MTU and the size specified.

Parameters

octets

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

Values 40 to 9198

timeout

Syntax

timeout *timeout*

no timeout

Context

config>service>sdp>keep-alive

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

Default

5

Parameters

timeout
Specifies the timeout time, in seconds.

Values 1 to 10

2.15.2.2 Show commands

- [Show service commands](#)
- [Show ETH-CFM commands](#)

2.15.2.2.1 Show service commands

customer

Syntax

customer [*customer-id*] [**site** *customer-site-name*]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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.

site customer-site-name
Specifies the customer site which is an anchor point for an ingress and egress virtual scheduler hierarchy.

Output

The following output is an example of service customer information, and [Table 13: Output fields: customer](#) describes the output fields.

Sample output

```

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

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

Customer-ID : 3
Contact      : Test
Description  : TiMetra Networks
Phone       : (345) 555-1212

Customer-ID : 6
Contact      : Test1
Description  : Epipe Customer
Phone       : (456) 555-1212

Customer-ID : 7
Contact      : Test2
Description  : VPLS Customer
Phone       : (567) 555-1212

Customer-ID : 274
Contact      : TestA
Description  : ABC Company
Phone       : 650 123-4567

Customer-ID : 94043
Contact      : Test Engineer on Duty
Description  : TEST Customer
Phone       : (789) 555-1212
-----
Total Customers : 8
-----
*A:ALA-12#
*A:ALA-12# show service customer 274
=====
Customer 274
=====
Customer-ID : 274
Contact      : Mssrs. Beaucoup
Description  : ABC Company
Phone       : 650 123-4567
-----
Multi Service Site
-----
Site          : west
Description    : (Not Specified)
=====
*A:ALA-12#

```

Table 13: Output fields: customer

Label	Description
Customer-ID	The ID that uniquely identifies a customer.
Contact	The name of the primary contact person.
Description	Generic information about the customer.
Phone	The phone/pager number to reach the primary contact person.
Total Customers	The total number of customers configured.
Site	Multi-service site name. A multi-service customer site is a group of SAPs with common origination and termination points.
Description	Displays information about a specific customer's multi-service site.
Assignment	The port ID, MDA, or card number, where the SAP's that are members of this multi- service site are defined.
I. Sched Pol	The ingress QoS scheduler policy assigned to this multi-service site.
E. Sched Pol	The egress QoS scheduler policy assigned to this multi-service site.
Service-ID	The ID that uniquely identifies a service.
SAP	Specifies the SAP assigned to the service.

fdb-mac

Syntax

fdb-mac [*ieee-address*] [**expiry**]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays the FDB entry for a specific MAC address.

Parameters

ieee-address

Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

expiry

Shows amount of time until MAC is aged out.

Output

The following output is an example of FDB entry information for a specific MAC address, and [Table 14: Output fields: FDB MAC](#) describes the output fields.

Sample output

```
*A:ALA-48# show service fdb-mac
=====
Service Forwarding Database
=====
ServId    MAC                Source-Identifier    Type/Age    Last Change
-----
103       12:34:56:78:90:0f  sap:1/1/7:0        Static      02/02/2009 09:27:57
700       90:30:ff:ff:ff:8f  cpm                Host        02/02/2009 09:27:57
-----
No. of Entries: 2
=====
*A:ALA-48#

*A:ALA-48# show service fdb-mac expiry
=====
Service Forwarding Database
=====
ServId    MAC                Source-Identifier    Type/Expiry    Last Change
-----
103       12:34:56:78:90:0f  sap:1/1/7:0        Static          02/02/2009 09:27:57
700       90:30:ff:ff:ff:8f  cpm                Host            02/02/2009 09:27:57
-----
No. of Entries: 2
=====
*A:ALA-48#
```

Table 14: Output fields: FDB MAC

Label	Description
ServId	Displays the configured service ID.
MAC	Displays the MAC address.
Source-Identifier	Displays the location where the MAC is defined.
Type/Expiry	Static - FDB entries created by management Learned - dynamic entries created by the learning process OAM - entries created by the OAM process

Label	Description
	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 time when the specific row entry was last change.

sdp

Syntax

sdp *sdp-id* **keep-alive-history**

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

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

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

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays SDP information.

If no optional parameters are specified, a summary SDP output for all SDPs is displayed.

Parameters

sdp-id

Specifies the SDP ID for which to display information.

Values 1 to 17407

Default All SDPs

far-end ip-address

Displays only SDPs matching with the specified far-end IP address.

Default SDPs with any far-end IP address

detail

Displays detailed SDP information.

Default SDP summary output

keep-alive-history

Displays the last fifty SDP keepalive events for the SDP.

Default SDP summary output

Output

The following output is an example of service SDP information, and [Table 15: Output fields: SDP](#) describes the output fields.

Sample output

```
*A:ALA-7210M# show service sdp
=====
Services: Service Destination Points
=====
SdpId      Adm MTU    Opr MTU    IP address      Adm  Opr      Del LSP Signal
-----
10         4462      4462      10.20.1.3       Up   Dn NotReady MPLS B TLDP
40         4462      1534      10.20.1.20      Up   Up        MPLS B TLDP
60         4462      1514      10.20.1.21      Up   Up        MPLS B TLDP
100        4462      4462      10.0.0.2        Down Down      MPLS B TLDP
500        4462      4462      10.20.1.50      Up   Dn NotReady MPLS B TLDP
-----
Number of SDPs : 5
=====
*A:ALA-7210M#

*7210SAS>show>service# sdp 1 detail

=====
Service Destination Point (Sdp Id : 1) Details
=====
-----
Sdp Id 1 -0.0.0.0
-----
Description          : (Not Specified)
SDP Id               : 1
Admin Path MTU       : 0
Far End              : 0.0.0.0
Tunnel Far End       : n/a
SDP Source            : manual
Oper Path MTU        : 0
Delivery             : MPLS
LSP Types            : None

Admin State          : Down
Signaling            : TLDP
Acct. Pol            : None
Last Status Change   : 11/04/2099 22:56:41
Last Mgmt Change     : 11/10/2099 15:56:44
Bw BookingFactor     : 100
Oper Max BW(Kbps)    : 0
Net-Domain           : default
Flags                : SdpAdminDown NoSysIPAddr
                    : TranspTunnDown

Mixed LSP Mode Information :
Mixed LSP Mode           : Enabled
P, BGP
Revert Time             : 200
Active LSP Type         : RSVP....also be LD
Revert Count Down      : n/a

KeepAlive Information :
```

```

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

```

```

-----
RSVP/Static LSPs
-----

```

```

Associated LSP List :
No LSPs Associated

```

```

=====
*7210SAS>show>service#

```

Table 15: Output fields: SDP

Label	Description
SDP Id	The SDP identifier.
Description	Displays a text string describing the SDP.
Admin Path MTU	Displays the desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented. The default value of zero indicates that the path MTU should be computed dynamically from the corresponding MTU of the tunnel.
Opr Path MTU	Displays the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented. To be able to bind this SDP to a specific service, the value of this object minus the control word size (if applicable) must be equal to or larger than the MTU of the service, as defined by its service MTU.
Far End	Displays the far end IP address.
Delivery	The type of delivery used by the SDP: MPLS.
IP address	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Adm Admin State	The desired state of the SDP.
Opr Oper State	The operating state of the SDP.
Flags	Specifies all the conditions that affect the operating status of this SDP.
Signal	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.

Label	Description
Signaling	
Last Status Change	The time of the most recent operating status change to this SDP.
Adv. NTU Over	Specifies whether the advertised MTU of a VLL spoke-SDP bind includes the 14-byte Layer 2 header, so that it is backward compatible with pre-2.0 software.
Last Mgmt Change	The time of the most recent management-initiated change to this SDP.
KeepAlive Information	This section displays Keepalive information.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	The length of the SDP echo request messages transmitted on this SDP.
Hello Timeout	The number of seconds to wait for an SDP echo response message before declaring a timeout.
Unmatched Replies	The number of SDP unmatched message replies timer expired.
Max Drop Count	The maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	The amount of time to wait before the keepalive operating status is eligible to enter the alive state.
TX Hello Msgs	The number of SDP echo request messages transmitted because the keepalive was administratively enabled or the counter was cleared.
Rx Hello Msgs	The number of SDP echo request messages received because the keepalive was administratively enabled or the counter was cleared.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field.
Lsp Name	Displays the LSP name.
Time Since Last Transaction	Displays the time of the last transaction.
Signaling	Specifies the signaling type.
Metric	Displays the metric to be used within the Tunnel Table Manager for decision making purposes. When multiple SDPs going to the

Label	Description
	same destination exist, this value is used as a tie-breaker by Tunnel Table Manager users like MP-BGP to select route with lower value.
Acct. Pol	Displays the policy to use to collect accounting statistics on this SDP. The value zero indicates that the agent should use the default accounting policy, if one exists.
Collect Stats	Specifies whether the agent collects accounting statistics for this SDP. When the value is true the agent collects accounting statistics on this SDP.
VLAN VC Etype	Displays the VLAN VC type.
BW Booking Factor	Specifies the value used to calculate the max SDP available bandwidth. The value specifies the percentage of the SDP max available bandwidth for VLL call admission. When the value of is set to zero (0), no new VLL spoke-SDP bindings with non-zero bandwidth are permitted with this SDP. Overbooking, >100% is allowed.
PBB Etype	Displays the Ethertype used in frames sent out on this SDP when specified as vlan for Provider Backbone Bridging frames.
Oper Max BW (Kbps)	Indicates the operational bandwidth in kilo-bits per seconds (Kbps) available for this SDP. The value is determined by the sum of the bandwidth of all the RSVP LSPs used by the SDP.
Avail BW (Kbps)	Indicates the bandwidth that is still free for booking by the SDP bindings on the SDP.
Net-Domain	Specifies the network-domain name configured on this SDP. The default value of this object is the default network-domain.
Egr Interface	Indicates whether all the egress network interfaces that can carry traffic on this SDP are associated with the network-domain configured on this SDP. Not applicable: Indicates that there is no egress network interface that can carry traffic on this SDP. Consistent: Indicates that the network-domains for all the egress network interfaces that can carry traffic on this SDP are consistent. Inconsistent: Indicates that the network-domain for one or more egress network interfaces that can carry traffic on this SDP are inconsistent.
Mixed LSP Mode	Indicates if the SDP is enabled to use mixed-mode-lsp.
Active LSP Type	Displays the LSP type that is currently active and in use to transport service packets. When multiple LSPs are configured

Label	Description
	under the SDP and enabled with the command 'mixed-mode-lsp', the active LSP could be one of the configured ones. It displays RSVP, if the LSP in use is of type RSVP LSP, LDP if the LSP in use is of type LDP LSP and BGP 3107, if LSP if of type RFC 3107 BGP Labeled route LSP.
Revert Time	Specifies the time to wait before reverting back from LDP to the configured LSPs, after having failed over to LDP.
Revert Count Down	Indicates the timer countdown before reverting back from LDP on this SDP. The timer countdown begins after the first configured LSP becomes active.
Flags	Displays all the conditions that affect the operating status of this SDP.
Class Forwarding	Indicates the admin state of class-based forwarding on this SDP. When the value is true, class-based forwarding is enabled.
EnforceDSTELspFc	Specifies whether service manager must validate with RSVP the support of the FC by the LSP.
Default LSP	Specifies the LSP ID that is used as a default when class-based forwarding is enabled on this SDP. This object must be set when enabling class-based forwarding.
Multicast LSP	Displays the LSP ID that all multicast traffic is forwarded on when class-based forwarding is enabled on this SDP. When this object has its default value, multicast traffic is forwarded on an LSP according to its forwarding class mapping.
Number of SDPs	The total number of SDPs displayed according to the criteria specified.

sdp-using

Syntax

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

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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

Specifies the virtual circuit identifier.

Values 1 to 4294967295

far-end ip-address

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

Default Services with any far-end IP address.

Output

The following output is an example of information about services using SDP, and [Table 16: Output fields: SDP-using](#) describes the output fields.

Sample output

```
*A:ALA-7210M# show service sdp-using 300
=====
Service Destination Point (Sdp Id : 300)
=====
SvcId      SdpId      Type Far End      Opr State I.Label  E.Label
-----
1          300:1      Spok 10.0.0.13      Up      131071  131071
2          300:2      Spok 10.0.0.13      Up      131070  131070
100        300:100    Spok 10.0.0.13      Up      131069  131069
101        300:101    Spok 10.0.0.13      Up      131068  131068
-----
Number of SDPs : 4
=====
*A:ALA-7210M#
```

Table 16: Output fields: SDP-using

Label	Description
Svc ID	The service identifier.
Sdp ID	The SDP identifier.
Type	Type of SDP: spoke
Far End	The far end address of the SDP.
Oper State	The operational state of the service.

Label	Description
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.

service-using

Syntax

service-using [**epipe**] [**vpls**] [**b-vpls**] [**m-vpls**] [**sdp** *sdp-id*] [**customer** *customer-id*]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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.

vpls

Displays matching VPLS instances.

b-vpls

Displays matching B-VPLS instances. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

sdp *sdp-id*

Displays only services bound to the specified SDP ID. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

Values 1 to 17407

Default Services bound to any SDP ID.

customer *customer-id*

Displays services only associated with the specified customer ID.

Values 1 to 2147483647

Default Services associated with a customer.

Output

The following output is an example of service information, and [Table 17: Output fields: service-using](#) describes the output fields.

Sample output

```
*7210SAS>show>service# service-using customer 1

=====
Services Customer 1
=====
ServiceId      Type      Adm  Opr  CustomerId Service Name
-----
1              VPLS      Up   Up   1
2              VPLS      Up   Up   1
3              VPLS      Up   Up   1
4              VPLS      Up   Up   1
2147483648     IES       Up   Down 1      _tmnx_InternalIesService
2147483649     intVpls   Up   Down 1      _tmnx_InternalVplsService
-----
Matching Services : 6
=====
*7210SAS>show>service#
```

Table 17: Output fields: service-using

Label	Description
Service Id	The service identifier.
Type	Specifies the service type configured for the service ID.
Adm	The desired state of the service.
Opr	The operating state of the service.
CustomerId	The ID of the customer who owns this service.
Service name	The name of the service.

eth-ring

Syntax

eth-ring [*status*]

eth-ring [*ring-index*] *hierarchy*

eth-ring *ring-index* [*path* {*a* | *b*}]

Context

show

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays the Ethernet rings information.

Parameters***status***

Displays the status information of the Ethernet rings configured on the system.

hierarchy

Displays eth-ring hierarchical relationships.

path {a|b}

Displays information related to the configured Ethernet rings.

ring-index

Specifies the ring index of the Ethernet ring.

Values 1to128

Output

The following outputs are examples of Ethernet ring information, and the associated tables describe the output fields:

- [Sample output — Standard, Table 18: Output fields: Ethernet ring.](#)
- [Sample output — Ethernet ring status, Table 19: Output fields: Ethernet ring status.](#)

Sample output — Standard

```
*A:NS1015C0821>show# eth-ring 10

=====
Ethernet Ring 10 Information
=====
Description      : (Not Specified)
Admin State      : Down          Oper State       : Down
Node ID          : 00:25:ba:03:48:04
Guard Time       : 5 deciseconds RPL Node          : rplNone
Max Revert Time  : 300 seconds   Time to Revert   : N/A
CCM Hold Down Time : 0 centiseconds CCM Hold Up Time : 20 deciseconds
Compatible Version : 2
APS Tx PDU       : N/A
Defect Status     :
Sub-Ring Type     : virtualLink   Interconnect-ID  : N/A

-----
Ethernet Ring Path Summary
-----
Path Port      Raps-Tag      Admin/Oper      Type            Fwd State
```

```

-----
a - - -/- - -
b - - -/- - -
=====
*A:NS1015C0821>show#

```

Table 18: Output fields: Ethernet ring

Label	Description
Description	The ring description
Admin State	Displays the administrative state
Oper State	Displays the operational state
Node ID	Displays the node identifier
Guard Time	Displays the configured guard time
Max Revert time	Displays the configured maximum revert time
CCM Hold down time	Displays the configured CCM Hold down time
APS TX PDU	Displays the APS TX PDU information
Defect Status	Displays the defect status
RPL Node	Displays the RPL node information
Time to revert	Displays the configured time to revert
CCM Hold Up Time	Displays the configured CCM Hold up time
Sub-Ring Type	Displays the sub-ring type information, the sub-ring type can be virtual link or on-virtual link.
Interconnect-ID	Displays the interconnect ID. The ID can be a ring-index ID or VPLS service ID.
Compatible Version	Displays the Ethernet ring version information.

Sample output — Ethernet ring status

```

*A:NS1015C0821>show# eth-ring status

=====
Ethernet Ring (Status information)
=====
Ring  Admin  Oper   Path Information      MEP Information
ID    State   State Path      Tag      State      Ctrl-MEP CC-Intvl Defects
-----
1      Up      Up     a - 1/1/1  100      Up         Yes      100ms  -----
      b - 1/1/2  100      Up         Yes      100ms  -----
10     Down   Down   a - N/A    -         -          -       -       -----
      b - N/A    -         -          -       -       -       -----

```

```

=====
Ethernet Tunnel MEP Defect Legend:
R = Rdi, M = MacStatus, C = RemoteCCM, E = ErrorCCM, X = XconCCM
*A:NS1015C0821>show#

```

Table 19: Output fields: Ethernet ring status

Label	Description
Ring Id	The ring identifier
Admin State	Displays the administrative state
Oper State	Displays the operational state
Path Information	
Path	Displays the path information
Tag	Displays the tag information
State	Displays the state of the path
MEP Information	
Ctrl-MEP	Displays the Ctrl-MEP information
CC-Intvl	Displays the Ctrl-Interval information
Defects	Displays the defects

pw-routing

Syntax

pw-routing {**local-prefix** | **static-route** | **paths** | **all**}

pw-routing route-table [**all-routes**]

pw-routing route-table summary

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays PW routing information at this 7210 node.

Parameters

local-prefix | static-route | paths | all

Displays details of the T-PE prefixes configured on this node, static routes from this node, explicit PW paths configured on this node, or all of these.

route-table [all-routes]

Displays the PW routing table on this node. If all-routes is specified, then the full routing table is displayed.

route-table summary

Displays a summary of the PW routing table for this node.

Output

The following output is an example of PW routing information.

Sample output

```
*A:Dut-C# show service pw-routing local-prefix
=====
Service PW Routing Information
=====
Service PW Routing Local-Prefix RD Information
=====
Local-Prefix          Route-Dist          Community          Adv-Bgp
-----
3:10.20.1.3           100:3              100:3              enabled
                      100:4              100:4              enabled
-----
Local-Prefix Entries found: 1
=====

*A:Dut-C# show service pw-routing static-route
=====
Service PW Routing Information
=====
Service PW Routing Static-Route Information
=====
Prefix                Next-Hop
-----
6:10.20.1.6/64        10.20.1.5
-----
Static Route Entries found: 1
=====

*A:Dut-C# show service pw-routing paths
=====
Service PW Routing Information
=====
Service PW Routing Path Information
=====
Path                  Adm    Hop IP Address
-----
path1_to_F            up     1   10.20.1.5
                      2   10.20.1.2
path1_to_F2           up     1   10.20.1.2
```

```

2 10.20.1.5
-----
Path Entries found: 2
=====

*A:Dut-C# show service pw-routing all
=====
Service PW Routing Information
=====
SPE-Address      : 3:10.20.1.3
Boot Timer       : 10 secs
Boot Timer Remain : 0 secs
Retry Timer      : 30 secs
Retry Count      : 30
=====

Service PW Routing Local-Prefix RD Information
=====
Local-Prefix      Route-Dist      Community      Adv-Bgp
-----
3:10.20.1.3      100:3      100:3      enabled
                  100:4      100:4      enabled
-----
Local-Prefix Entries found: 1
=====

Service PW Routing Static-Route Information
=====
Prefix            Next-Hop
-----
6:10.20.1.6/64    10.20.1.5
-----
Static Route Entries found: 1
=====

Service PW Routing Path Information
=====
Path              Adm      Hop IP Address
-----
path1_to_F        up        1  10.20.1.5
                  2  10.20.1.2
path1_to_F2        up        1  10.20.1.2
                  2  10.20.1.5
-----
Path Entries found: 2
=====

*A:Dut-C# show service pw-routing route-table all-routes
=====
Service PW L2 Routing Information
=====
AII-Type2/Prefix-Len      Next-Hop      Owner      Age
Route-Distinguisher      Community      Best
-----
3:10.20.1.3:0/64          10.20.1.3     local      00h32m08s
0:0                        0:0           yes
3:10.20.1.3:1/96          10.20.1.3     host       00h32m08s
0:0                        0:0           yes
3:10.20.1.3:2/96          10.20.1.3     host       00h32m08s
0:0                        0:0           yes

```

```
3:10.20.1.3:3/96      10.20.1.3      host      00h32m08s
0:0                  0:0            yes
3:10.20.1.3:4/96      10.20.1.3      host      00h32m08s
0:0                  0:0            yes
3:10.20.1.3:5/96      10.20.1.3      host      00h32m08s
0:0                  0:0            yes
3:10.20.1.3:6/96      10.20.1.3      host      00h32m08s
0:0                  0:0            yes
3:10.20.1.3:7/96      10.20.1.3      host      00h32m08s
0:0                  0:0            yes
3:10.20.1.3:8/96      10.20.1.3      host      00h32m08s
0:0                  0:0            yes
3:10.20.1.3:9/96      10.20.1.3      host      00h32m08s
0:0                  0:0            yes
3:10.20.1.3:10/96     10.20.1.3      host      00h32m07s
0:0                  0:0            yes
6:10.20.1.6:0/64      10.20.1.5      static    00h07m33s
0:0                  0:0            yes
6:10.20.1.6:0/64      10.20.1.5      bgp       00h31m34s
100:6                100:6          no
-----
Entries found: 13
=====

*A:Dut-C# show service pw-routing route-table summary
=====
Service PW L2 Routing Summary
=====
Source           Active
-----
BGP               1
Static            1
Host              10
Local             3
-----
Total             15
=====
```

pw-template

Syntax
pw-template

Context
show>service

Platforms
Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description
This command displays pseudowire template information.

Output

The following output is an example of PW template information.

Sample output

```
*A:Dut-B# show service pw-template 1
=====
PW Template Information
=====
PW Tmpl Id       : 1
Use Provisioned Sdp : enabled          VcType           : vlan
Acctg Policy     : default             Collect Stats    : disabled
Mac-Learning     : enabled             Mac-Ageing      : enabled
Discard Unkn Src : disabled            Limit MacMove   : blockable
Mac-Pinning      : disabled            Vlan VcTag      : 4095
MAC Address Limit : no limit           Rest Prot Src Mac: disabled
Auto Learn Mac Prot : disabled         RestProtSrcMacAct: disable
Block On Peer Fault : disabled

SHG
Name             :
Description      : (Not Specified)
Rest Prot Src Mac : disabled           Rest Unprot Dst  : disabled
Auto Learn Mac Prot : disabled         RestProtSrcMacAct: disable

Egress
Mac FilterId     : none                Ip FilterId      : none
Ipv6 FilterId    : none                QoS NetPlcyId   : none
Port RedirectQGrp : none                Instance Id     : none

Ingress
Mac FilterId     : none                Ip FilterId      : none
Ipv6 FilterId    : none                QoS NetPlcyId   : none
Fp RedirectQGrp  : none                Instance Id     : none

IGMP
Fast Leave       : disabled            Import Plcy      : none
Last Memb Intvl  : 10 deci-secs        Max Nbr Grps    : 0
Send Queries     : disabled
Version          : 3

Force VlanVc Fwd : disabled            Control Word     : disabled
Hash Label       : disabled            Hash Lbl Sig Cap : disabled
Last Changed     : 02/12/2013 22:11:49

-----
Included SDP-Groups
-----
red
-----
```

saii-type2-using

Syntax

saii-type2-using *global-id[:prefix[:ac-id]]*

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays the SDP used by a spoke-sdp-fec with a specified FEC129 Type 2 SAII.

Parameters

global-id[:prefix[:ac-id]]
Specifies the switch-point information using SAII-Type2.

Values	
	<global-id[:prefix*> : <global-id>[:<prefix>[:<ac-id>]]
global-id	1..4294967295
prefix	a.b.c.d 1..4294967295
ac-id	1..4294967295

Output

The following output is an example of information about an SDP used by a spoke-SDP FEC with a specified FEC129 type 2 SAII, and [Table 20: Output fields: SDP](#) describes the output fields.

Sample output

```
*A:Dut-E# show service sai-type2-using 3:10.20.1.3:1
=====
Service Switch-Point Information
=====
SvcId      Oper-SdpBind      SAII-Type2
-----
2147483598 17407:4294967195   3:10.20.1.3:1
-----
Entries found: 1
=====
```

Table 20: Output fields: SDP

Label	Description
SvcId	Displays the service ID.

spoke-sdp-fec-using

Syntax

spoke-sdp-fec-using [**spoke-sdp-fec-id** *spoke-sdp-fec-id*] [**saii-type2** *global-id:prefix:ac-id*] [**taii-type2** *global-id:prefix:ac-id*] [**path** *name*] [**expired**] **taii-type2-using** *global-id[:prefix[:ac-id]]*

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays the SDPs used by spoke-sdp-fecs at this node.

Output

The following output is an example of information about SDPs being used by spoke-SDP FECs.

Sample output

```
*A:Dut-C# show service spoke-sdp-fec-using
=====
Service Spoke-SDP-Fec Information
=====
SvcId SpokeSdpFec Oper-SdpBind SAII-Type2
Path TAII-Type2
-----
1 1 17407:4294967245 3:10.20.1.3:1
n/a 6:10.20.1.6:1
2 2 17407:4294967247 3:10.20.1.3:2
n/a 6:10.20.1.6:2
3 3 17407:4294967248 3:10.20.1.3:3
n/a 6:10.20.1.6:3
4 4 17407:4294967249 3:10.20.1.3:4
n/a 6:10.20.1.6:4
5 5 17407:4294967250 3:10.20.1.3:5
n/a 6:10.20.1.6:5
6 6 17407:4294967251 3:10.20.1.3:6
n/a 6:10.20.1.6:6
7 7 17407:4294967252 3:10.20.1.3:7
n/a 6:10.20.1.6:7
8 8 17407:4294967253 3:10.20.1.3:8
n/a 6:10.20.1.6:8
9 9 17407:4294967254 3:10.20.1.3:9
n/a 6:10.20.1.6:9
10 10 17407:4294967255 3:10.20.1.3:10
n/a 6:10.20.1.6:10
-----
Entries found: 10
=====
```

taii-type2-using**Syntax**

taii-type2-using *global-id[:prefix[:ac-id]]*

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays switch-point information using TAI.

Parameters

global-id[:prefix[:ac-id]]
Specifies the switch-point information using SAI-Type2.

Values	
	<global-id[:prefix*> : <global-id>[:<prefix>[:<ac-id>]]
global-id	1..4294967295
prefix	a.b.c.d 1 to 4294967295
ac-id	1 to 4294967295

Output

The following output is an example of information about a switch-point using TAI Type 2.

Sample output

```
*A:Dut-E# show service taii-type2-using 6:10.20.1.6:1
=====
Service Switch-Point Information
=====
SvcId      Oper-SdpBind      TAII-Type2
-----
2147483598 17407:4294967195 6:10.20.1.6:1
-----
Entries found: 1
=====
```

2.15.2.2.2 Show ETH-CFM commands

eth-cfm

Syntax

eth-cfm

Context

show

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context display eth-cfm information.

association

Syntax

association [*ma-index*] [**detail**]

Context

show>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays eth-cfm association information.

Parameters

ma-index

Specifies the maintenance association (MA) index.

Values 1 to 4294967295

detail

Displays more information for the eth-cfm association.

Output

Show eth-cfm association command output

The following output is an example of eth-cfm association information, and [Table 21: Output fields: ETH-CFM association](#) describes output fields.

Sample output

```
A:dut-b# show eth-cfm association
```

CFM Association Table				
Md-index	Ma-index	Name	CCM-interval	Bridge-id
1	1	a1	1	1
1	2	a2	1	2
2	1	a1	1	2

```
2          2          a2          1          1
=====
A:dut-b#
```

Table 21: Output fields: ETH-CFM association

Label	Description
Md-index	Displays the maintenance domain (MD) index.
Ma-index	Displays the maintenance association (MA) index.
Name	Displays the part of the maintenance association identifier which is unique within the maintenance domain name.
CCM-interval	Displays the CCM transmission interval for all MEPs in the association.
Bridge-id	Displays the bridge-identifier value for the domain association.
MHF Creation	Displays the MIP half function (MHF) for the association.
Primary VLAN	Displays the primary bridge-identifier VLAN ID.
Num Vids	Displays the number of VIDs associated with the VLAN.
Remote Mep Id	Displays the remote maintenance association end point (MEP) identifier

cfm-stack-table

Syntax

```
cfm-stack-table [{all-ports}] [level 0..7] [direction <down>]
```

Context

```
show>eth-cfm
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays stack-table information. This stack-table is used to display the various management points MEPs and MIPs that are configured on the system. This can be Service based. The various options allow the operator to be specific. If no parameters are include then the entire stack-table is displayed.

Parameters

- port *port-id*

Displays the bridge port or aggregated port on which MEPs or MHFs are configured.
- vlan *vlan-id*

Displays the associated VLAN ID.
- level

Display the MD level of the maintenance point.

Values0 to 7
- direction down

Displays the direction in which the MP faces on the bridge port.

Output

Show eth-cfm CFM stack table command output

The following output is an example of eth-cfm CFM stack table information, and [Table 22: Output fields: CFM stack table](#) describes the fields.

Sample output

```
*A:7210SAS>show>eth-cfm# cfm-stack-table

=====
CFM SAP Stack Table
=====
Sap          Level Dir  Md-index  Ma-index  Mep-id Mac-address
-----
1/1/18:100   7      Up    7         100       1      00:25:ba:0d:21:13
=====

=====
CFM Ethernet Tunnel Stack Table
=====
Eth-tunnel   Level Dir  Md-index  Ma-index  Mep-id Mac-address
-----
No Matching Entries
=====

=====
CFM SDP Stack Table
=====
Sdp          Level Dir  Md-index  Ma-index  Mep-id Mac-address
-----
No Matching Entries
=====

=====
CFM Virtual Stack Table
=====
Service      Level Dir  Md-index  Ma-index  Mep-id Mac-address
-----
No Matching Entries
=====
*A:7210SAS>show>eth-cfm#
```

Table 22: Output fields: CFM stack table

Label	Description
Sap	Displays associated SAP IDs.
Sdp	Displays the SDP binding for the bridge.
Level Dir	Displays the MD level of the maintenance point.
Md-index	Displays the maintenance domain (MD) index.
Ma-index	Displays the maintenance association (MA) index.
Mep-id	Displays the integer that is unique among all the MEPs in the same MA.
Mac-address	Displays the MAC address of the MP.

domain

Syntax

domain [*md-index*] [**association** *ma-index* | **all-associations**] [**detail**]

Context

show>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays domain information.

Parameters

md-index

Displays the index of the MD to which the MP is associated, or 0, if none.

association *ma-index*

Displays the index to which the MP is associated, or 0, if none.

all-associations

Displays all associations to the MD.

detail

Displays detailed domain information.

Output

Show eth-cfm Domain Command Output

The following output is an example of eth-cfm domain information, and [Table 23: Output Fields: ETH-CFM domain](#) describes the output fields.

Sample output

```
A:dut-b# show eth-cfm domain

=====
CFM Domain Table
=====
Md-index   Level Name                                     Format
-----
1           6    d1                                     charString
2           7    d2                                     charString
=====
A:dut-b#
```

Table 23: Output Fields: ETH-CFM domain

Label	Description
Md-index	Displays the Maintenance Domain (MD) index value.
Level	Displays an integer identifying the Maintenance Domain Level (MD Level). Higher numbers correspond to higher Maintenance Domains, those with the greatest physical reach, with the highest values for customers' CFM PDUs. Lower numbers correspond to lower Maintenance Domains, those with more limited physical reach, with the lowest values for CFM PDUs protecting single bridges or physical links.
Name	Displays a generic Maintenance Domain (MD) name.
Format	Displays the type of the Maintenance Domain (MD) name. Values include dns, mac, and string.

mep

Syntax

```
mep mep-id domain md-index association ma-index [loopback] [linktrace]
mep mep-id domain md-index association ma-index remote-mepid mep-id | all-remote-mepids
mep mep-id domain md-index association ma-index eth-test-results [remote-peer mac-address]
mep mep-id domain md-index association ma-index one-way-delay-test [remote-peer mac-address]
mep mep-id domain md-index association ma-index two-way-delay-test [remote-peer mac-address]
mep mep-id domain md-index association ma-index two-way-slm-test [remote-peer macaddress]
```

Context

show>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays Maintenance Endpoint (MEP) information.



Note:

- The **show eth-cfm mep mep-id domain md-id association ma-id** command does not display CCM ERROR, CCM XCON frames in the output.
- The **show eth-cfm mep mep-id domain md-id association ma-id remote-mep rmep-id** command does not display some TLVs details.

Parameters

mep-id

Displays the integer that is unique among all the MEPs in the same MA.

domain md-index

Displays the index of the MD to which the MP is associated, or 0, if none.

association ma-index

Displays the index to which the MP is associated, or 0, if none.

loopback

Displays loopback information for the specified MEP.

linktrace

Displays linktrace information for the specified MEP.

remote-mepid mep-id

Includes specified remote mep-id information for specified the MEP.

all-remote-mepids

Includes all remote mep-id information for the specified MEP.

eth-test-results

Includes eth-test-result information for the specified MEP.

one-way-delay-test

Includes one-way-delay-test information for the specified MEP.

two-way-delay-test

Includes two-way-delay-test information for the specified MEP.

two-way-slm-test

Includes two-way-slm-test information for the specified MEP.

remote-peer mac-address

Includes specified remote mep-id information for the specified MEP.

Output

The following outputs are examples of MEP information.

Sample output

```
A:dut-b# show eth-cfm mep 1 domain 1 association 1 linktrace
-----
Mep Information
-----
Md-index          : 1                Direction       : Down
Ma-index          : 1                Admin           : Enabled
MepId             : 1                CCM-Enable      : Enabled
IfIndex           : 35946496         PrimaryVid      : 1
FngState          : fngReset          ControlMep      : False
LowestDefectPri   : macRemErrXcon     HighestDefect   : none
Defect Flags      : None
Mac Address       : 00:25:ba:01:c3:6a  CcmLtmPriority  : 7
CcmTx             : 0                CcmSequenceErr  : 0
Eth-1Dm Threshold : 3(sec)
Eth-Ais           : Disabled
Eth-Tst           : Disabled
CcmLastFailure Frame:
None
XconCcmFailure Frame:
None
-----
Mep Linktrace Message Information
-----
LtRxUnexplained   : 0                LtNextSequence  : 2
LtStatus          : False             LtResult        : False
TargIsMepId       : False             TargMepId       : 0
TargMac           : 00:00:00:00:00:00  TTL             : 64
EgressId          : 00:00:00:25:ba:01:c3:6a  SequenceNum     : 1
LtFlags           : useFDBOnly
-----
Mep Linktrace Replies
-----
SequenceNum       : 1                ReceiveOrder    : 1
Ttl               : 63                Forwarded       : False
LastEgressId      : 00:00:00:25:ba:01:c3:6a  TerminalMep    : True
NextEgressId      : 00:00:00:25:ba:00:5e:bf  Relay          : rlyHit
ChassisIdSubType  : unknown value (0)
ChassisId:
None
ManAddressDomain:
None
ManAddress:
None
IngressMac        : 00:25:ba:00:5e:bf        Ingress Action  : ingOk
IngrPortIdSubType : unknown value (0)
IngressPortId:
None
EgressMac         : 00:00:00:00:00:00        Egress Action   : egrNoTlv
EgrPortIdSubType  : unknown value (0)
EgressPortId:
None
Org Specific TLV:
None
A:dut-b#
A:dut-b#

A:dut-b# show eth-cfm mep 1 domain 1 association 1 loopback
-----
```

Mep Information

```

-----
Md-index      : 1                      Direction      : Down
Ma-index      : 1                      Admin          : Enabled
MepId         : 1                      CCM-Enable     : Enabled
IfIndex       : 35946496              PrimaryVid     : 1
FngState      : fngReset              ControlMep     : False
LowestDefectPri : macRemErrXcon        HighestDefect   : none
Defect Flags   : None
Mac Address    : 00:25:ba:01:c3:6a      CcmLtmPriority  : 7
CcmTx         : 0                      CcmSequenceErr : 0
Eth-1Dm Threshold : 3(sec)
Eth-Ais       : Disabled
Eth-Tst       : Disabled
CcmLastFailure Frame:
None
XconCcmFailure Frame:
None
-----

```

Mep Loopback Information

```

-----
LbRxReply     : 1                      LbRxBadOrder   : 0
LbRxBadMsdu   : 0                      LbTxReply      : 0
LbSequence    : 2                      LbNextSequence : 2
LbStatus      : False                  LbResultOk     : True
DestIsMepId   : False                  DestMepId      : 0
DestMac       : 00:00:00:00:00:00      SendCount      : 0
VlanDropEnable : True                  VlanPriority    : 7
Data TLV:
None
A:dut-b#

```

*A:dut-b# show eth-cfm mep 1 domain 4 association 4 two-way-delay-test remote-peer 00:25:ba:00:5e:bf

Eth CFM Two-way Delay Test Result Table

```

=====
Peer Mac Addr      Delay (us)      Delay Variation (us)
-----
00:25:ba:00:5e:bf  507             507
=====

```

*A:dut-b#

*A:dut-b# show eth-cfm mep 1 domain 4 association 4 two-way-delay-test

Eth CFM Two-way Delay Test Result Table

```

=====
Peer Mac Addr      Delay (us)      Delay Variation (us)
-----
00:25:ba:00:5e:bf  507             507
=====

```

*A:dut-b#

*A:dut-a# show eth-cfm mep 2 domain 4 association 4 eth-test-results remote-peer 00:25:ba:01:c3:6a

Eth CFM ETH-Test Result Table

```

=====
FrameCount      Current      Accumulate
ErrBits          ErrBits

```

```

Peer Mac Addr      ByteCount      CrcErrs      CrcErrs
-----
00:25:ba:01:c3:6a 6              0              0
384                0              0
=====
*A:dut-a#

*A:dut-a# show eth-cfm mep 2 domain 4 association 4 eth-test-results

=====
Eth CFM ETH-Test Result Table
=====
Peer Mac Addr      FrameCount      Current      Accumulate
ByteCount          ErrBits         ErrBits
CrcErrs            CrcErrs
-----
00:25:ba:01:c3:6a 6              0              0
384                0              0
=====

*A:dut-a# show eth-cfm mep 2 domain 4 association 4 one-way-delay-test remote-
peer 00:25:ba:01:c3:6a

=====
Eth CFM One-way Delay Test Result Table
=====
Peer Mac Addr      Delay (us)      Delay Variation (us)
-----
00:25:ba:01:c3:6a 402             402
=====

*A:dut-a#

*A:dut-a# show eth-cfm mep 2 domain 4 association 4 one-way-delay-test

=====
Eth CFM One-way Delay Test Result Table
=====
Peer Mac Addr      Delay (us)      Delay Variation (us)
-----
00:25:ba:01:c3:6a 402             402
=====

*A:dut-a#

```

Show output for two-way-slm-test

```

*A:7210SAS# show eth-cfm mep 1 domain 7 association 100 two-way-slm-test
=====
Eth CFM Two-way SLM Test Result Table (Test-id: 1)
=====
Peer Mac Addr      Remote MEP      Count      In Loss      Out Loss      Unack
-----
00:25:ba:0d:1e:12 2              1          0            0            0
=====
*A:7210SAS#

```

connection-profile

Syntax

connection-profile [conn-prof-id] [associations]

Context

show

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays connection profile information.

Parameters

conn-prof-id

Specifies the connection profile ID.

Values 1 to 8000

associations

Displays the SAP and the service ID that use this connection profile.

Output

The following output is an example of connection-profile information, and [Table 24: Output fields: connection profile](#) describes output fields.

Sample output — Connection profile

```
*7210SAS>show# connection-profile

=====
Connection Profile Summary Information
=====
CP Index  Number of HasRange
         Members
-----
1          0      Yes
2          0      Yes
3          0      Yes
5          0      Yes
6          0      Yes
100        0      Yes
200        0      Yes
300        0      Yes
400        0      Yes
500        0      Yes
600        0      Yes
700        0      Yes
```

```
800      0      Yes
900      0      Yes
=====
*7210SAS>show#
```

Show output for connection-profile associations

```
*A:7210SAS>show# connection-profile associations

=====
Connection Profile Summary Information
=====
CP Index  Number of HasRange
          Members
-----
1          0          No
=====
*A:7210SAS>show#
```

Table 24: Output fields: connection profile

Label	Description
CP Index	Identifies the connection-profile.
Number of Members	Indicates the number of ATM connection profile members not applicable for 7210.
HasRange	Indicates whether VLAN range is configured.

2.15.2.3 Tools commands

2.15.2.3.1 Tools perform commands

```
tools
```

Syntax
tools

Context
root

Platforms
Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context enable useful tools for debugging purposes.

Parameters**dump**

Enables dump tools for the various protocols.

perform

Enables tools to perform specific tasks.

perform**Syntax**

perform

Context

tools

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context enable tools to perform specific tasks.

service**Syntax**

services

Context

tools>perform

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure tools for services.

id

Syntax

id *service-id*

Context

tools>perform>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure tools for a specific service.

Parameters

service-id

Specify an existing service ID.

Values 1 to 2147483647

endpoint

Syntax

endpoint *endpoint-name*

Context

tools>perform>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures tools for a specific VLL service endpoint.

Parameters

endpoint-name

Specify an existing VLL service endpoint name.

force-switchover

Syntax

force-switchover *sdp-id:vc-id*

no force-switchover

force-switchover spoke-sdp-fec [1..4294967295]

Context

tools>perform>service>id>endpoint

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

Parameters

sdp-id:vc-id

Specify an existing spoke-SDP for the service.

spoke-sdp-fec

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

Values 1 to 4294967295

Output

The following output is an example of force switchover information.

Sample Output

```
*A:Dut-B# show service id 1 endpoint
```

```
=====
Service 1 endpoints
=====
```

```
Endpoint name : mcep-t1
Description : (Not Specified)
Revert time : 0
Act Hold Delay : 0
Ignore Standby Signaling : false
Suppress Standby Signaling : false
Block On Mesh Fail : true
Multi-Chassis Endpoint : 1
MC Endpoint Peer Addr : 10.1.1.3
Psv Mode Active : No
Tx Active : 221:1(forced)
Tx Active Up Time : 0d 00:00:17
Revert Time Count Down : N/A
Tx Active Change Count : 6
```



```
Last Tx Active Change : 02/14/2009 00:17:32
```

```
-----  
Members  
-----
```

```
Spoke-sdp: 221:1 Prec:1 Oper Status: Up
```

```
Spoke-sdp: 231:1 Prec:2 Oper Status: Up  
=====
```

```
*A:Dut-B#
```

eval-pw-template

Syntax

eval-pw-template

Context

tools>perform>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command reevaluates the pseudowire template policy.

Parameters

policy-id

Specifies the pseudowire template policy.

eval-expired-fec

Syntax

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

eval-expired-fec all

Context

tools>perform>service>pw-routing

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

spoke-sdp-fec-release

Syntax

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

Context

tools>perform>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command is used to clear the MS-PW bindings associated with particular SAll or TAll on an S-PE.

3 VLL services

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

3.1 Epipe services

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


3.1.1 Epipe services overview

An Epipe service is a Layer 2 point-to-point service where the customer data is encapsulated and transported across a service provider network. An Epipe service is completely transparent to the subscriber data and protocols. The Epipe service does not perform any MAC learning. A local Epipe service consists of two SAPs on the same node, whereas a distributed Epipe service consists of two SAPs on different nodes.

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

The following figure shows Epipe/VLL customer service.

Figure 17: Epipe/VLL service



OSSG021

3.1.2 Epipe with PBB



Note:

Epipes with PBB are supported only on 7210 SAS-T operating in the network mode.

A PBB tunnel may be linked to an Epipe to a B-VPLS. MAC switching and learning is not required for the point-to-point service (all packets ingressing the SAP are PBB encapsulated and forwarded to the PBB

tunnel to the backbone destination MAC address, and all the packets ingressing the B-VPLS destined for the ISID are PBB de-encapsulated and forwarded to the Epipe SAP).

A fully specified backbone destination address must be provisioned for each PBB Epipe instance to be used for each incoming frame on the related I-SAP. If the backbone destination address is not found in the B-VPLS FDB, packets may be flooded through the B-VPLSs

All B-VPLS constructs may be used, including B-VPLS resiliency and OAM. Not all generic Epipe commands are applicable when using a PBB tunnel.

3.1.3 Processing packets received with more than two tags on a QinQ SAP in Epipe service



Note:

This section applies to all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

The 7210 SAS platforms operating in access-uplink mode can process and forward packets with more than two tags. See [Configuration notes](#) for restrictions on the use of SAPs by platforms operating in access-uplink mode.

To forward packets with two or more tags using a QinQ SAP, an Epipe service type is available for use when 7210 SAS devices are operating in network mode. This service allows for configuration of a QinQ SAP as one endpoint and the following service entities as the other endpoint:

- MPLS spoke-SDP with **vc-type** set to **vlan**
 - The vc-vlan-tag must match the inner-tag VLAN ID value specified in the QinQ SAP.
- dot1q SAP
 - The VLAN value configured for the dot1q SAP must match the inner-tag VLAN ID value of the QinQ SAP.
- QinQ SAP
 - The inner VLAN tag of both QinQ SAPs configured in the service must be the same.

In the forward direction, the device will process the packet as follows:

- If the packet is received on a QinQ SAP, assign an incoming packet to this service based on matching the outermost two tags in the packet header (that is, the first two tags in the packet header). It will strip only the outermost tag (single tag) on ingress and forward the rest to the other endpoint in the service.
- If the other endpoint the packet is sent out of is an MPLS SDP, MPLS encapsulation is added.
- If the other endpoint that the packet is sent out of is a dot1q SAP, the packet is forwarded as is, without any egress VLAN checks. It is expected that the operator will ensure that the inner tag of the packet matches the dot1q VLAN value.
- If the other endpoint that the packet is sent out of is another QinQ SAP (for example, Q1.Q2 SAP), another tag (that is, Q2 tag) is added to the packet and sent out of the QinQ SAP.

In the reverse direction, the device will process the packet as follows:

- If the packet is received on an MPLS SDP, the vc-vlan tag is retained as is and the VLAN tag corresponding to the outermost tag configured for the QinQ SAP (that is, the other endpoint) is added to the packet. The system does not match the vc-vlan tag received in the packet with the configured value

(that is, the inner tag of the QinQ SAP). It is expected that the operator will configure both end of the service appropriately to ensure that only appropriate packets enter the service.

- If the packet is received on a dot1q SAP, the outermost tag is not stripped and the VLAN tag corresponding to the outermost tag configured for the QinQ SAP is added to the packet.
- If the packet is received on a QinQ SAP, assign an incoming packet to this service based on matching the outermost two tags in the packet header (that is, the first two tags in the packet header). It will strip only the outermost tag (single tag) on ingress. The VLAN tag corresponding to the outermost tag configured for the QinQ SAP (that is, the other endpoint) is added to the packet, which is forwarded from the QinQ SAP.

Therefore, the device processes packets received with two or more tags using the MPLS SDP or a dot1q SAP while classifying on the QinQ SAP ingress using two tags.

3.1.3.1 Feature support, configuration notes, and restrictions

An **svc-sap-type** value **qinq-inner-tag-preserve** is available for configuring the service. This must be used when creating a new Epipe service if this functionality is needed (for example, **epipe 10 svc-sap-type qinq-inner-tag-preserve create**):

- This service is available only in network mode.
- Epipe service created with the parameter **svc-sap-type** set to **qinq-inner-tag-preserve** will allow for only one QinQ SAP and only one SDP of **vc-type vlan**. The system will not allow the use of any other SAP in this new service; that is, null SAP, Q1. * SAP, 0.* SAP, and so on, are not allowed for configuration in this service. The SDP cannot be of **vc-type ether**.
- The user can configure the **vlan-vc** tag value for the SDP, the dot1q SAP VLAN tag value, and the inner tag VLAN value of a QinQ SAP to match the VLAN ID value of the inner tag specified in the Q1.Q2 SAP configured in the service. For example, if the SAP is 1/1/10:Q1.Q2, the **vlan-vc** tag must be set to Q2, the dot1q SAP VLAN value must be Q2, and the inner tag of another QinQ SAP must be set to Q2.

If any other value, other than the QinQ SAP inner tag is configured for the **vlan-vc** tag or dot1q SAP VLAN value, or for the inner tag of the QinQ SAP, it will be errored out by the software. If the **vlan-vc** tag value is not configured, it defaults to use the inner VLAN tag value. It is highly recommended that the customer configure the **vlan-vc-tag** value to match the VLAN ID value of the inner tag configured for the QinQ SAP, to avoid misconfiguration.

- Existing QoS and ACL functionality for the Epipe service entities are available, with the following exceptions:
 - If the packet is received with more than two tags, IP match-criteria cannot be used with SAP ingress QoS classification and ACLs (both ingress and egress ACLs).
 - If the packet is received with more than two tags, the Ethertype value in the mac-criteria cannot be used with SAP ingress QoS classification and ACLs (both ingress and egress ACLs).
 - Dot1p bits from the outermost tag (that is, Q1 VLAN tag, if the SAP is 1/1/10:Q1.Q2) will be used for SAP ingress classification. Dot1p bits of the outermost tag will be marked on egress, if marking is enabled on the egress port. The Dot1p bit value of the **vc-vlan-tag** is not used to mark the Dot1p bits of the outermost VLAN tag, when the packets is exiting the QinQ SAP.
- OAM tools
 - MPLS OAM tools such as **vccv-ping**, **vccv-trace**, and so on, are supported for the SDPs.
 - Accounting and Statistics for the service entities (for example, SAP and SDP) will be available as before.

- CFM/Y.1731 tools are supported. Up and Down MEP is supported on the SAPs and the SDPs configured in the Epipe service.
- The following redundancy mechanisms available in Epipe service are supported when using MPLS SDP:
 - Epipe PW redundancy
 - MC-LAG based protection for access SAPs using the service type (along with using PW redundancy)

3.1.3.2 Configuration example of Epipe services for processing of packets received with more than two tags on a QinQ SAP

Example

The following is a sample output of an Epipe service with the vlan-vc-tag value configured to match the inner tag specified in the Q1.Q2 SAP in the service.

```
*A:7210SAS>config>service# info
-----
epipe 10 svc-sap-type qinq-inner-tag-preserve customer 1 create
      sap 1/1/3:10.45 create
      exit
      spoke-sdp 111:69 vc-type vlan create
      vlan-vc-tag 45
      exit
      no shutdown
-----
```

Example

The following is a sample output of an Epipe service with QinQ SAP and dot1q SAP. The dot1q SAP (1/1/4:45) VLAN value 45, matches the inner tag VLAN value specified with QinQ SAP (1/1/3:10.45).

```
*A:7210>config>service# info
-----
epipe 10 svc-sap-type qinq-inner-tag-preserve customer 1 create
      sap 1/1/3:10.45 create
      no shutdown
      exit
      sap 1/1/4:45 create
      no shutdown
      exit
      no shutdown
      exit
-----
```

Example

The following is a sample output of an Epipe service with two QinQ SAPs. The inner tag of both QinQ SAPs matches and is set to a value of 45.

```
*A:7210>config>service# info
-----
-----
epipe 10 svc-sap-type qinq-inner-tag-preserve customer 1 create
```

```

sap 1/1/3:10.45 create
    no shutdown
exit
sap 1/1/4:200.45 create
    no shutdown
exit
    no shutdown
exit
-----
-----

```

3.1.4 Epipe operational state decoupling

An Epipe service transitions to an operational state of down when only a single entity SAP or binding is active and the operational state of the mate is down or displays an equivalent state. The default behavior does not allow operators to validate the connectivity and measure performance metrics. With this feature, an option is provided to allow operators to validate the connectivity and measure performance metrics of an Epipe service before the customer handoff. The operator can also maintain performance and continuity measurement across their network regardless of the connectivity between the terminating node and the customer.

If the SAP between the operator and the customer enters a oper down state, the Epipe remains operationally up, so the results can continue to be collected uninterrupted. The operator receives applicable port or SAP alerts and alarms. This option is available only for the customer-facing SAP failures. If a network-facing SAP or spoke-SDP, fails the operational state of the Epipe service is set to down. That is, there is no option to hold the service in an up state, if a network component fails.

The following functionality is supported:

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

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

3.2 Pseudowire switching



Note:

The 7210 SAS-T can only be configured as T-PE nodes.

The 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE can be configured as either T-PE or S-PE nodes.

The pseudowire switching feature provides the user with the ability to create a VLL service by cross-connecting two spoke-SDPs. This feature allows the scaling of VLL and VPLS services in a large network in which the otherwise full mesh of PE devices would require thousands of Targeted LDP (T-LDP) sessions per PE node.

Services with one SAP and one spoke-SDP are usually created on the PE; however, the target destination of the SDP is the pseudowire switching node instead of what is usually the remote PE.

The pseudowire switching node acts in a passive role with respect to signaling of the pseudowires. It waits until one or both of the PEs sends the label mapping message before relaying it to the other PE. This is because it needs to pass the interface parameters of each PE to the other.

A pseudowire switching point TLV is inserted by the switching pseudowire to record its system address when relaying the label mapping message. This TLV is useful in a few situations:

- It allows for troubleshooting of the path of the pseudowire, especially if multiple pseudowire switching points exist between the two PEs.
- It helps in loop detection of the T-LDP signaling messages where a switching point would receive back a label mapping message it had already relayed.
- It is inserted in pseudowire status notification messages when they are sent end-to-end or from a pseudowire switching node toward a destination PE.

Pseudowire OAM is supported for the manual switching pseudowires and allows the pseudowire switching node to relay end-to-end pseudowire status notification messages between the two PEs. The pseudowire switching node can generate a pseudowire status and send it to one or both of the PEs by including its system address in the pseudowire switching point TLV. This allows a PE to identify the origin of the pseudowire status notification message.


In the following pseudowire service switching node example, the user configures a regular Epipe VLL service PE1 and PE2. These services each consist 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 an Epipe VLL service on the pseudowire switching node using the two SDPs.

```
|7210 PE1 (Epipe)|---sdp 2:10---|7210 PW SW (Epipe)|---sdp 7:15---|7210 PE2 (Epipe)|
```


3.2.1 Pseudowire switching with protection

Pseudowire switching scales VLL and VPLS services over a multi-area network by removing the need for a full mesh of targeted LDP sessions between PE nodes. The following figure shows the use of pseudowire redundancy to provide a scalable and resilient VLL service across multiple IGP areas in a provider network.

Figure 18: VLL Resilience with pseudowire redundancy and switching



In the network in the preceding figure, PE nodes act as leading nodes and pseudowire switching nodes act as followers for the purpose of pseudowire signaling. A switching node will need to pass the SAP interface parameters of each PE to the other. T-PE1 sends a label mapping message for the Layer 2 FEC to the peer pseudowire switching node; for example, S-PE1.

It will include the SAP interface parameters, such as MTU, in the label mapping message. S-PE1 checks the FEC against the local information and if a match exists, it appends the optional pseudowire switching point TLV to the FEC TLV in which it records its system address. T-PE1 then relays the label mapping message to S-PE2. S-PE2 performs similar operations and forwards a label mapping message to T-PE2.

The same procedures are followed for the label mapping message in the reverse direction; for example, from T-PE2 to T-PE1. S-PE1 and S-PE2 will affect the spoke-SDP cross-connect only when both directions of the pseudowire have been signaled and matched.

The pseudowire switching TLV is useful in a few situations. First, it allows for troubleshooting of the path of the pseudowire, especially if multiple pseudowire switching points exist between the two T-PE nodes. Second, it helps in loop detection of the T-LDP signaling messages where a switching point receives back

a label mapping message it already relayed. Finally, it can be inserted in pseudowire status messages when they are sent from a pseudowire switching node toward a destination PE.

Pseudowire status messages can be generated by the T-PE nodes. Pseudowire status messages received by a switching node are processed and then passed on to the next-hop. An S-PE node appends the optional pseudowire switching TLV, with its system address added to it, to the FEC in the pseudowire status notification message, only if it originated the message or the message was received with the TLV in it. Otherwise, the message was originated by a T-PE node and the S-PE should process and pass the message without changes except for the VC ID value in the FEC TLV.

3.2.2 Pseudowire switching behavior

In the network in [Figure 18: VLL Resilience with pseudowire redundancy and switching](#), PE nodes act as leading nodes and pseudowire switching nodes act as followers for the purpose of pseudowire signaling. A switching node must pass the SAP interface parameters of each PE to the other. T-PE1 sends a label mapping message for the Layer 2 FEC to the peer pseudowire switching node; for example, S-PE1.

It 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 similar operations and forwards a label mapping message to T-PE2.

The same procedures are followed for the label mapping message in the reverse direction; for example, from T-PE2 to T-PE1. S-PE1 and S-PE2 will affect the spoke-SDP cross-connect only when both directions of the pseudowire have been signaled and matched.


The merging of the received T-LDP status notification message and the local status for the spoke-SDPs from the service manager at a PE complies with the following rules:

- When the local status for both spokes is up, the S-PE passes any received SAP or SDP-binding generated status notification message unchanged; for example, the status notification TLV is unchanged but the VC ID in the FEC TLV is set to the value of the pseudowire segment to the next hop.
- When the local operational status for any of the spokes is down, the S-PE always sends SDP-binding down status bits, regardless of whether the received status bits from the remote node indicated SAP up/down or SDP-binding up/down.

3.2.2.1 Pseudowire switching TLV

The following figure shows the format of the pseudowire switching TLV.

Figure 19: Pseudowire switching TLV format



hw0507

- **PW sw TLV Length**

This field specifies the total length of all the following pseudowire switching point TLV fields in octets.

- **Type**

This field encodes how the Value field is interpreted.

- **Length**

This field specifies the length of the Value field in octets.

- **Value**

The octet string of Length octets encodes information to be interpreted as specified by the Type field.

The format of the pseudowire switching point sub-TLVs is as follows:

- **pseudowire ID of last pseudowire segment traversed**

This sub-TLV type contains a pseudowire ID in the format of the pseudowire ID.

- **pseudowire switching point description string**

This is an optional description text string up to 80 characters.

- **IP address of pseudowire switching point**

- **IPv4 or IPv6 address of pseudowire switching point**

This is an optional sub-TLV.

- **MH VCCV capability indication**

3.2.2.2 Static-to-dynamic pseudowire switching

When one segment of the pseudowire cross-connect at the S-PE is static while the other is signaled using T-LDP, the S-PE operates much like a T-PE from a signaling perspective and as an S-PE from a data plane perspective.

The S-PE signals a label mapping message as soon as the local configuration is complete. The control word C-bit field in the pseudowire FEC is set to the value configured on the static spoke-SDP.

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

It is possible for end nodes of a static pseudowire segment to be misconfigured. In that case, an S-PE or T-PE node may be receiving packets with the wrong encapsulation. Therefore, an invalid payload will be forwarded over the pseudowire or the SAP, respectively. Also, if the S-PE or T-PE node is expecting the control word in the packet encapsulation, and the received packet comes with no control word but the first nibble below the label stack is 0x0001, the packet may be mistaken for a VCCV OAM packet and may be forwarded to the CPM. In that case, the CPM will perform a check of the IP header fields, such as version, IP header length, and checksum. If any of this fails, the VCCV packet will be discarded.

3.2.3 Pseudowire redundancy

Pseudowire redundancy provides the ability to protect a pseudowire with a preprovisioned pseudowire and to switch traffic over to the secondary standby pseudowire in case of a SAP and network failure condition. Usually, pseudowires are redundant because of the SDP redundancy mechanism. For example, if the SDP is an RSVP LSP and is protected by a secondary standby path or by fast reroute (FRR) paths, the

pseudowire is also protected. However, there are a couple of applications in which SDP redundancy does not protect the end-to-end pseudowire path:


- There are two different destination PE nodes for the same VLL service. The main use case is the provision of dual-homing of a CPE or access node to two PE nodes located in different POPs. The other use case is the provision of a pair of active and standby BRAS nodes, or active and standby links to the same BRAS node, to provide service resiliency to broadband service subscribers.
- The pseudowire path is switched in the middle of the network and the 7210 SAS pseudowire switching node fails.

Pseudowire and VPLS link redundancy extends link-level resiliency for pseudowires and VPLS to protect critical network paths against physical link or node failures. These innovations enable the virtualization of redundant paths across the metro or core IP network to provide seamless and transparent fail-over for point-to-point and multi-point connections and services. When deployed with multi-chassis LAG, the path for return traffic is maintained through the pseudowire or VPLS switchover, which enables carriers to deliver "always on" services across their IP/MPLS networks.

3.2.3.1 VLL resilience with two destination PE nodes

The following figure shows the application of pseudowire redundancy to provide Ethernet VLL service resilience for broadband service subscribers accessing the broadband service on the service provider BRAS.

Figure 20: VLL resilience



OSSG115-7210M

If the Ethernet SAP on PE2 fails, PE2 notifies PE1 of the failure by either withdrawing the primary pseudowire label it advertised or by sending a pseudowire status notification with the code set to indicate a SAP defect. PE1 will receive it and will immediately switch its local SAP to forward over the secondary standby spoke-SDP. To avoid black holing of in-flight packets during the switching of the path, PE1 will accept packets received from PE2 on the primary pseudowire while transmitting over the backup pseudowire.

When the SAP on PE2 is restored, PE2 updates the new status of the SAP by sending a new label mapping message for the same pseudowire FEC or by sending a pseudowire status notification message indicating that the SAP is back up. PE1 then starts a timer and reverts back to the primary at the expiry

of the timer. By default, the timer is set to 0, which means PE1 reverts immediately. A special value of the timer (infinity) will mean that PE1 should never revert to the primary pseudowire.

The behavior of the pseudowire redundancy feature is the same if PE1 detects or is notified of a network failure that brought the spoke-SDP operational status to down. The following are the events that will cause PE1 to trigger a switchover to the secondary standby pseudowire:

1. T-LDP peer (remote PE) node withdrew the pseudowire label.
2. T-LDP peer signaled a FEC status indicating a pseudowire failure or a remote SAP failure.
3. T-LDP session to the peer node times out.
4. SDP binding and VLL service went down as a result of a network failure condition, such as the SDP to peer node going operationally down.

The 7210 SAS routers support the ability to configure multiple secondary standby pseudowire paths. For example, PE1 uses the value of the user-configurable precedence parameter associated with each spoke-SDP to select the next available pseudowire path after the failure of the current active pseudowire (whether it is the primary or one of the secondary pseudowires). However, the revertive operation always switches the path of the VLL back to the primary pseudowire. There is no revertive operation between secondary paths, meaning that the path of the VLL will not be switched back to a secondary pseudowire of higher precedence when the latter comes back up again.

-The 7210 SAS routers support the ability for a user-initiated manual switchover of the VLL path to the primary or any of the secondary be supported to divert user traffic in case of a planned outage, such as in node upgrade procedures.

3.2.4 Dynamic multi-segment pseudowire routing



Note:

T-PE functionality is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.


S-PE functionality is only supported on 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE.

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

Dynamic Multi-Segment Pseudowire (MS-PW) routing enables a complete MS-PW to be established, while only requiring per-pseudowire configuration on the T-PEs. No per-pseudowire configuration is required on the S-PEs. End-to-end signaling of the MS-PW is achieved using T-LDP, while multi-protocol BGP is used to advertise the T-PEs, allowing dynamic routing of the MS-PW through the intervening network of S-PEs. Dynamic MS-PWs are described in IETF draft-ietf-pwe3-dynamic-ms-pw-13.txt.


The following figure shows the operation of dynamic MS-PWs.

Figure 21: Dynamic MS-PW operation



The FEC 129 Attachment Individual Identifier (All) type 2 structure shown in the following figure is used to identify each individual pseudowire endpoint.


Figure 22: MS-PW addressing using FEC 129 all Type 2



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

This new All type is mapped into the MS-PW BGP NLRI (a new BGP AFI of L2VPN, and SAFI for network layer reachability information for dynamic MS-PWs. As soon as a new T-PE is configured with a local prefix address of global id:prefix, pseudowire routing will proceed to advertise this new address to all the other T-PEs and S-PEs in the network, as shown in the following figure.

Figure 23: Advertisement of PE addresses by PW routing




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

Next, in steps 2 to 5, MP-BGP will use the NLRI for the MS-PW routing SAFI to advertise the location of the new T-PE to all the other PEs in the network. Alternatively, static routes may be configured on a per T-PE/S-PE basis to accommodate non-BGP PEs in the solution.

As a result, pseudowire routing tables for all the S-PEs and remote T-PEs are populated with the next hop to be used to reach T-PE2.


VLL services can then be established, as shown in the following figure.

Figure 24: Signaling of dynamic MS-PWs using T-LDP



In step 1 and 1', the T-PEs are configured with the local and remote endpoint information, Source AII (SAIL), Target AII (TAIL). On the 7210 SAS, the AIs are locally configured for each spoke-SDP, according to the model shown in the following figure. Therefore, the 7210 SAS provides for a flexible mapping of AII to SAP. That is, the values used for the AII are through local configuration, and it is the context of the spoke-SDP that binds it to a specific SAP.

Figure 25: Mapping of All to SAP



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

This signaling hop is then used to choose the T-LDP session to the chosen next-hop S-PE. Signaling proceeds through each subsequent S-PE using similar matching procedures to determine the next signaling hop. Otherwise, if a subsequent S-PE does not support dynamic MS-PW routing, and therefore uses a statically configured PW segment, the signaling of individual segments follows the procedures already implemented in the PW switching feature.




Note:

BGP can install a PW All route in the PW routing table with ECMP next hops. However, when LDP needs to signal a PW with matching TAll, it will choose only one next hop from the available ECMP next hops. PW routing supports up to four ECMP paths for each destination.

The signaling of the forward path ends when the PE matches the TAll in the label mapping message with the SAll of a spoke-SDP bound to a local SAP. The signaling in the reverse direction can now be initiated, which follows the entries installed in the forward path. The PW routing tables are not consulted for the reverse path. This ensures that the reverse direction of the PW follows exactly the same set of S-PEs as the forward direction.

This solution can be used in either a MAN-WAN environment or in an inter-AS/inter-provider environment, as shown in the following figure.

Figure 26: VLL using dynamic MS-PWs, Inter-AS scenario

**Note:**

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

3.2.4.1 Pseudowire routing

Each S-PE and T-PE has a pseudowire routing table that contains a reference to the T-LDP session to use to signal to a set of next-hop S-PEs to reach a specific T-PE (or the T-PE if that is the next hop). For VLLs, this table contains aggregated All type 2 FECs and may be populated with routes that are learned through MP-BGP or that are statically configured.

MP-BGP is used to automatically distribute T-PE prefixes using the new MS-PW NLRI, or static routes can be used. The MS-PW NLRI is composed of a length, an eight-byte RD, a four-byte global ID, a four-byte local prefix, and (optionally) a four-byte AC ID. Support for the MS-PW address family is configured in CLI under **config>router>bgp>family ms-pw**.

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

To enable support for dynamic MS-PWs on a 7210 SAS node to be used as a T-PE or S-PE, a single, globally unique, S-PE ID, known as the S-PE address, is first configured under **config>service>pw-routing** on each 7210 SAS to be used as a T-PE or S-PE. The S-PE address has the format global-id:prefix. It is not possible to configure any local prefixes used for pseudowire routing or to configure spoke-SPDs using dynamic MS-PWs at a T-PE unless an S-PE address has already been configured. The S-PE address is used as the address of a node used to populate the switching point TLV in the LDP label mapping message and the pseudowire status notification sent for faults at the S-PE.

Each T-PE is also be configured with the following parameters:

- 1. Global ID**

This is a 4-byte identifier that uniquely identifies an operator or the local network.

- 2. Local prefix**

One or more local (Layer 2) prefixes (up to a maximum of 16), which are formatted in the style of a 4-octet IPv4 address. A local prefix identifies a T-PE or S-PE in the PW routing domain.

- 3. For each local prefix, at least one 8-byte route distinguisher can be configured. It is also possible to configure an optional BGP community attribute.**

For each local prefix, BGP then advertises each global ID/prefix tuple and unique RD and community pseudowire using the MS-PW NLRI, based on the aggregated FEC129 All Type 2 and the Layer 2 VPN/ PW routing AFI/SAFI 25/6, to each T-PE/S-PE that is a T-LDP neighbor, subject to local BGP policies.

The dynamic advertisement of each of these pseudowire routes is enabled for each prefix and RD, using the **advertise-bgp** command.

Example: Exporting MS-PW routes in MP-BGP

An export policy is also required to export MS-PW routes in MP-BGP. This can be done using a default policy, such as the following:

```
*A:lin-123>config>router>policy-options# info
-----
      policy-statement "ms-pw"
        default-action accept
        exit
      exit
-----
```

However, the preceding would export all routes. A recommended choice is to enable filtering per-family, as follows:

```
*A:lin-123>config>router>policy-options# info
-----
      policy-statement "to-mspw"
        entry 1
          from
            family ms-pw
          exit
          action accept
          exit
        exit
      exit
-----
```

The following command is then added in the **config>router>bgp** context:

```
export "to-mspw"
```

The local preference for iBGP and BGP communities can be configured under such a policy.

3.2.4.1.1 Static routing

In addition to support for BGP routing, static MS-PW routes may also be configured using the **config>services>pw-routing>static-route** command. Each static route comprises the target T-PE global ID and prefix, and the IP address of the T-LDP session to the next-hop S-PE or T-PE that should be used.

If a static route is set to 0, this represents the default route. If a static route exists to a specific T-PE, this is used in preference to any BGP route that may exist.

3.2.4.1.2 Explicit paths

A set of default explicit paths to a remote T-PE or S-PE prefix may be configured on a T-PE under **config>services>pw-routing** using the **path name** command. Explicit paths are used to populate the explicit route TLV used by MS-PW T-LDP signaling. Only strict (fully qualified) explicit paths are supported.

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

3.2.4.2 Configuring VLLs using dynamic MS-PW

One or more spoke-SDPs may be configured for distributed Epipe VLL services. Dynamic MS-PWs use FEC129 (also known as the Generalized ID FEC) with All type 2 to identify the pseudowire, as opposed to FEC128 (also known as the PW ID FEC) used for traditional single segment pseudowires and for pseudowire switching. FEC129 spoke-SDPs are configured under the **spoke-sdp-fec** command in the CLI.

FEC129 All type 2 uses a SAll and a TAll to identify the end of a pseudowire at the T-PE. The SAll identifies the local end, while the TAll identifies the remote end. The SAll and TAll are each structured as follows:

- **global ID**

The global ID is a 4-byte identifier that uniquely identifies an operator or the local network.

- **prefix**

This is 4-byte prefix, which should correspond to one of the local prefixes assigned under pw-routing.

- **AC ID**

The AC ID is a 4-byte identifier for this end of the pseudowire. This should be locally unique within the scope of the global-id:prefix.

3.2.4.2.1 Active/passive T-PE selection

Dynamic MS-PWs use single-sided signaling procedures with double-sided configuration; a fully qualified FEC must be configured at both endpoints. That is, one T-PE (the source T-PE: ST-PE) of the MS-PW initiates signaling for the MS-PW, while the other end (the terminating T-PE: TT-PE) passively waits for the label mapping message from the far-end. The TT-PE only responds with a label mapping message to set up the opposite direction of the MS-PW when it receives the label mapping from the ST-PE. By default, the 7210 SAS will determine which T-PE is the ST-PE (the active T-PE) and which is the TT-PE (the passive T-PE) automatically, based on comparing the SAll with the TAll as unsigned integers. The T-PE with SAll>TAll assumes the active role. However, it is possible to override this behavior using the **signaling {master | auto}** command in the **spoke-sdp-fec** context. If master is selected at a specific T-PE, it will assume the active role. If a T-PE is at the endpoint of a spoke-SDP that is bound to an VLL SAP, and single-sided autoconfiguration is used, that endpoint is always passive. Therefore, signaling master should only be used when it is known that the far end will assume a passive behavior.

3.2.4.2.2 Automatic endpoint configuration

Automatic endpoint configuration allows the configuration of an endpoint without specifying the TAll associated with that spoke-SDP FEC. It allows a single-sided provisioning model where an incoming label

mapping message with a TAIL that matches the SAIL of that spoke-SDP can be automatically bound to that endpoint. This is useful in scenarios where a service provider must separate the service configuration from the service activation phase.

Automatic endpoint configuration is supported and required for Epipe VLL spoke-SDP FEC endpoints bound to a VLL SAP. It is configured using the **spoke-sdp-fec auto-config** command, and excludes the TAIL from the configuration. When autoconfiguration is used, the node assumes passive behavior with regards to T-LDP signaling. Therefore, the far-end T-PE must be configured for signaling master for that spoke-SDP FEC.

3.2.4.2.3 Selecting a path for an MS-PW

Path selection for signaling occurs in the outbound direction (ST-PE to TT-PE) for an MS-PW. In the TT-PE to ST-PE direction, a label mapping message follows the reverse of the path of the outgoing label mapping.

A node can use explicit paths, static routes, or BGP routes to select the next hop S-PE or T-PE. The order of preference used in selecting these routes is:

1. explicit path
2. static route
3. BGP route

To use an explicit path for an MS-PW, an explicit path must have been configured in the **config>services>pw-routing>path** *path-name* context. The user must then configure the corresponding **path** *path-name* in the **spoke-sdp-fec** context.

If an explicit path name is not configured, the TT-PE or S-PE will perform a longest match lookup for a route (static if it exists, and BGP if not) to the next-hop S-PE or T-PE to reach the TAIL.

Pseudowire routing chooses the MS-PW path in terms of the sequence of S-PEs to use to reach a specific T-PE. It does not select the SDP to use on each hop, which is instead determined at signaling time. When a label mapping is sent for a specific pseudowire segment, an LDP SDP will be used to reach the next-hop S-PE/T-PE if such an SDP exists. If not, and an RFC 3107 labeled BGP SDP is available, then that will be used. Otherwise, the label mapping will fail and a label release will be sent.

3.2.4.2.4 Pseudowire templates


Dynamic MS-PWs support the use of the pseudowire template for specifying generic pseudowire parameters at the T-PE. The pseudowire template to use is configured in the **spoke-sdp-fec>pw-template-bind** *policy-id* context. Dynamic MS-PWs do not support the provisioned SDPs specified in the pseudowire template.

3.2.4.3 Pseudowire redundancy

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

The following figure shows the use of pseudowire redundancy.

Figure 27: Pseudowire redundancy



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

- Each MS-PW in the redundant set must have a unique SAIL/TAI set and is signaled separately. The primary pseudowire is configured in the **spoke-sdp-fec>primary** context.
- Each MS-PW in the redundant set should use a diverse path (from the point of view of the S-PEs traversed) from every other MS-PW in that set, if path diversity is possible in a specific network topology. There are a number of possible ways to achieve this:
 - Configure an explicit path for each MS-PW.
 - Allow BGP routing to automatically determine diverse paths using BGP policies applied to different local prefixes assigned to the primary and standby MS-PWs.
 - Provide path diversity for each primary pseudowire through the use of a BGP route distinguisher.

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



Note:

Because the SDP ID is determined dynamically at signaling time, it cannot be used as a tie breaker to choose the primary MS-PW between multiple MS-PWs of the same precedence. The user should therefore explicitly configure the precedence values to determine which MS-PW is active in the final selection.

3.2.4.4 VCCV OAM for dynamic MS-PWs

The primary difference between dynamic MS-PWs and those using FEC128 is support for FEC129 All type 2. As in PW switching, VCCV on dynamic MS-PWs requires the use of the VCCV control word on the pseudowire. Both the **vccv-ping** and **vccv-trace** commands support dynamic MS-PWs.

3.2.4.5 VCCV-Ping on dynamic MS-PWs

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

If the SAIL:TAIL is entered by the user in the **vccv-ping** command, those values are used for the VCCV-ping echo request. However, their order is reversed before being sent, so that they match the order for the downstream FEC element for an S-PE, or the locally configured SAIL:TAIL for a remote T-PE of that MS-PW.



Note:

If SAIL:TAIL is entered in addition to the spoke-sdp-fec-id, the system will verify the entered values against the values stored in the context for that spoke-sdp-fec-id.

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

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



Note:

The use of spoke-sdp-fec-id in vccv-ping is only applicable at T-PE nodes, because it is not configured for a specific MS-PW at S-PE nodes.

3.2.4.6 VCCV-Trace on dynamic MS-PWs

The 7210 SAS supports the MS-PW path trace mode of operation for VCCV-trace, as per pseudowire switching, but using FEC129 All type 2. As in the case of VCCV-ping, the SAIL:TAIL used in the VCCV-trace echo request message sent from the T-PE or S-PE, from which the **vccv-trace** command is executed, is specified by the user or derived from the context of the MS-PW.




Note:

The use of spoke-sdp-fec-id in **vccv-trace** is only applicable at T-PE nodes, since it is not configured for a specific MS-PW at S-PE nodes

3.2.5 Example dynamic MS-PW configuration

The following figure shows how to configure dynamic MS-PWs for a VLL service between a set of 7210 SAS nodes. The network consists of two 7210 T-PEs and two 7210 SAS nodes in the role of S-PEs. Each 7210 SAS peers with its neighbor using LDP and BGP.

Figure 28: Dynamic MS-PW example



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

Finally, pseudowire routing must be configured on each node. This includes an S-PE address for every participating node, and one or more local prefixes on the T-PEs. MS-PW paths and static routes may also be configured. When this routing and signaling infrastructure is established, spoke-SDP FECs can be configured on each of the T-PEs.

Example: Sample configuration for T-PE1

The following is a sample configuration for T-PE1.

```
config
router
  ldp
    targeted-session
      peer 10.20.1.5
    exit
  exit
  policy-options
    begin
      policy-statement "exportMsPw"
        entry 10
          from
            family ms-pw
          exit
          action accept
        exit
      exit
    exit
  exit
  commit
exit
bgp
  family ms-pw
  connect-retry 1
  min-route-advertisement 1
```

```

        export "exportMsPw"
        rapid-withdrawal
        group "ebgp"
            neighbor 10.20.1.5
                multihop 255
                peer-as 200
            exit
        exit
    exit
config
  service
    pw-routing
      spe-address 3:10.20.1.3
      local-prefix 3:10.20.1.3 create
      exit
      path "path1_to_F" create
          hop 1 10.20.1.5
          hop 2 10.20.1.2
          no shutdown
      exit
    exit
    epipe 1 customer 1 vpn 1 create
        description "Default epipe
            description for service id 1"
        service-mtu 1400
        service-name "XYZ Epipe 1"
        sap 2/1/1:1 create
        exit
        spoke-sdp-fec 1 fec 129 aii-type 2 create
            retry-timer 10
            retry-count 10
            saii-type2 3:10.20.1.3:1
            taii-type2 6:10.20.1.6:1
            no shutdown
        exit
    no shutdown
    exit

```

Example: Sample configuration for T-PE2

The following is a sample configuration for T-PE2.

```

config
  router
    ldp
      targeted-session
        peer 10.20.1.2
        exit
      exit
    ...
    policy-options
      begin
        policy-statement "exportMsPw"
          entry 10
            from
              family ms-pw
            exit
            action accept
            exit
          exit
        exit
      exit

```



```

        commit
    exit

    bgp
        family ms-pw
        connect-retry 1
        min-route-advertisement 1
        export "exportMsPw"
        rapid-withdrawal
        group "ebgp"
            neighbor 10.20.1.2
                multihop 255
                peer-as 300
        exit
    exit
exit

config
service
    pw-routing
        spe-address 6:10.20.1.6
        local-prefix 6:10.20.1.6 create
        exit
        path "path1_to_F" create
            hop 1 10.20.1.2
            hop 2 10.20.1.5
            no shutdown
        exit
    exit
    epipe 1 customer 1 vpn 1 create
        description "Default epipe
        description for service id 1"
service-mtu 1400
    service-name "XYZ Epipe 1"
    sap 1/1/3:1 create
    exit
    spoke-sdp-fec 1 fec 129 aii-type 2 create
        retry-timer 10
        retry-count 10
        saii-type2 6:10.20.1.6:1
        taii-type2 3:10.20.1.3:1
        no shutdown
    exit
    no shutdown
exit

```

Example: Sample configuration for S-PE1

The following is a sample configuration for S-PE1.

```

config
router
    ldp
        targeted-session
            peer 10.20.1.3
            exit
            peer 10.20.1.2
            exit
        exit
    ...
    bgp
        family ms-pw

```

```

        connect-retry 1
        min-route-advertisement 1
        rapid-withdrawal
        group "ebgp"
            neighbor 10.20.1.2
                multihop 255
                peer-as 300
            exit
            neighbor 10.20.1.3
                multihop 255
                peer-as 100
            exit
        exit
    exit
service
    pw-routing
        spe-address 5:10.20.1.5
    exit

```

Example: Sample configuration for S-PE2

The following is a sample configuration for S-PE2.

```

config
router
    ldp
        targeted-session
            peer 10.20.1.5
            exit
            peer 10.20.1.6
            exit
        exit
    ...
    bgp
        family ms-pw
        connect-retry 1
        min-route-advertisement 1
        rapid-withdrawal
        group "ebgp"
            neighbor 10.20.1.5
                multihop 255
                peer-as 200
            exit
            neighbor 10.20.1.6
                multihop 255
                peer-as 400
            exit
        exit
    exit
service
    pw-routing
        spe-address 2:10.20.1.2
    exit

```

3.3 Master-slave operation



Note:


The **standby-signaling-master** command is supported on all 7210 SAS platforms as described in this document, except for those operating in access-uplink mode. The **standby-signaling-slave** command is not supported. In the following section, references to the **standby-signaling-slave** command are only used to describe the complete solution. The 7210 SAS platforms can only be used where the **standby-signaling-master** command is used in the example.

This section describes a mechanism in which one end on a pseudowire (the "master") dictates the active PW selection, which is followed by the other end of the PW (the "slave"). This mechanism and associated terminology is specified in RFC 6870.

This section describes master-slave pseudowire redundancy. This redundancy adds the ability for the remote peer to react to the pseudowire standby status notification, even if only one spoke-SDP terminates on the VLL endpoint on the remote peer, by blocking the transmit (Tx) direction of a VLL spoke-SDP when the far-end PE signals standby. This solution enables the blocking of the Tx direction of a VLL spoke-SDP at both master and slave endpoints when standby is signalled by the master endpoint. This satisfies a majority of deployments where bidirectional blocking of the forwarding on a standby spoke-SDP is required.

The following figure shows the operation of master-slave pseudowire redundancy. In this scenario, an Epipe service is provided between CE1 and CE2. CE2 is dual-homed to PE2 and PE3, and therefore PE1 is dual-homed to PE2 and PE3 using Epipe spoke-SDPs. The objective of this feature is to ensure that only one pseudowire is used for forwarding in both directions by PE1, PE2, and PE3, in the absence of a native dual homing protocol between CE2 and PE2/PE3, such as MC-LAG. In normal operating conditions (the SAPs on PE2 and PE3 toward CE2 are both up and there are no defects on the ACs to CE2), PE2 and PE3 cannot choose which spoke-SDP to forward on based on the status of the AC redundancy protocol.

Figure 29: Master-slave pseudowire redundancy



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

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

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

That the specific spoke-SDP has been blocked will be signaled to the LDP peer through the pseudowire status bit (PW FWD standby (0x20)). This will prevent traffic being sent over this spoke-SDP by the remote peer, but only in case that remote peer supports and reacts to pseudowire status notification. Previously, this applied only if the spoke-SDP terminated on an IES, VPRN, or VPLS.

Note that although master-slave operation provides bidirectional blocking of a standby spoke-SDP during steady-state conditions, it is possible that the Tx directions of more than one slave endpoint can be active for transient periods during a fail-over operation. This is because of slave endpoints transitioning a spoke-SDP from standby to active receiving and processing a pseudowire preferential forwarding status message before those transitioning a spoke-SDP to standby.

This transient condition is most likely when a forced switch-over is performed, or the relative preferences of the spoke-SDPs are changed, or the active spoke-SDP is shutdown at the master endpoint. During this period, loops of unknown traffic may be observed. Fail-overs because of common network faults that can occur during normal operation, or a failure of connectivity on the path of the spoke-SDP or the SAP, would not result in such loops in the datapath.


3.3.1 Operation of master-slave pseudowire redundancy with existing scenarios

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

3.3.1.1 VLL resilience path

The following figure shows a VLL resilience path example. An sample configuration follows.

Figure 30: VLL resilience

**Note:**

A **revert-time** value of zero (default) means that the VLL path will be switched back to the primary immediately after it comes back up.

Example: Sample configuration for PE1

The following is a sample configuration for PE1.

```
configure service epipe 1
  endpoint X
  exit
  endpoint Y
  revert-time 0
  standby-signaling-master
  exit
  sap 1/1/1:100 endpoint X
  spoke-sdp 1:100 endpoint Y
precedence primary
  spoke-sdp 2:200 endpoint Y
precedence 1
```

Example: Sample configuration for PE2

The following is a sample configuration for PE2.

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

Example: Sample configuration for PE3

The following is a sample configuration for PE3.

```
configure service epipe 1
  endpoint X
```

```

exit
sap 3/3/3:300 endpoint X
spoke-sdp 2:200
standby-signaling-slave

```

3.3.2 VLL resilience for a switched PW path




Note:

T-PE functionality is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

S-PE functionality is only supported on 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE.

The following figure shows VLL resilience for a switched pseudowire path example. A sample configuration follows.

Figure 31: VLL resilience with pseudowire switching



sw0505

Example: Sample configuration of T-PE1

The following is a sample configuration of T-PE1.

```

configure service epipe 1
  endpoint X
  exit
  endpoint Y
  revert-time 100
  standby-signaling-master
  exit
  sap 1/1/1:100 endpoint X
  spoke-sdp 1:100 endpoint Y
    precedence primary
  spoke-sdp 2:200 endpoint Y
    precedence 1
  spoke-sdp 3:300 endpoint Y
    precedence 1

```

Example: Sample configuration of T-PE2

The following is a sample configuration of T-PE2.

```
configure service epipe 1
  endpoint X
  exit
  endpoint Y
  revert-time 100
  standby-signaling-slave
  exit
  sap 2/2/2:200 endpoint X
  spoke-sdp 4:400 endpoint Y
    precedence primary
  spoke-sdp 5:500 endpoint Y
    precedence 1
  spoke-sdp 6:600 endpoint Y
    precedence 1
```

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

Example: Sample configuration of S-PE1


The following is a sample configuration of S-PE1.

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

3.3.3 Access mode resilience Using MC-LAG and pseudowire redundancy

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

Figure 32: Access node resilience using MC-LAG and PW redundancy



OSSG116

In this application, a pseudowire status bit of active or standby indicates the status of the SAP in the MC-LAG instance in the 7210 SAS aggregation node. All spoke-SDPs are of secondary type and there is no use of a primary pseudowire type in this mode of operation.

Node A is in the active state according to its local MC-LAG instance, and therefore advertises active status notification messages to both its peer pseudowire nodes; for example, nodes C and D. Node D performs the same operation. Node B is in the standby state according to the status of the SAP in its local MC-LAG instance, and therefore advertises standby status notification messages to both nodes C and D. Node C performs the same operation.

The 7210 SAS node selects a pseudowire as the active path for forwarding packets when both the local pseudowire status and the received remote pseudowire status indicate active status. However, a 7210 SAS device in standby status according to the SAP in its local MC-LAG instance is capable of processing packets for a VLL service received over any of the pseudowires that are up. This is to avoid black holing of user traffic during transitions.

The 7210 SAS standby node forwards these packets to the active node via the Inter-Chassis Backup (ICB) pseudowire for this VLL service. An ICB is a spoke-SDP used by an MC-LAG node to back up an MC-LAG SAP during transitions. The same ICB can also be used by the peer MC-LAG node to protect against network failures causing the active pseudowire to go down.


Note that at configuration time, the user specifies a precedence parameter for each of the pseudowires that are part of the redundancy set as described in the application. A 7210 SAS node uses this to select which pseudowire to forward packet to in case both pseudowires show active/active for the local/remote status during transitions.

Only VLL service of type Epipe is supported in this application. Also, ICB spoke-SDP can only be added to the SAP side of the VLL cross-connect if the SAP is configured on a MC-LAG instance.

3.3.4 VLL resilience for a switched pseudowire path

The following figure shows the use of both pseudowire redundancy and pseudowire switching to provide a resilient VLL service across multiple IGP areas in a provider network.

Figure 33: VLL resilience with PW redundancy and PW switching



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

As in the application in [VLL resilience for a switched pseudowire path](#), the T-PE1 node switches the path of a VLL to a secondary standby pseudowire, in the case of a network-side failure causing the VLL binding status to be down or if T-PE2 notified it that the remote SAP went down. This application requires that

pseudowire status notification messages generated by either a T-PE node or a S-PE node be processed and relayed by the S-PE nodes.

**Note:**

It is possible that the secondary pseudowire path terminates on the same target PE as the primary; for example, T-PE2. This provides protection against network-side failures, but not against a remote SAP failure.

When the target destination PE for the primary and secondary pseudowires is the same, T-PE1 will usually not switch the VLL path onto the secondary pseudowire upon receipt of a pseudowire status notification indicating that the remote SAP is down. This occurs because the status notification is sent over both the primary and secondary pseudowires. However, the status notification on the primary pseudowire may arrive earlier than the one on the secondary pseudowire because of the differential delay between the paths. This will cause T-PE1 to switch the path of the VLL to the secondary standby pseudowire and remain there until the status notification is cleared. At that time, the VLL path is switched back to the primary pseudowire because of the revertive behavior operation. The path will not switch back to a secondary path when it comes up, even if it has a higher precedence than the currently active secondary path.


3.4 Pseudowire redundancy service models

This section describes the various MC-LAG and pseudowire redundancy scenarios, and the algorithm used to select the active transmit object in a VLL endpoint.

3.4.1 Redundant VLL service model

To implement pseudowire redundancy, a VLL service accommodates more than a single object on the SAP side and on the spoke-SDP side. The following figure shows the model for a redundant VLL service based on the concept of endpoints.

Figure 34: Redundant VLL endpoint objects



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

To add more objects, up to two explicitly named endpoints may be created per VLL service. The endpoint name is locally significant to the VLL service. They are referred to as endpoint X and endpoint Y as shown in the preceding figure.

The preceding figure is only an example and the Y endpoint can also have a SAP and/or an ICB spoke-SDP. The following describes the four types of endpoint objects supported and the rules used when associating them with an endpoint of a VLL service:

- **SAP**

There can only be a maximum of one SAP per VLL endpoint.

- **primary spoke-SDP**

The VLL service always uses this pseudowire and only switches to a secondary pseudowire when it is down. The VLL service switches the path to the primary pseudowire when it is back up. The user can configure a timer to delay reverting back to primary or to never revert. There can only be a maximum of one primary spoke-SDP per VLL endpoint.

- **secondary spoke-SDP**

There can be a maximum of four secondary spoke-SDPs per endpoint. The user can configure the precedence of a secondary pseudowire to indicate the order in which a secondary pseudowire is activated.

- **inter-chassis backup (ICB) spoke-SDP**

This is a special pseudowire used for MC-LAG and pseudowire redundancy applications. Forwarding between ICBs is blocked on the same node. The user has to explicitly indicate that the spoke-SDP is an ICB at creation time. There are a few scenarios, as follows, where the user can configure the spoke-SDP as an ICB or as a regular spoke-SDP on a specific node. The CLI for those cases will indicate both options.

A VLL service endpoint can only use one active object to transmit at a time, but can receive from all endpoint objects.

An explicitly named endpoint can have a maximum of one SAP and one ICB. When a SAP is added to the endpoint, only one more object of type ICB spoke-SDP is allowed. The ICB spoke-SDP cannot be added to the endpoint if the SAP is not part of a MC-LAG instance. A SAP that is not part of a MC-LAG instance cannot be added to an endpoint that already has an ICB spoke-SDP.

An explicitly named endpoint, which does not have a SAP object, can have a maximum of four spoke-SDPs and can include any of the following:

- a single primary spoke-SDP
- one or many secondary spoke-SDPs with precedence
- a single ICB spoke-SDP

3.4.2 T-LDP status notification handling rules

Using [Figure 34: Redundant VLL endpoint objects](#) as a reference, the following are the rules for generating, processing, and merging T-LDP status notifications in VLL service with endpoints. Note that any allowed combination of objects as specified in [Redundant VLL service model](#) can be used on endpoints "X" and "Y". The following sections refer to the specific combination objects in [Figure 34: Redundant VLL endpoint objects](#) as an example to describe the more general rules.

3.4.2.1 Processing endpoint SAP active/standby status bits

The advertised admin forwarding status of active/standby reflects the status of the local LAG SAP in the MC-LAG application. If the SAP is not part of a MC-LAG instance, the forwarding status of active is always advertised.

When the SAP in endpoint "X" is part of a MC-LAG instance, a node must send the T-LDP forwarding status bit of "SAP active/standby" over all "Y" endpoint spoke-SDPs, except the ICB spoke-SDP, whenever this status changes. The status bit sent over the ICB is always zero (active by default).

When the SAP in endpoint "X" is not part of a MC-LAG instance, the forwarding status sent over all "Y" endpoint spoke-SDPs should always be set to zero (active by default).

3.4.2.2 Processing and merging

Endpoint X is operationally up if at least one of its objects is operationally up. It is down if all its objects are operationally down.

If the SAP in endpoint X transitions locally to the down state, or receives a SAP down notification by SAP-specific OAM signal, the node must send T-LDP SAP down status bits on the Y endpoint ICB spoke-SDP only.



Note:

The Ethernet SAP does not support the SAP OAM protocol. All other SAP types cannot exist on the same endpoint as an ICB spoke-SDP because a non-Ethernet SAP cannot be part of an MC-LAG instance.

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

If the ICB spoke-SDP in endpoint X receives T-LDP SDP-binding down status bits or pseudowire not forwarding status bits, the node saves this status and takes no further action. The saved status is used for selecting the active transmit endpoint object.

If all objects in endpoint X have any of the following happen, the node must send status bits of SAP down over all Y endpoint spoke-SDPs, including the ICB:

- transition locally to down state
- receive a SAP down notification by remote T-LDP status bits or by SAP-specific OAM signal
- receive status bits of SDP-binding down
- receive status bits of pseudowire not forwarding

Endpoint Y is operationally up if at least one of its objects is operationally up. It is down if all its objects are operationally down.

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

If a spoke-SDP in endpoint Y, including the ICB spoke-SDP, receive any of the following, the node saves this status and takes no further action:

- T-LDP SAP down status bits
- T-LDP SDP-binding down status bits
- status bits of pseudowire not forwarding

The saved status is used for selecting the active transmit endpoint object.

If all objects in endpoint Y, except the ICB spoke-SDP, have any of the following happen, the node must send status bits of SDP-binding down over the X endpoint ICB spoke-SDP only:

- transition locally to down state
- receive T-LDP SAP down status bits
- receive T-LDP SDP-binding down status bits
- receive status bits of pseudowire not forwarding

If all objects in endpoint Y have any of the following happen, the node must send status bits of the SDP-binding down over the X endpoint ICB spoke-SDP, and must send a SAP down notification on the X endpoint SAP by the SAP-specific OAM signal, if applicable:

- transition locally to down state
- receive T-LDP SAP down status bits
- receive T-LDP SDP-binding down status bits
- receive status bits of pseudowire not forwarding

An Ethernet SAP does not support signaling status notifications.

3.5 Epipe configuration for MPLS-TP



Note:

MPLS-TP PWs are supported in Epipe services.

MPLS-TP is only supported on 7210 SAS-T operating in the network mode.

The following subsections describe how SDPs and spoke-SDPs are used with MPLS-TP LSPs and static PWs with MPLS-TP OAM.

3.5.1 SDPs

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

The following syntax shows the MPLS-TP options:

```
config
service
    sdp
        no description
        network-domain "default"
        signaling off
        far-end node-id 0.0.0.43 global-id 4294967295
        no mixed-lsp-mode
        no ldp
        no bgp-tunnel
        lsp "unnumberedLSP"
        no vlan-vc-etype
        no pbb-etype
        no path-mtu
        no adv-mtu-override
        keep-alive
            shutdown
            hello-time 10
            hold-down-time 10
            max-drop-count 3
            timeout 5
            no message-length
        exit
        no metric
        no collect-stats
        no accounting-policy
        binding
            no port
        exit
        no shutdown
    -----
    *A:7210SAS>config>service>sdp#
```

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

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

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

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

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

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

Signaling TLDP or BGP is blocked if:

- **far-end node-id** or **global-id** is configured
- **control-channel-status** is enabled on any spoke-SDP (or mate VC-switched spoke)
- **pw-path-id** is configured on any spoke-SDP (or mate VC-switched spoke)

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

- **class-forwarding**
- **tunnel-far-end**
- **mixed-lsp-mode**
- **keep-alive**
- LDP or BGP-tunnel
- **adv-mtu-override**

3.5.2 VLL spoke-SDP configuration

The 7210 SAS-T can only be a T-PE. MPLS-TP OAM related commands are applicable to spoke-SDPs configured under all services supported by MPLS-TP pseudowires. All commands and functions that are applicable to spoke-SDPs in the current implementation are supported, except for those that explicitly depend on an LDP session on the SDP or as described in this section. Likewise, all existing functions on a specific service SAP are supported if the spoke-SDPs that it are matched to is MPLS-TP.

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

A spoke-SDP bound to an SDP with the mpls-tp keyword cannot be **no shutdown** unless the ingress label, the egress label, the control word, and the PW path ID are configured, as follows:

```
*7210SAS>config>service>epipe# info
-----
      sap 1/1/10:1.111 create
      exit
      spoke-sdp 1:111 create
```

```

[no] hash-label ingress
      vc-label 2111
exit
egress
      vc-label 2111
exit
control-word
pw-path-id
      agi 0:111
      saii-type2 4294967295:0.0.0.42:111
      taii-type2 4294967295:0.0.0.43:111
exit
no shutdown
exit
no shutdown
-----
*7210SAS>config>service>epipe#

```

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

The **control-channel-status** command enables static PW status signaling. This is valid for any spoke-SDP where **signaling none** is configured on the SDP (for example, where T-LDP signaling is not in use). The refresh timer is specified in seconds, from 10 to 65535, with a default of 0 (off). This value can only be changed if **control-channel-status** is **shutdown**.

Commands that rely on PW status signaling are allowed if **control-channel-status** is configured for a spoke-SDP bound to an SDP with signaling off, but the system will use control channel status signaling rather than T-LDP status signaling. The ability to configure control channel status signaling on a specific spoke-SDP is determined by the credit-based algorithm described in [Credit-based algorithm](#). The **control-channel-status** configuration for a particular PW only counts against the credit-based algorithm if it is in a **no shutdown** state and has a non-zero refresh timer.



Note:

Shutdown of a service will cause the static PW status bits for the corresponding PW to be set.

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

The system accepts the **node-id** of the **pw-path-id** SAI or TAI being entered in either 4-octet IP address format <a.b.c.d> or unsigned integer format.

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

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

- network mode D
- SDP signaling is off
- **control-word** is enabled (**control-word** is disabled by default)
- service type Epipe or VPLS
- mate SDP signaling is off for VC-switched services
- an MPLS-TP **node-id** or **global-id** is configured under the **config>router>mpls>mpls-tp** context; this is required for OAM to provide a reply address

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

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

- network mode D
- SDP signaling is off
- **control-word** is enabled (**control-word** is disabled by default)
- the service type is Epipe or VPLS interface
- mate SDP signaling is off (in VC-switched services)
- **pw-status-signaling** is enabled
- **pw-path-id** is configured for this spoke

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

The control channel status request mechanism is enabled when the **request-timer** *timer* parameter is non-zero. When enabled, this overrides the normal RFC-compliant refresh timer behavior. The refresh timer value in the status packet defined in RFC 6478 is always set to zero.

The refresh timer in the sending node is taken from the request timer. The two mechanisms are not compatible with each other. One node sends a request timer while the other is configured for a refresh timer. In a specific node, the request timer can only be configured with both acknowledgment and refresh timers disabled.

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

Use the following syntax to configure control channel status requests:

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

- This parameter determines the interval at which PW status messages, including a reliable delivery TLV, with the request bit set are sent. This cannot be enabled if **refresh-timer** is not equal to 0. **retry-timer**: 3-60s.
- This parameter determines the timeout interval if no response to a PW status is received. This defaults to 0 when no retry-timer. timeout-multiplier <value> - 3-15.
- If a requesting node does not hear back after the **retry-timer** value times the multiplier value, the node must assume that the peer is down. This defaults to zero (0) when no retry-timer.

3.5.3 Credit-based algorithm

To constrain the CPU resources consumed processing control channel status messages, the system should implement a credit-based mechanism. If a user enables control channel status on a PW[n], a specific number of credits (c_n) are consumed from a CPM-wide pool of max_credit credits. The number of credits consumed is inversely proportional to the configured refresh timer (the first three messages at 1 s intervals do not count against the credit). If the current_credit ≤ 0, the control channel status signaling cannot be configured on a PW (but the PW can still be configured, and **no shutdown**).

The following is an example algorithm:

If refresh timer > 0, c_n = 65535 / refresh_timer

Else $c_n = 0$.

For $n=1$, $current\_credit[n] = max\_credits - c_n$

Else $current\_credit[n] = current\_credit[n-1] - c_n$

By default, if a PE with a non-zero refresh timer configured does not receive control channel status refresh messages for 3.5 times the specified timer value, the PE will time out and assume a PW status of zero. A proprietary optional extension to the [RFC6478] protocol should be implemented to enable a node to resolve such a stale PW status condition by requesting the status from the far-end node in specific cases.

3.6 VLAN range for SAPs in an Epipe service

The 7210 SAS VLAN ranges provide a mechanism to group a range of VLAN IDs as a single service entity. This allows the operator to provide the service treatment (forwarding, ACL, QoS, accounting, and others) to the group of VLAN IDs as a whole.



Note:

Grouping a range of VLAN IDs to a SAP is supported only for Virtual Leased Line (VLL) Ethernet services.

3.6.1 Processing behavior for SAPs using VLAN ranges in access-uplink mode

The access SAPs that specify VLAN range values using connection profile (also known as dot1q range SAPs) are allowed in Epipe service and in VPLS service. For more information about functionality supported, see [VLAN range SAPs feature support and restrictions](#). The system allows one range SAP in an Epipe service, and fails any attempt to configure more than one. Range SAPs can only be configured on access ports. The other endpoint in the Epipe service has to be a Q.* SAP in access-uplink mode. The processing and forwarding behavior for packets received on range SAPs are listed as follows:

- No VLAN tags are removed/stripped on ingress of an access dot1q SAP configured to use VLAN ranges. A single tag (Q1) is added to the frame when it is forwarded out of the Q1.* access-uplink SAP.
- When a packet is received on the access-uplink Q1.* SAP, the outermost tag is removed and the packet is forwarded out of the access dot1q range SAP. The system does not check if the inner VLAN tag matches the VLANs IDs (both range and individual values specified in the connection profile) of the dot1q access SAPs configured in the service.
- The dot1q range SAP can be supported in a service with **svc-sap-type** set to **dot1q-range**.

3.6.2 VLAN range SAPs feature support and restrictions

The following restrictions apply to VLAN range SAPs:

- The access SAPs that specify VLAN range values (using the connection profile) are only allowed in an Epipe service. The system allows only one range SAP in an Epipe service, and fails any attempt to configure more than one. Range SAPs can only be configured on access ports.
- In access-uplink mode, the dot1q range SAP is only allowed to be configured in a service with **svc-sap-type** set to **dot1q-range**. In network mode, the dot1q range SAP is allowed to be configured in a service with **svc-sap-type** set to **any**.

- The access SAPs using VLAN range values are only allowed for dot1q encapsulation ports or LAGs. A connection profile is used to specify either range of VLAN IDs or individual VLANs to be grouped together in a single SAP.
- A connection profile is used to specify either range of VLAN IDs or individual VLANs to be grouped together in a single SAP.
- Multiple connection profiles can be used per port or LAG, as long as the VLAN value specified by each of them does not overlap. The number of VLAN ranges available per port/LAG is limited. The available number must be shared among all the SAPs on the port/LAG.
- The connection profile associated with a SAP cannot be modified. To modify a connection profile, it must be removed from all SAPs that are using it.

3.6.3 Processing behavior for SAPs using VLAN ranges in network operating mode

The access SAPs that specify VLAN range values (using the connection profile) are only allowed in an Epipe service. The system allows only one range SAP in an Epipe service, and fails any attempt to configure more than one. Range SAPs can only be configured on access ports. The other endpoint in the Epipe service has to be a Q.* access SAP or a spoke-SDP (PW) in network mode.

The spoke-SDP processing and forwarding behavior for packets received on range SAPs are as follows:

- No VLAN tags are removed/stripped on ingress of the access dot1q SAPs using VLAN range connection profile.
- When the other endpoint in the service is configured to be an Q1.* access SAP, 7210 adds another tag to the packet and forwards it out of that SAP.
- If the other endpoint in the service is configured to be a spoke-SDP whose **vc-type** is set to **vc-ether**, the 7210 SAS adds the appropriate MPLS PW and LSP encapsulations and forwards it out of the SDP.
- In the reverse direction, when the other endpoint is a Q1.* SAP and a packet is received on it, the 7210 SAS removes the outermost VLAN tag and forwards the packet out of the access dot1q SAP using VLAN ranges.
- When the other endpoint is a spoke-SDP (whose **vc-type** is set to **vc-ether**), the 7210 SAS removes the MPLS PW and LSP encapsulation and forwards the packet out of the access dot1q SAP using VLAN ranges.
- The system does not check if the VLAN in the packet matches the VLAN IDs of the dot1q access SAPs configured in the service.

ACLs, QoS, hardware, accounting, and mirroring are supported as follows:

- For ACLs, filter policies are supported on SAP ingress. In 7210 SAS-T operating in access-uplink mode, IP criteria and MAC criteria-based filter policy is supported with access SAPs.

For more information about ACLs on range SAPs, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.

- In 7210 SAS devices operating in network mode, only MAC criteria-based filter policy is supported with access SAPs.

For more information about ACLs on range SAPs, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.

- For QoS, ingress classification and hierarchical metering are supported for SAP ingress. On the 7210 SAS-T, egress per port queues and shaping are supported. On the 7210 SAS-Mxp, egress per SAP queues and shaping are supported:

- SAP ingress classification criteria that is available for use with VLAN range SAPs is similar to that available for other SAPs supported in an Epipe service. Dot1p-based ingress classification uses the dot1p bits in the outermost VLAN tag for matching. On access egress, dot1p received from the SDP (on a network port) from another access port is preserved.
- The amount of hardware resources (such as CAM entries used for matching in QoS classification and ACL matching, meters used in SAP ingress policy, and others), consumed by a single range SAP, are equivalent to the amount of resources consumed by a single SAP that specifies a single VLAN ID for service identification. That is, the hardware has the ability to match a range of VLAN values, and therefore uses X resources for a SAP using a VLAN range, instead of $X \times n$, where n is the number of VLANs specified in the range and X is the amount of QoS or ACL resources needed.
- Ingress accounting support is similar to the support available for other SAPs in an Epipe service. The count of packets or octets received from individual VLANs configured in the connection profile is not available. No support for Egress SAP statistics and accounting is available.
- Mirroring is supported. In network mode, the use of service resiliency mechanisms, such as MC-LAG and Epipe PW redundancy, is supported.

3.7 VLL service considerations

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

3.7.1 SDPs

The most basic SDPs must have the following:

- locally unique SDP identification (ID) number
- system IP address of the originating and far-end routers
- SDP encapsulation type, MPLS

3.7.2 SAP encapsulations

The Epipe service is designed to carry Ethernet frame payloads, so it can provide connectivity between any two SAPs that pass Ethernet frames. The following SAP encapsulations are supported on the Epipe service:

- Ethernet null
- Ethernet dot1q
- QinQ

While different encapsulation types can be used, encapsulation mismatch can occur if the encapsulation behavior is not understood by connecting devices and they are unable to send and receive the expected traffic. For example, if the encapsulation type on one side of the Epipe is dot1q and the other is null, tagged traffic received on the null SAP will potentially be double tagged when it is transmitted out of the dot1q SAP.

3.7.3 QoS policies

The following QoS policies can be applied to an Epipe service:

- **traffic management**

Traffic management of Ethernet VLLs is achieved through the application of ingress QoS policies to SAPs and access egress QoS policies applied to the port. All traffic management is forwarding-class aware, and the SAP ingress QoS policy identifies the forwarding class based on the rules configured to isolate and match the traffic ingressing on the SAP. Forwarding classes are determined based on the Layer 2 (dot1p, MAC) or Layer 3 (IP, DSCP) fields of contained packets, and this association of forwarding class at ingress will determine both the queuing and the Dot1P bit setting of packets on the Ethernet VLL on the egress.

- **SAP ingress classification and policing**

The traffic at the SAP ingress is classified and metered according to the SLA parameters. All the traffic ingressing on the SAP is classified to a particular forwarding class. All the forwarding class is metered through and marked in-profile or out-of-profile, based on the meter parameters.

When applied to 7210 SAS Epipe services, service ingress QoS policies only create the unicast queues defined in the policy. The multipoint queues are not created on the service. Both Layer 2 or Layer 3 criteria can be used in the QoS policies for traffic classification in a service.

- **egress network dot1p marking**

Marking of IEEE dot1p bits in VLAN tags is as per the FC-to-dot1p map. See the default network QoS policy in the *7210 SAS-Mxp, R6, R12, S, Sx, T Quality of Service Guide* for more information. This marking is applied at the port level on access ports and access-uplink ports.

- **Ingress network classification**

Ingress network classification is based on the dot1p bits in the outer VLAN tag received on the access-uplink port. Dot1p-to-FC mapping is based on the network ingress QoS policy.

3.7.4 Filter policies

The 7210 SAS Epipe services can have a single filter policy associated on both ingress and egress. Both MAC and IP filter policies can be used on Epipe services.

3.7.5 MAC resources

Epipe services are point-to-point Layer 2 VPNs capable of carrying any Ethernet payloads. Although an Epipe is a Layer 2 service, the 7210 SAS Epipe implementation does not perform any MAC learning on the service, so Epipe services do not consume any MAC hardware resources.

3.8 Configuring a VLL service with CLI

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

3.8.1 Common configuration tasks

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

- Associate the service with a customer ID.
- Define SAP parameters:
 - (Optional) Select ingress QoS policies (configured in the **config>qos** context).
 - (Optional) Select accounting policy (configured in the **config>log** context).
- Define spoke-SDP parameters.



Note:

Spoke-SDP parameters are only supported on 7210 SAS platforms operating in the network mode.

- Enable the service.

3.8.2 Configuring VLL components

This section provides VLL configuration examples for the VLL services.

3.8.2.1 Creating an Epipe service in network mode

Use the following syntax to create an Epipe service.

```
config>service# epipe service-id [customer customer-id][create][vpn vpn-id]
description description-string
no shutdown
```

Example

The following is a sample Epipe configuration output:

```
A:ALA-1>config>service# info
-----
...
      1101 customer 1 vpn 1101 create
      description "Default epipe description for service id 1101"
      no shutdown
    exit
-----
A:ALA-1>config>service#
```

3.8.2.2 Creating an Epipe service in access-uplink mode



Note:

This section only applies to 7210 SAS-T operating in access-uplink mode.

Use the following syntax to create an Epipe service.

```
config>service# epipe service-id [customer customer-id] [create] [svc-sap-type {null-star
| dot1q | dot1q-preserve | any}] [customer- vid vlan-id] description description-string no
shutdown
```

```
A:ALA-1>config>service# info
-----
...
    epipe 500 customer 1 svc-sap-type null-star create
        description "Local Epipe Service with NULL SVC_SAP_TYPE"
        no shutdown
    exit
-----
A:ALA-1>config>service#
```

3.8.2.2.1 Configuring Epipe SAP parameters

A default QoS policy is applied to each ingress SAP. Additional QoS policies can be configured in the **config>qos** context. Filter policies are configured in the **config>filter** context and explicitly applied to a SAP. There are no default filter policies.

3.8.2.2.1.1 Epipe SAP

```
config>service# epipe service-id [customer customer-id]
    sap sap-id
    accounting-policy policy-id
    collect-stats
    description description-string
    no shutdown
    egress
        filter {ip ip-filter-name | mac mac-filter-name}
    ingress
        filter {ip ip-filter-name | mac mac-filter-name}
        qos policy-id
```

3.8.2.2.1.2 Local Epipe SAPs

To configure a basic local Epipe service, enter the **sap sap-id** command twice with different port IDs in the same service configuration.

By default, QoS policy ID 1 is applied to ingress service SAPs. Existing filter policies or other existing QoS policies can be associated with service SAPs on ingress and egress ports.

Ingress and egress SAP parameters can be applied to local and distributed Epipe service SAPs.

Example

The following is a sample of SAP configurations for local Epipe service 500 on SAP 1/1/2 and SAP 1/1/3 on ALA-1:

```
A:ALA-1>config>service# epipe 500 customer 5 create
config>service>epipe$ description "Local epipe service
config>service>epipe# sap 1/1/2 create
```

```

config>service>epipe>sap? ingress
config>service>epipe>sap>ingress# qos 20
config>service>epipe>sap>ingress# filter ip 1
config>service>epipe>sap>ingress# exit
config>service>epipe>sap# egress
config>service>epipe>sap 1/1/3 create
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# qos 555
config>service>epipe>sap>ingress# filter ip 1
config>service>epipe>sap>ingress# exit
config>service>epipe>sap# no shutdown
config>service>epipe>sap# exit

A:ALA-1>config>service# info
-----
...
    epipe 500 customer 5 create
        description "Local epipe service"
        sap 1/1/2 create
            ingress
                qos 20
                filter ip 1
            exit
        exit
        sap 1/1/3 create
            ingress
                qos 555
                filter ip 1
            exit
        exit
        no shutdown
    exit
-----
A:ALA-1>config>service#

```

3.8.2.2.1.3 Distributed Epipe service



Note:

SDPs are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

To configure a distributed Epipe service, you must configure service entities on the originating and far-end nodes. You must use the same service ID on both ends (for example, Epipe 5500 on ALA-1 and Epipe 5500 on ALA-2). The **spoke-sdp** *sdp-id:vc-id* must match on both sides. A distributed Epipe consists of two SAPs on different nodes.

By default, QoS policy ID 1 is applied to ingress service SAPs. On egress, QoS policies are associated with a port. Existing filter policies can be associated with service SAPs on ingress and egress.

Meters (defined in SAP-ingress policies) can be applied on ingress, which is associated with SAPs. Scheduler policies can be applied on egress, which is associated with a port.

Ingress and egress SAP parameters can be applied to local and distributed Epipe service SAPs.

For SDP configuration information, see [Configuring an SDP](#). For SDP binding information, see [Configuring SDP bindings](#).

Example

The following example configures a distributed service between ALA-1 and ALA-2:

```
A:ALA-1>epipe 5500 customer 5 create
config>service>epipe$ description "Distributed epipe service to east coast"
config>service>epipe# sap 221/1/3:21 create
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# qos 555
config>service>epipe>sap>ingress# filter ip 1
config>service>epipe>sap>ingress# exit
config>service>epipe>sap# no shutdown
config>service>epipe>sap# exit
config>service>epipe#

A:ALA-2>config>service# epipe 5500 customer 5 create
config>service>epipe$ description "Distributed epipe service to west coast"
config>service>epipe# sap 441/1/4:550 create
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# filter ip 1020
config>service>epipe>sap>ingress# exit
config>service>epipe>sap# egress
config>service>epipe>sap>egress# filter ip 6
config>service>epipe>sap>egress# exit
config>service>epipe>sap# no shutdown
config>service>epipe#
```

Example

The following is a sample of the SAP configurations for ALA-1 and ALA-2:

```
A:ALA-1>config>service# info
-----
...
    epipe 5500 customer 5 vpn 5500 create
        description "Distributed epipe service to east coast"
        sap 221/1/3:21 create
            ingress
                qos 555
                filter ip 1
            exit
        exit
    exit
...
-----
A:ALA-1>config>service#

A:ALA-2>config>service# info
-----
...
    epipe 5500 customer 5 vpn 5500 create
        description "Distributed epipe service to west coast"
        sap 441/1/4:550 create
            ingress
                qos 654
                filter ip 1020
            exit
        exit
    exit
...
-----
```

```
-----
A:ALA-2>config>service#
```

3.8.2.2.1.4 Configuring ingress SAP parameters

By default, QoS policy ID 1 is applied to ingress service SAPs. Existing filter policies or other existing QoS policies can be associated with service SAPs on ingress and egress ports.

Ingress and egress SAP parameters can be applied to local and distributed Epipe service SAPs.

Example

The following is a sample of SAP ingress and egress parameters:

```
ALA-1>config>service# epipe 5500
config>service>epipe# sap 1/1/3:21
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# qos 555
config>service>epipe>sap>ingress# filter ip 1
config>service>epipe>sap>ingress# exit
config>service>epipe>sap#
```

Example

The following is a sample Epipe SAP ingress configuration output:

```
A:ALA-1>config>service#
-----
...
    epipe 5500 customer 5 vpn 5500 create
        description "Distributed epipe service to east coast"
        sap 1/1/3:21 create
            ingress
                qos 555
                filter ip 1
            exit
        exit
        no shutdown
    exit
-----
A:ALA-1>config>service#
```

3.8.2.3 Creating an Epipe service for 7210 SAS-Mxp with range SAPs

Example

The following is a sample of **connection-profile** used to configure a range of SAPs and an Epipe configuration output using the connection profile:

```
*A:7210SAS>config>connprof# info
-----
    ethernet
        ranges 0 2804-2805 2810-2811 2813 2832-2839
    exit
-----

*A:7210SAS>config>service>epipe# info
```

```

-----
description "Default epipe description for service id 292"
sap 1/1/4:292.* create
    description "Default sap description for service id 292"
    exit
exit
sap 1/1/9:cp-292 create
    description "Default sap description for service id 292"
    exit
exit
no shutdown
-----

```

3.8.2.4 Configuring default QinQ SAPs for Epipe transit traffic in a ring scenario in access-uplink mode




Note:

Default QinQ SAPs are only supported on 7210 SAS platforms operating in the access-uplink mode.

In the following figure, 7210-1 is used to deliver some services to customers connected to the device, and additionally it needs to pass through transit from other nodes on the ring (for example, traffic from 7210-2 to 7210-3 or from 7210-2 to 7750-SR onto the core network).

Figure 35: Default QinQ SAP for transit traffic in a ring scenario



Without default QinQ SAPs, the user would need to configure a service on 7210-1, with access-uplink SAPs for each service originating on some other node in the ring. With support for default QinQ SAPs, all traffic that does not need to be delivered to any customer service configured on 7210-1 can be switched using the Epipe service.

Example

The following is a sample configuration output:

```
ALA-1>config>service# epipe 8 customer 1 svc-sap-type null-star create
      sap 1/1/5:*. * create
        statistics
        ingress
        exit
      exit
    exit
  sap 1/1/6:*. * create
    statistics
    ingress
    exit
  exit
exit
no shutdown
exit
```

3.8.2.5 Configuring SDP bindings




Note:

SDPs are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

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.

A spoke-SDP is treated like the equivalent of a traditional bridge "port" where flooded traffic received on the spoke-SDP is replicated on all other "ports" (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

Figure 36: SDPs — unidirectional tunnels



Use the following syntax to create a spoke-SDP binding with an Epipe service.

```
config>service# epipe service-id [customer customer-id]
      spoke-sdp sdp-id:vc-id [vc-type {ether | vlan}]
```

```

vlan-vc-tag 0..4094
egress
    filter {ip ip-filter-id}
    vc-label egress-vc-label
ingress
    filter {ip ip-filter-id}
    vc-label ingress-vc-label
no shutdown

```

The following shows the command usage to bind an Epipe service between ALA-1 and ALA-2. This example assumes that the SAPs have already been configured (see [Distributed Epipe service](#)).

```

A:ALA-1>config>service# epipe 5500
config>service>epipe# spoke-sdp 2:123
config>service>epipe>spoke-sdp# egress
config>service>epipe>spoke-sdp>egress# vc-label 5500
config>service>epipe>spoke-sdp>egress# exit
config>service>epipe>spoke-sdp# ingress
config>service>epipe>spoke-sdp>ingress# vc-label 6600
config>service>epipe>spoke-sdp>ingress# exit
config>service>epipe>spoke-sdp# no shutdown

ALA-2>config>service# epipe 5500
config>service>epipe# spoke-sdp 2:456
config>service>epipe>spoke-sdp# egress
config>service>epipe>spoke-sdp>egress# vc-label 6600
config>service>epipe>spoke-sdp>egress# exit
config>service>epipe>spoke-sdp# ingress
config>service>epipe>spoke-sdp>ingress# vc-label 5500
config>service>epipe>spoke-sdp>ingress# exit
config>service>epipe>spoke-sdp# no shutdown

```

The following is a sample SDP binding for the Epipe service between ALA-1 and ALA-2 configuration output:

```

A:ALA-1>config>service# info
-----
...
    epipe 5500 customer 5 vpn 5500 create
        description "Distributed epipe service to east coast"
        sap 1/1/3:21 create
            ingress
                qos 555
                filter ip 1
            exit
        exit
        spoke-sdp 2:123 create
            ingress
                vc-label 6600
            exit
            egress
                vc-label 5500
            exit
        exit
        no shutdown
    exit
...
-----
A:ALA-1>config>service#

A:ALA-2>config>service# info

```

```

-----
...
exit
    epipe 5500 customer 5 vpn 5500 create
        description "Distributed epipe service to west coast"
        sap 441/1/4:550 create
            ingress
                qos 654
                filter ip 1020
            exit
        exit
    spoke-sdp 2:456 create
        ingress
            vc-label 5500
        exit
        egress
            vc-label 6600
        exit
    exit
    no shutdown
exit
...
-----
A:ALA-2>config>service#

```

3.8.3 Using spoke-SDP control words



Note:

SDPs are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

The **control-word** command provides the option to add a control word as part of the packet encapsulation for PW types for which the control word is optional. On the 7210 SAS, an option is provided to enable it for Ethernet PW (Epipe). The control word might be needed because when ECMP is enabled on the network, packets of a specific PW may be spread over multiple ECMP paths if the hashing router mistakes the PW packet payload for an IPv4 or IPv6 packet. This occurs when the first nibble following the service label corresponds to a value of 4 or 6.

The control word negotiation procedures described in Section 6.2 of RFC 4447 are not supported, and therefore the service will only come up if the same C-bit value is signaled in both directions. If a spoke-SDP is configured to use the control word, but the node receives a label mapping message with a C-bit clear, the node releases the label with an "Illegal C-bit" status code per Section 6.1 of RFC 4447. As soon as the user enables control of the remote peer, the remote peer withdraws its original label and sends a label mapping with the C-bit set to 1 and the VLL service is up in both nodes.

When the control word is enabled, VCCV packets also include the VCCV control word. In that case, the VCCV CC type 1 (OAM CW) is signaled in the VCCV parameter in the FEC. If the control word is disabled on the spoke-SDP, the router alert label is used. In that case, VCCV CC type 2 is signaled. Note that for a multi-segment PW (MS-PW), the CC type 1 is the only type supported, and therefore the control word must be enabled on the spoke-SDP to be able to use VCCV-ping and VCCV-trace.

Example

The following is a sample spoke-SDP control word configuration output:

```

-Dut-B>config>service>epipe# info
-----

```

```


description "Default epipe description for service id 2100"
sap 1/2/7:4 create
description "Default sap description for service id 2100"
exit
spoke-sdp 1:2001 create
control-word
exit
no shutdown
-----
*A:ALA-Dut-B>config>service>epipe#
To disable the control word on spoke-sdp 1:2001:
*A:ALA-Dut-B>config>service>epipe# info
-----
description "Default epipe description for service id 2100"
sap 1/2/7:4 create
description "Default sap description for service id 2100"
exit
spoke-sdp 1:2001 create
exit
no shutdown
-----
*A:ALA-Dut-B>config>service>epipe#

```

3.8.4 Configuring VLL resilience

The following figure shows an example to create VLL resilience. Note that the zero revert-time value means that the VLL path will be switched back to the primary immediately after it comes back up.

Figure 37: VLL resilience



OSSG246

Example: PE1 configuration output

The following is a sample configuration output on PE1:

```

*A:ALA-48>config>service>epipe# info
-----
endpoint "x" create
exit
endpoint "y" create
exit
spoke-sdp 1:100 endpoint "y" create
precedence primary
exit

```

```


    spoke-sdp 2:200 endpoint "y" create
        precedence 1
    exit
    no shutdown
-----
*A:ALA-48>config>service>epipe#

```

3.8.5 Configuring VLL resilience for a switched pseudowire path

Figure 38: VLL resilience with pseudowire switching shows VLL resilience with pseudowire switching.

Figure 38: VLL resilience with pseudowire switching



Example: T-PE1 configuration output

The following is a sample configuration output on TPE1:

```

*A:ALA-48>config>service>epipe# info
-----
    endpoint "x" create
    exit
    endpoint "y" create
    exit
    sap 1/1/1:100 endpoint "x" create
    exit
    spoke-sdp 1:100 endpoint "y" create
        precedence primary
    exit
    spoke-sdp 2:200 endpoint "y" create
        precedence 1
    exit
    spoke-sdp 3:300 endpoint "y" create
        precedence 1
    exit
    no shutdown
-----
*A:ALA-48>config>service>epipe#

```

Example: T-PE2 configuration output

The following is a sample configuration output on TPE2:

```

*A:ALA-49>config>service>epipe# info

```



```

-----
    endpoint "x" create
    exit
    endpoint "y" create
        revert-time 100
    exit
    spoke-sdp 4:400 endpoint "y" create
        precedence primary
    exit
    spoke-sdp 5:500 endpoint "y" create
        precedence 1
    exit
    spoke-sdp 6:600 endpoint "y" create
        precedence 1
    exit
    no shutdown
-----
*A:ALA-49>config>service>epipe#

```

Example: S-PE1 configuration output

The following is a sample configuration output on S-PE1:

```

*A:ALA-50>config>service>epipe# info
-----
...
    spoke-sdp 1:100 create
    exit
    spoke-sdp 4:400 create
    exit
    no shutdown
-----
*A:ALA-49>config>service>epipe#

```

3.8.6 Service management tasks

This section provides information about VLL service management tasks.

3.8.6.1 Modifying Epipe service parameters

The following shows the command usage to add an accounting policy to an existing SAP.

Example:

```

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

```

Example: SAP configuration output

The following is a sample SAP configuration output:

```

ALA-1>config>service# info
-----
    epipe 2 customer 6 vpn 2 create
        description "Distributed Epipe service to east coast"
        sap 1/1/3:21 create

```

```

        accounting-policy 14
        exit
        no shutdown
    exit
-----
ALA-1>config>service#

```

3.8.6.2 Disabling an Epipe service

Use the following syntax to shut down an Epipe service without deleting the service parameters.

```

config>service> epipe service-id
shutdown

```

Example:

```

config>service# epipe 2
config>service>epipe# shutdown
config>service>epipe# exit

```

3.8.6.3 Re-enabling an Epipe service

Use the following syntax to re-enable an Epipe service that was shut down.

```

config>service# epipe service-id
no shutdown

```

Example:

```

config>service# epipe 2
config>service>epipe# no shutdown
config>service>epipe# exit

```

3.8.6.4 Deleting an Epipe service

Perform the following steps before deleting an Epipe service.

1. Shut down the SAP.
2. Delete the SAP.
3. Shut down the service.

Use the following syntax to delete an Epipe service.

```

config>service
[no] epipe service-id
shutdown
[no] sap sap-id
shutdown

```

Example:

```

config>service# epipe 2
config>service>epipe# sap 1/1/3:21
config>service>epipe>sap# shutdown

```

```

config>service>epipe>sap# exit
config>service>epipe# no sap 1/1/3:21
config>service>epipe# epipe 2
config>service>epipe# shutdown
config>service>epipe# exit
config>service# no epipe 2

```

3.9 VLL services command reference

3.9.1 Command hierarchies

- [VLL service configuration commands](#)
 - Epipe Service Configuration Commands
 - [Epipe global commands \(for access-uplink operating mode\)](#)
 - [Epipe global commands \(for network operating mode\)](#)
 - [Epipe SAP configuration commands](#)
 - [Epipe SAP meter override commands](#)
 - [Epipe SAP QoS and filter commands \(for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE \(standalone and standalone-VC\), and 7210 SAS-Sx 10/100GE \(standalone\) devices\)](#)
 - [Epipe SAP QoS and filter commands \(for 7210 SAS-Mxp\)](#)
 - [Epipe SAP statistics commands](#)
 - [Epipe spoke-SDP configuration commands](#)
 - Connection profile commands
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

3.9.1.1 VLL service configuration commands

3.9.1.1.1 Epipe global commands (for access-uplink operating mode)

```

config
- service
- epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {null-
star | dot1q-preserve | any | dot1q-range}] [customer-vid vlan-id]
- no epipe service-id
- description description-string
- no description
- sap sap-id [create] [no-endpoint]
- sap sap-id [create] endpoint endpoint-name
- no sap sap-id
- [no] shutdown

```

3.9.1.1.2 Epipe global commands (for network operating mode)

```

config
- service
- epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {any |
qinq-inner-tag-preserve}] [customer-vid vlan-id] [pbb-epipe]
- no epipe service-id
- description description-string
- no description
- [no] endpoint endpoint-name [create]
- active-hold-delay active-endpoint-delay
- no active-hold-delay
- revert-time [revert-time | infinite]
- no revert-time
- standby-signaling-master
- [no] standby-signaling-master
- sap sap-id [create] [no-endpoint]
- sap sap-id [create] endpoint endpoint-name
- no sap sap-id
- service-mtu octets
- no service-mtu
- [no] service-mtu-check
- [no] shutdown
- spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] [no-endpoint]
- spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] endpoint
- no spoke-sdp sdp-id[:vc-id]
- spoke-sdp-fec spoke-sdp-fec-id [fec fec-type] [aii-type aii-type] [create]
- spoke-sdp-fec spoke-sdp-fec-id no-endpoint
- spoke-sdp-fec spoke-sdp-fec-id [fec fec-type] [aii-type aii-type] [create]
endpoint name [icb]
- no spoke-sdp-fec spoke-sdp-fec-id

```

3.9.1.1.3 Epipe SAP configuration commands

```

config
- service
- epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {any |
qinq-inner-tag-preserve}] [customer-vid vlan-id] [pbb-epipe]
- epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {null-
star | dot1q-preserve | any | dot1q-range}] [customer-vid vlan-id]
- no epipe service-id
- sap sap-id [create] [no-endpoint]
- sap sap-id [create] endpoint endpoint-name
- no sap sap-id
- accounting-policy acct-policy-id
- no accounting-policy acct-policy-id
- [no] cflowd
- [no] collect-stats
- description description-string
- no description
- eth-cfm
- [no] mep mep-id domain md-index association ma-index [direction {up |
down}] primary-vlan-enable
- ais-enable [no]
- [no] client-meg-level [[level [level ...]]]
- [no] interval {1 | 60}
- [no] priority priority-value
- [no] ccm-enable
- [no] ccm-ltm-priority priority

```

```

- [no] description
- [no] eth-test-enable
  - [no] bit-error-threshold bit-errors
  - [no] test-pattern {all-zeros | all-ones} [crc-enable]
- [no] fault-propagation-enable {use-if-tlv | suspendccm}
- low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon |
xcon | noXcon}
- [no] mac-address mac-address
- [no] one-way-delay-threshold seconds
- [no] shutdown
- mip [mac mac address]
- mip default-mac
- no mip
- ethernet
  - [no] llf
- [no] ignore-oper-down
- [no] shutdown

```

3.9.1.1.4 Epipe SAP meter override commands

```

config
- service
  - epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {any |
qinq-inner-tag-preserve}] [customer-vid vlan-id] [pbb-epipe]
  - epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {null-
star | dot1q-preserve | any | dot1q-range}] [customer-vid vlan-id]
  - no epipe service-id
    - sap sap-id [create] [no-endpoint]
    - sap sap-id [create] endpoint endpoint-name
    - no sap sap-id
      - ingress
        - meter-override
          - meter meter-id [create]
          - no meter meter-id
            - adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
            - cbs size [kbits | bytes | kbytes]
            - no cbs
            - mbs size [kbits | bytes | kbytes]
            - no mbs
            - mode mode
            - no mode
            - rate cir cir-rate [pir pir-rate]
            - no rate

```

3.9.1.1.5 Epipe SAP statistics commands

```

config
- service
  - epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {any |
qinq-inner-tag-preserve}] [customer-vid vlan-id] [pbb-epipe]
  - epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {null-
star | dot1q-preserve | any | dot1q-range}] [customer-vid vlan-id]
  - no epipe service-id
    - sap sap-id [create] [no-endpoint]
    - sap sap-id [create] endpoint endpoint-name
    - no sap sap-id
      - statistics
        - ingress

```

- **counter-mode** {in-out-profile-count | forward-drop-count}
- [no] **drop-count-extra-vlan-tag-pkts**

3.9.1.1.6 Epipe spoke-SDP configuration commands



Note:

Spoke-SDP commands are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

```

config
- service]
- epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {any |
qinq-inner-tag-preserve}] [customer-vid vlan-id] [pbb-epipe]
- no epipe service-id
- spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] [no-endpoint]
- spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] endpoint
- no spoke-sdp sdp-id[:vc-id]
- accounting-policy acct-policy-id
- no accounting-policy
- [no] collect-stats
- [no] control-word
- control-channel-status
- acknowledgment
- no acknowledgment
- refresh-timer seconds
- no refresh-timer
- request-timer request-timer request-timer-secs retry-timer retry-timer-
secs timeout-multiplier multiplier
- [no] description
- [no] egress
- [no] vc-label egress-vc-label
- eth-cfm
- [no] mep mep-id domain md-index association ma-index [direction {up |
down}]
- [no] ais-enable
- [no] client-meg-level [[level [level ...]]
- [no] interval {1 | 60}
- [no] priority priority-value
- [no] ccm-enable
- [no] ccm-ltm-priority priority
- [no] description
- [no] eth-test-enable
- [no] bit-error-threshold bit-errors
- [no] test-pattern {all-zeros | all-ones} [crc-enable]
- [no] fault-propagation-enable {use-if-tlv | suspendccm}
- low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon |
xcon | noXcon}
- [no] mac-address mac-address
- [no] one-way-delay-threshold seconds
- [no] shutdown
- mip [mac mac address]
- mip default-mac
- no mip
- [no] force-vlan-vc-forwarding
- hash-label
- hash-label [signal-capability]
- no hash-label
- [no] ingress
- [no] vc-label egress-vc-label
- precedence [precedence-value] primary

```

```

- no precedence
- [no] pw-path-id
  - agi attachment-group-identifier
  - no agi
  - no saii-type2
  - saii-type2 global-id:node-id:ac-id
  - no taii-type2
  - taii-type2 global-id:node-id:ac-id
- no pw-status-signaling
- pw-status-signaling
- [no] shutdown
- vlan-vc-tag 0..4094
- no vlan-vc-tag [0..4094]
- spoke-sdp-fec spoke-sdp-fec-id [fec fec-type] [aii-type aii-type] [create]
- spoke-sdp-fec spoke-sdp-fec-id no-endpoint
- spoke-sdp-fec spoke-sdp-fec-id [fec fec-type] [aii-type aii-type] [create]
endpoint name [icb]
- no spoke-sdp-fec spoke-sdp-fec-id
  - [no] auto-config
  - path name
  - no path
  - precedence prec-value
  - precedence primary
  - no precedence
  - pw-template-bind policy-id
  - no pw-template-bind
  - retry-count retry-count
  - no retry-count
  - retry-timer retry-timer
  - no retry-timer
  - saii-type2 global-id:prefix:ac-id
  - no saii-type2
  - [no] shutdown
  - taii-type2 global-id:prefix:ac-id
  - no taii-type2

```

3.9.1.1.7 Epipe SAP QoS and filter commands (for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), and 7210 SAS-Sx 10/100GE (standalone) devices)

```

config
- service
  - epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {any |
    qinq-inner-tag-preserve}] [customer-vid vlan-id] [pbb-epipe]
  - epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {null-
    star | dot1q-preserve | any | dot1q-range}] [customer-vid vlan-id]
  - no epipe service-id
  - sap sap-id [create] [no-endpoint]
  - sap sap-id [create] endpoint endpoint-name
  - no sap sap-id
  - egress
    - aggregate-meter-rate rate-in-kbps [burst burst-in-kbits] [enable-stats]
    - no aggregate-meter-rate
    - filter [ip ip-filter-id]
    - filter [ipv6 ipv6 -filter-id]
    - filter [mac mac-filter-id]
    - no filter [ip ip-filter-id] [ipv6 ipv6 -filter-id] [mac mac-filter-id]
  - ingress
    - aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
    - no aggregate-meter-rate
    - filter [ip ip-filter-id]
    - filter [ipv6 ipv6-filter-id]

```

```

- filter [mac mac-filter-id]
- no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
- qos policy-id
- no qos

```

3.9.1.1.8 Epipe SAP QoS and filter commands (for 7210 SAS-Mxp)

```

config
- service
- epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {any |
qinq-inner-tag-preserve}] [customer-vid vlan-id] [pbb-epipe]
- no epipe service-id
- sap sap-id [create] [no-endpoint]
- sap sap-id [create] endpoint endpoint-name
- no sap sap-id
- egress
- agg-rate-limit [cir cir-rate] [pir pir-rate]
- no agg-rate-limit
- filter [ip ip-filter-id]
- filter [ipv6 ipv6 -filter-id]
- filter [mac mac-filter-id] (app
- no filter [ip ip-filter-id] [ipv6 ipv6 -filter-id] [mac mac-filter-id]
- qos policy-id
- no qos
- ingress
- aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
- no aggregate-meter-rate
- filter [ip ip-filter-id]
- filter [ipv6 ipv6-filter-id]
- filter [mac mac-filter-id]
- no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
- qos policy-id [enable-table-classification]
- no qos

```

3.9.1.1.9 Connection profile commands

```

config
- connection-profile conn-prof-id [create]
- no connection-profile conn-prof-id
- description description-string
- no description
- ethernet
- no ranges
- ranges vlan ranges [vlan ranges...(up to 32 max)]

```

3.9.1.2 Show commands

```

show
- service
- egress-label start-label [end-label]
- id service-id
- all
- base
- endpoint [endpoint-name]
- epipe

```



```

- labels
- sap sap-id [detail]
- sdp [sdp-id | far-end ip-addr] [detail]
- split-horizon-group [group-name]
- stp [detail]]
- saii-type2-using global-id[:prefix[:ac-id]]
- sap-using [sap sap-id]
- sap-using [ingress | egress] filter filter-id
- sap-using [ingress] qos-policy qos-policy-id
- service-using [epipe] [vpls] [mirror] [b-vpls] [i-vpls] [m-vpls] [sdp sdp-id]
[customer customer-id]
- sdp [sdp-id | far-end ip-address] [detail | keep-alive-history]
- spoke-sdp-fec-using [spoke-sdp-fec-id spoke-sdp-fec-id] [saii-type2 global-
id:prefix:ac-id] [taii-type2 global-id:prefix:ac-id] [path name] [expired]
- taii-type2-using global-id[:prefix[:ac-id]]
show
- connection-profile [conn-prof-id] [associations]

```

3.9.1.3 Clear commands

```

clear
- service
- id service-id
- statistics
- id service-id
- counters
- sap sap-id {all | cem | counters | stp | l2pt}

```

3.9.2 Command descriptions

- [VLL service configuration commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

3.9.2.1 VLL service configuration commands

- [Generic commands](#)
- [VLL global commands](#)
- [VLL SAP commands](#)
- [Service filter and QoS policy commands](#)
- [Epipe Service SAP Statistics Commands](#)
- [VLL SDP Commands](#)
- [Connection profile commands](#)

3.9.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>service>epipe

config>service>epipe>sap

config>service>epipe>spoke-sdp (not supported in access-uplink operating mode)

config>connection-profile

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

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

Default

No description associated with the configuration context.

Parameters

string

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

shutdown

Syntax

[no] **shutdown**

Context

config>service>epipe

```
config>service>epipe>sap
config>service>epipe>spoke-sdp (not supported in access-uplink operating mode)
config>service>epipe>sap>eth-cfm>mep
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

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

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

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

3.9.2.1.2 VLL global commands

epipe

Syntax

```
epipe service-id [customer customer-id] [create] [vpn vpn-id] [svc-sap-type {any | qinq-inner-tag-preserve}] [customer-vid vlan-id] [pbb-epipe]
```

```
no epipe service-id
```

Context

```
config>service
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command configures an Epipe service instance. This command is used to configure a point-to-point Epipe service. An Epipe connects two endpoints defined as Service Access Points (SAPs). In a local service, the SAPs may be defined in one 7210 SAS node and in distributed service the SAPs may be defined on two different 7210 SAS nodes.

On 7210 SAS, platforms operating in network mode, both local and distributed services are supported.

MAC learning and filtering are not supported on an Epipe service.

When a service is created, the **customer** keyword and *customer-id* parameter must be specified to associate the service with a customer. The *customer-id* value must already exist having been created using the **customer** command in the service context. When 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.

When a service is created, the use of the **customer** *customer-id* command is optional for navigating into the service configuration context. Editing a service with the incorrect *customer-id* value specified results in an error.

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 and all instances of SAPs, mesh, or spokes have been deleted from the service.

Default

No Epipe services exist until they are explicitly created with this command.

Parameters

service-id

The unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values *service-id*: 1 to 2147483647
 svc-name: 64 characters maximum

customer *customer-id*

Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn *vpn-id*

Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN ID. If this parameter is not specified, the VPN ID uses the same service ID number.

Values 1 to 2147483647

Default null (0)

svc-sap-type

Specifies the service type and the allowed SAPs in the service. When the **pbb-epipe** keyword is not configured, a plain Epipe service can be configured with SAPs and SDPs.

Values **any** - Specifies that all supported SAPs are allowed in the service. See [section QinQ SAP Configuration Restrictions for 7210 SAS platforms in network operating mode](#) for information about restrictions related to QinQ SAPs.

qinq-inner-tag-preserve - Specifies that an Epipe service processes and forwards packets received with 3 or more tags on a QinQ SAP. See [Processing packets received with more than two tags on a QinQ SAP in Epipe service](#) for more information.

Default any

create

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

pbb-epipe

Specifies a PBB Epipe service, which allows the software to allocate the appropriate resources for PBB Epipe. Only SAPs can be configured with a PBB Epipe service.



Note:

PBB and the **pbb-epipe** keyword are supported only on 7210 SAS-T operating in the network mode.

epipe

Syntax

epipe *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**svc-sap-type** {**null-star** | **dot1q-preserve** | **any** | **dot1q-range**}] [**customer-vid** *vlan-id*]
no epipe *service-id*

Context

config>service

Platforms

Only supported on 7210 SAS platforms configured in the access-uplink operating mode

Description

This command configures an Epipe service instance. This command is used to configure a point-to-point Epipe service. An Epipe connects two endpoints defined as SAPs. In a local service, the SAPs may be defined in one 7210 SAS node and in distributed service the SAPs may be defined on two different 7210 SAS nodes.



Note:

Distributed services are only supported on 7210 SAS platforms operating in the network mode. 7210 SAS platforms operating in access-uplink mode only support local SAP-to-SAP service.

MAC learning and filtering are not supported on an Epipe service.

When a service is created, the **customer** keyword and *customer-id* parameter must be specified to associate the service with a customer. The *customer-id* parameter must already exist having been created using the **customer** command in the service context. When 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.

When a service is created, the use of the **customer** *customer-id* command is optional for navigating into the service configuration context. Editing a service with the incorrect *customer-id* value specified results in an error.

The **no** form of this command deletes the Epipe service instance with the specified *service-id* value. The service cannot be deleted until it has been shut down and all instances of SAPs have been deleted from the service.

Default

No Epipe services exist until they are explicitly created with this command.

Parameters

service-id

Specifies the unique service identification number or string identifying the service in the service domain. The ID must be unique to the service and may not be used for any other service of any type. The *service-id* must be the same value for every 7210 SAS device on which this service is defined.

Values *service-id*: 1 to 2147483647
 svc-name: 64 characters maximum

customer *customer-id*

Specifies the customer ID to be associated with the service. This parameter is required on service creation and is optional for service editing or deletion.

Values 1 to 2147483647

vpn vpn-id

Specifies the VPN ID , which allows you to identify virtual private networks (VPNs). If this parameter is not specified, the VPN ID uses the same number as the service ID.

Values 1 to 2147483647

Default null (0)

svc-sap-type

Specifies the service type and the SAPs allowed in the service.

Values **null-star** - Specifies the allowed SAP in the service, which can be null SAP, dot1q default SAP, Q.* SAP, or default QinQ SAP (also known as *.* SAP).

dot1q-preserve - Specifies that the allowed access SAPs in the service are dot1q. The dot1q ID is not stripped after packets match the SAP.

dot1q-range - Specifies that the access SAP in the service can use VLAN ranges as the SAP tags. The VLAN ranges are configured using the **configure>connection-profile** CLI command. On ingress

of the access dot1q SAP using VLAN ranges, the outermost tag is not removed before forwarding.

any - Specifies that the SAPs allowed in the service are defined as shown in [Table 8: SAP and service combinations for 7210 SAS-T in access-uplink mode](#). See the [SAP configuration notes for 7210 SAS platforms in access-uplink operating mode](#) section for more information.

Default any

customer-vid *vlan-id*

Specifies the dot1q VLAN ID to be used while creating the local dot1q SAP for the **svc-sap-type** command with the **dot1q-preserve** option specified.

Values 1 to 4094

create

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

endpoint

Syntax

[no] **endpoint** *endpoint-name*

Context

config>service>epipe

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures a service endpoint.

Parameters

endpoint-name

Specifies an endpoint name.

active-hold-delay

Syntax

active-hold-delay *active-hold-delay*

no active-hold-delay

Context

config>service>epipe>endpoint

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies that the node delays sending the change in the T-LDP status bits for the VLL endpoint when the MC-LAG transitions the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby** or when any object in the endpoint. For example, SAP, ICB, or regular spoke-SDP, transitions from up to down operational state.

By default, when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby**, the node sends immediately new T-LDP status bits indicating the new value of "standby" over the spoke-SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

There is no delay applied to the VLL endpoint status bit advertisement when the MC-LAG transitions the LAG subgroup which hosts the SAP from "standby" to "active" or when any object in the endpoint transitions to an operationally up state.

Default

0 — A value of zero means that when the MC-LAG transitioned the LAG subgroup which hosts the SAP for this VLL endpoint from **active** to **standby**, the node sends immediately new T-LDP status bits indicating the new value of **standby** over the spoke-SDPs which are on the mate-endpoint of the VLL. The same applies when any object in the endpoint changes an operational state from up to down.

Parameters

active-hold-delay

Specifies the active hold delay in 100s of milliseconds.

Values 0 to 60

revert-time

Syntax

revert-time [*revert-time* | **infinite**]

no revert-time

Context

config>service>epipe>endpoint

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the time to wait before reverting back to the primary spoke-SDP defined on this service endpoint, after having failed over to a backup spoke-SDP.

The **no** form of this command resets the timer to the default value of 0.

Parameters

revert-time

Specify 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>vll>endpoint

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the pseudowire standby bit (value 0x00000020) to be sent to T-LDP peer for each spoke SPD of the endpoint that is selected as a standby.

This command is mutually exclusive with a VLL mate SAP created on an MC-LAG or ICB. It is also mutually exclusive with vc-switching.

Default

no standby-signaling-master

service-mtu

Syntax

service-mtu *octets*

no service-mtu

Context

config>service>epipe

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the service payload (Maximum Transmission Unit – MTU), in bytes, for the service. This MTU value overrides the service-type default MTU. The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP binding operational state within the service.

The service MTU and a SAP service delineation encapsulation overhead (that is, 4 bytes for a dot1q tag) is used to derive the required MTU of the physical port or channel on which the SAP was created. If the required payload is larger than the port or channel MTU, the SAP is placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP is able to transition to the operative 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.

Note: To disable service MTU check execute the command **no service-mtu-check**. Disabling service MTU check allows the packets to pass to the egress if the packet length is lesser than or equal to the MTU configured on the port.

Default

VPLS: 1514

The following table displays MTU values (in octets) for specific VC types.

Table 25: MTU values for specific VC types

VC-type	Example service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
VPLS	1514	1500
VPLS (with preserved dot1q)	1518	1504
VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (QinQ with preserved bottom Qtag)	1518	1504

Parameters

octets

Specifies the size of the MTU in octets, expressed as a decimal integer.

Values 1 to 9194

service-name

Syntax

service-name *service-name*

no service-name

Context

config>service>epipe

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures an optional service name, up to 64 characters, which adds a name identifier to a specific service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7210 SAS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a specific service when it is initially created.

Parameters

service-name

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

service-mtu-check

Syntax

[no] service-mtu-check

Context

config>service>epipe

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

The **no** form of this command disables the service MTU check.

Disabling service MTU check allows the packets to pass to the egress if the packet length is lesser than or equal to the MTU configured on the port. The length of the packet sent from a SAP is limited only by

the access port MTU. In case of a pseudowire the length of a packet is limited by the network port MTU (including the MPLS encapsulation).

**Note:**

If TLDP is used for signaling, the configured value for service-mtu is used during a pseudowire setup.

Default

enabled

3.9.2.1.3 VLL SAP commands

```
sap
```

Syntax

```
sap sap-id [create] [no-endpoint]
```

```
sap sap-id [create] endpoint endpoint-name
```

```
no sap sap-id
```

Context

```
config>service>epipe
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates a SAP within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the 7210 device. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP does 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.

In a single physical port only one SAP can belong to one service. Multiple SAPs can be defined over a physical port but each of these SAPs should belong to different service. This is true only for access-uplink mode. That is, for network mode, multiple SAPs on the same port can belong to the same service.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port. Additionally, in access-uplink mode, SAPs can be defined also on access-uplink port. Access-uplink SAPs are network facing SAPs representing Dot1q or QinQ tunnels used to transport traffic toward the service nodes.

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 is discarded.

The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The following encapsulations are supported:

- Ethernet access SAPs support null, dot1q
- Ethernet access-uplink SAPs support only QinQ encapsulation.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP are also deleted.

Default

No SAPs are defined.

Special Cases

Default SAPs

A default SAP has the following format: port-id:*. This type of SAP is supported only on Ethernet MDAs and its creation is allowed only in the scope of Layer 2 services (Epipe and VPLS).

Parameters

sap-id

Specifies the physical port identifier portion of the SAP. See [Common CLI command descriptions](#) for command syntax.

create

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

no-endpoint

Keyword used to remove the association of a SAP or a spoke-SDP with an explicit endpoint name. This keyword is not supported on platforms operating in the access-uplink operating mode.

endpoint-name

Specifies to add a SAP endpoint association, up to 32 characters. This parameter is not supported on platforms configured in the access-uplink operating mode.

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

config>service>epipe>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

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

Default accounting policy.

Parameters

acct-policy-id

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

Values 1 to 99

cflowd

Syntax

[no] cflowd

Context

config>service>epipe>sap

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE

Description

This command enables cflowd to collect traffic flow samples through a service interface (SAP) for analysis. When cflowd is enabled on an Ethernet service SAP, the Ethernet traffic can be sampled and processed by the system's cflowd engine and exported to IPFIX collectors with the **I2-ip** template enabled.

Cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement. When cflowd is enabled at the SAP level, all packets forwarded by the interface are subjected to analysis according to the cflowd configuration.

**Note:**

See the "Configuration notes" section in the "Cflowd" chapter of the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide* for more information about the sampling of packets.

The **no** form of this command disables cflowd to collect traffic flow samples through a SAP.

Default

no cflowd

collect-stats**Syntax**

[no] **collect-stats**

Context

config>service>epipe>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command enables accounting and statistical data collection for either the SAP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued, the statistics are still accumulated by the cards. However, the CPU 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

no collect-stats

description**Syntax**

description *description-string*

no description

Context

config>service>epipe>sap

config>service>epipe>spoke-sdp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command defines an ASCII string associated with egress-multicast-group-name.

The **no** form of this command removes an existing description string from egress-multicast-group.

Parameters

description-string

Specifies the description as an ASCII string of up to 80 characters. Only printable 127 bit ASCII characters are allowed. If the string contains spaces, the string must be specified with beginning and ending quotes.

Values An ASCII string up to 80 characters.

eth-cfm

Syntax

eth-cfm

Context

config>service>epipe>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure ETH-CFM parameters.

mep

Syntax

mep *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {**up** | **down**}] **primary-vlan-enable**

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

Context

config>service>epipe>sap>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command provisions the maintenance endpoint (MEP).

For more information about ETH-CFM support, see the *7210 SAS-Mxp, R6, R12, S, Sx, T OAM and Diagnostics Guide*.

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

Parameters

mep-id

Specifies the maintenance association end point identifier.

Values 1 to 8191

md-index

Specifies the maintenance domain (MD) index value.

Values 1 to 4294967295

ma-index

Specifies the MA index value.

Values 1 to 4294967295

direction up | down

Indicates the direction in which the maintenance association (MEP) faces on the bridge port. Direction is not supported when a MEP is created directly under the `vp1s>eth-cfm` construct (vMEP).

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

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

primary-vlan-enable

Provides a method for linking the with the primary VLAN configured under the bridge-identifier for the MA. This is only allowed if the mhf-creation method is static. MIPs cannot be changed from or to primary VLAN functions without first being deleted. This must be configured as part of the creation step and can only be changed by deleting the MEP and recreating it. Primary VLANs are only supported under Ethernet SAPs.

This parameter is only supported on 7210 SAS-T (network operating mode), 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), and 7210 SAS-Mxp (network operating mode).

ais-enable

Syntax

[no] ais-enable

Context

config>service>epipe>sap>eth-cfm>mep
config>service>epipe>spoke-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables the generation and the reception of AIS messages.

client-meg-level

Syntax

client-meg-level [[/level [/level ...]]
no client-meg-level

Context

config>service>epipe>sap>eth-cfm>mep>ais-enable
config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the client maintenance entity group (MEG) levels to use for AIS message generation. Up to 7 levels can be provisioned with the restriction that the client MEG level must be higher than the local MEG level.

Parameters

<i>level</i>	Specifies the client MEG level.
Values	1 to 7
Default	1

interval

Syntax

interval {1 | 60}

no interval

Context

config>service>vpls>epipe>eth-cfm>mep>ais-enable

config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the transmission interval of AIS messages in seconds.

Parameters

1 | 60

Specifies the transmission interval of AIS messages in seconds.

Default 1

priority

Syntax

priority *priority-value*

no priority

Context

config>service>vpls>epipe>eth-cfm>mep>ais-enable

config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the priority of AIS messages originated by the node.

Parameters

priority-value

Specifies the priority value of the AIS messages originated by the node.

ccm-enable

Syntax

[no] ccm-enable

Context

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

config>service>epipe>spoke-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables the generation of CCM messages.

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

ccm-ltm-priority

Syntax

ccm-ltm-priority *priority*

no ccm-ltm-priority

Context

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

config>service>epipe>spoke-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

The **no** form of this 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

eth-test-enable

Syntax

[no] **eth-test-enable**

Context

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

config>service>epipe>spoke-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

For ETH-test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands:

oam eth-cfm eth-test *mac-address mep mep-id domain md-index association ma-index* [**priority** *priority*] [**data-length** *data-length*]

A check is done for both the provisioning and test to ensure that the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP indicates the problem.

bit-error-threshold

Syntax

bit-error-threshold *errors*

no bit-error-threshold

Context

config>service>epipe>sap>eth-cfm>mep>eth-test-enable

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command specifies the threshold value of bit errors.

test-pattern

Syntax

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

no test-pattern

Context

config>service>epipe>sap>eth-cfm>mep>eth-test-enable

config>service>epipe>spoke-sdp>eth-cfm>mep>eth-test-enable (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the test pattern for eth-test frames.

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

Parameters

all-zeros

Specifies to use all zeros in the test pattern.

all-ones

Specifies to use all ones in the test pattern.

crc-enable

Generates a CRC checksum.

Default all-zeros

fault-propagation-enable

Syntax

fault-propagation-enable {use-if-tlv | suspend-ccm}

no fault-propagation-enable

Context

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

config>service>epipe>spoke-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the fault propagation for the MEP.

Parameters

- use-if-tlv*
Specifies to use the interface TLV.
- suspend-ccm*
Specifies to suspend the continuity check messages.

low-priority-defect

Syntax

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

Context

- config>service>epipe>sap>eth-cfm>mep
- config>service>epipe>spoke-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

Default

macRemErrXcon

Values	allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
	macRemErrXcon	Only DefMACstatus, DefRemoteCCM, Def ErrorCCM, and DefXconCCM
	remErrXcon	Only DefRemoteCCM, DefErrorCCM, and Def XconCCM
	errXcon	Only DefErrorCCM and DefXconCCM
	xcon	Only DefXconCCM; or
	noXcon	No defects DefXcon or lower are to be reported

mac-address

Syntax

mac-address *mac-address*

no mac-address

Context

config>service>epipe>spoke-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

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

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the MAC address of the MEP.

The **no** form of this command reverts the MAC address of the MEP back to that of the port (if the MEP is on a SAP) or the bridge (if the MEP is on a spoke).

Parameters

mac-address

Specifies the MAC address of the MEP.

Values	6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MEP. Must be unicast. Using the all zeros address is equivalent to the no form of this command.

agg-rate-limit

Syntax

agg-rate-limit *agg-rate*

no agg-rate-limit

Context

config>service>epipe>sap>egress

config>service>vppls>sap>egress

config>service>ies>sap>egress

config>service>vprr>sap>egress (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command defines a maximum total rate for all egress queues on a service SAP.

The SAP aggregate rate can be used only if SAP based scheduling mode is configured at the port level. It is not supported in FC-based scheduling mode.

When configured in SAP-based scheduling mode, the egress port scheduler distributes the available bandwidth to all the SAPs configured on the port, up to the configured aggregate rate for the SAP.

The **no** form of this command removes the aggregate rate limit from the SAP.

Parameters

agg-rate

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

Values 1 to 40000000, max

tod-suite

Syntax

tod-suite *tod-suite-name*

no tod-suite

Context

config>service>epipe>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command applies a time-based policy (filter or QoS policy) to the service SAP. The suite name must already exist in the **config>cron** context.

Default

no tod-suite

Parameters

tod-suite-name

Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP. The suite can be applied to more than one SAP.

one-way-delay-threshold

Syntax

one-way-delay-threshold *seconds*

Context

config>service>vpls>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command enables/disables eth-test functionality on MEP.

Parameters

seconds

Specifies the one way delay threshold in seconds.

Values 0 to 600

Default 3

mip

Syntax

mip [**mac** *mac-address*]

mip default-mac

no mip

Context

config>service>epipe>sap>eth-cfm

config>service>epipe>spoke-sdp>eth-cfm (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command allows Maintenance Intermediate Points (MIPs) to be created if mhf-creation for the MA is configured using the default option.

Parameters

mac-address

Specifies the MAC address of the MIP.

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MIP. The MAC must be unicast. Using the all zeros address is equivalent to the no form of this command.

default-mac

Using the no command deletes the MIP. If the operator needs to change the mac back to the default mac without having to delete the MIP and reconfiguring this command is useful.

Default no mip

ethernet

Syntax

ethernet

Context

config>service>epipe>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command configures Ethernet properties in this SAP.

llf

Syntax

[no] llf

Context

config>service>epipe>sap>ethernet

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command enables Link Loss Forwarding (LLF) on an Ethernet port. It provides an end-to-end OAM fault notification for Ethernet VLL service. LLF on an Ethernet port brings down the port when there is a

local fault on the pseudowire or service, or a remote fault on the SAP or pseudowire, signaled with label withdrawal or TLDP status bits. It ceases when the fault disappears.

The Ethernet port must be configured for null encapsulation.

The **no** form of this command disables LLF.

ignore-oper-down

Syntax

[no] ignore-oper-down

Context

config>service>epipe>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command enables the user to configure the optional command for a specific SAP to ignore the transition of the operational state to down when a SAP fails. Only a single SAP in an Epipe may use this option.

Default

no ignore-oper-down

3.9.2.1.4 Service filter and QoS policy commands

egress

Syntax

egress

Context

config>service>epipe>spoke-sdp (not supported in access-uplink mode)

config>service>epipe>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

Commands in this context configure egress SAP parameters.

force-vlan-vc-forwarding

Syntax

[no] force-vlan-vc-forwarding

Context

config>service>epipe>spoke-sdp

config>service>vpls>spoke-sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command forces vc-vlan-type forwarding in the datapath for spoke which have either vc-type. This command is not allowed on vlan-vc-type SDPs.

The **no** version of this command sets default behavior.

Default

Per default this feature is disabled

hash-label

Syntax

hash-label [signal-capability]

no hash-label

Context

config>service>epipe>spoke-sdp

Platforms

7210 SAS-Mxp, 7210 SAS-R6 IMM-b, 7210 SAS-R6 IMM-c, 7210 SAS-R12 IMM-b, 7210 SAS-R12 IMM-c, and 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC)

Description

This command enables the use of a hash label on a VLL or VPLS service bound to LDP or RSVP SDP, using the autobind mode with the **ldp**, **rsvp-te**, or **mpls** options. When this command is enabled, the ingress datapath is modified such that the result of the hash on the packet header is communicated to the egress datapath for use as the value of the label field of the hash label. Only the hash-2 parameters

are used to compute the hash label, even if the SDP is over a lag (with **load-balancing** set as **hash-1** or **hash-2**) or a port. The egress datapath adds the hash label at the bottom of the stack (BoS) and sets the S-bit to one.



Note:

On the 7210 SAS, the hash label is not used on the local node for ECMP hashing and LAG hashing. It is available for use by LSR nodes, through which the traffic flows and which are capable of using the labels for hashing.

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

The TTL of the hash label is set to a value of 0.

The user enables the signaling of the hash label capability under a VLL spoke-SDP interface, a VPLS spoke-SDP interface, or a VPLS mesh-SDP interface by configuring the **signal-capability** option. In this case, the decision to insert the hash label on the user and control plane packets by the local PE is solely determined by the outcome of the signaling process and can override the local PE configuration. The following procedures apply when the **hash-label** command and the **signal-capability** option are enabled on the local PE.

- The 7210 SAS local PE inserts the Flow Label Interface Parameters sub-TLV with T=1 and R=1 in the PW ID FEC element in the label mapping message for that spoke SDP or mesh SDP.
- If the remote PE does not send the Flow Label sub-TLV in the PW ID FEC element, or sends a Flow Label sub-TLV in the PW ID FEC element with T=FALSE and R=FALSE, the local node disables the hash label capability. Consequently, the local PE node does not insert a hash label in user and control plane packets it forwards on the spoke SDP or mesh SDP. It also drops user and control plane packets received from remote PE if they include a hash label. The latter may also be caused by a remote 7210 SAS PE that does not support the **hash-label** command, has the **hash-label** command enabled but does not support the **signal-capability** option, or does support both options but the user did not enable them because of a misconfiguration.
- If the remote PE sends the Flow Label sub-TLV in the PW ID FEC element with T=TRUE and R=TRUE, the local PE enables the hash label capability. The local PE inserts a hash label in user and control plane packets it forwards on the spoke-SDP or mesh SDP. It also accepts user and control plane packets from the remote PE with or without a hash label.
 - If the **hash-label** command was enabled on the local configuration of the spoke SDP or mesh SDP at the remote PE, the pseudowire packets received by the local PE have the hash label included. These packets must be dropped. The only way to solve this is to disable the **signaling-capability** option on the local node, which results in the insertion of the hash label by both PE nodes.
 - If the **hash-label** command is not supported or was not enabled on the local configuration of the spoke SDP or mesh SDP at the remote PE, the pseudowire received by the local PE does not have the hash label included.

The user can enable or disable the **signal-capability** option in CLI as required. When doing so, the router must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the PW ID FEC element.

**Note:**

- This feature is supported only for VLL and VPLS services. It is not supported for VPRN services. It is also not supported on multicast packets forwarded using RSVP P2MP LSPs or mLDP LSPs in both the base router instance and in the multicast VPN (mVPN) instance.
- On the 7750 SR and possibly other vendor implementations, to allow applications where the egress LER infers the presence of the hash label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the hash label. That is, the value of the hash label is always in the range of 524,288 to 1,048,575 and does not overlap with the signaled or static LSP and signaled or static service label ranges. This also guarantees that the hash label does not match a value in the reserved label range. This is not supported on the 7210 SAS for service traffic (for MPLS OAM traffic the MSB bit is set). That is, 7210 SAS devices do not set the MSB in the hash label value for service traffic. Consequently, the user must ensure that both ends are correctly configured to either process hash labels or disable them.

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

Default

no hash-label

Parameters**signal-capability**

Enables the signaling and negotiation of the hash label between the local and remote PE nodes.

ingress

Syntax

ingress

Context

config>service>epipe>sap

config>service>epipe>sap>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure ingress SAP Quality of Service (QoS) policies.

If no sap-ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing.

aggregate-meter-rate

Syntax

aggregate-meter-rate *rate-in-kbps* [**burst** *burst-in-kbits*] [**enable-stats**]

no aggregate-meter-rate

Context

config>service>vpls>sap>egress

config>service>epipe>sap>egress (not supported on 7210 SAS-Mxp)

config>service>ies>sap>egress

config>service>vprn>sap>egress (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the SAP egress aggregate policer. The rate (PIR) of the SAP egress aggregate policer must be specified by the user. The user can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the traffic sent out of the SAP and determines if the packet is either forwarded or dropped.

An option is provided to associate a set of two counters to calculate total forwarded packets and octets, and total dropped packets and octets. When the counter is enabled, the resources required increases to twice the resources required when the counter is not used. If enable-stats keyword is specified during the creation of the meter, the counter is allocated by software (if available). To free up the counter and relinquish its use, the user can use the **no aggregate-meter-rate** command, and then recreate the meter again using the aggregate-meter rate command.

If egress frame-based accounting is used, the SAP egress aggregate meter rate accounts for the Ethernet frame overhead. The system accounts for 12 bytes of IFG and 8 bytes of start delimiter. Frame-based counting does not affect the count of octets maintained by the counter (if in use).



Note:

- Before enabling this command for a SAP, resources must be allocated to this feature from the egress-internal-tcam resource pool using the command **configure>system>resource-profile>egress-internal-tcam>egress-sap-aggregate-meter**. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for more information.
- The egress aggregate meter is not FC aware. The forward and drop decisions are taken based on the order the packets are sent out of the SAP by the egress port scheduler.

The **no** form of this command removes the egress aggregate policer from use.

Default

no aggregate-meter-rate

Parameters

rate-in-kbps

Specifies the rate in kilobits per second.

Values 1 to 20000000 | max

Default max

burst <burst-in-kilobits>

Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 to 2146959

Default 512

enable-stats

Specifies if counter to count forwarded and dropped count must be allocated or not. If this keyword is used while configuring the meter, counter is allocated.

aggregate-meter-rate

Syntax

aggregate-meter-rate *rate-in-kbps* [**burst** *burst-in-kbits*]

no aggregate-meter-rate

Context

config>service>vpls>sap>ingress

config>service>epipe>sap>ingress

config>service>ies>sap>ingress

config>service>vprn>sap>ingress (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including platforms configured in the access-uplink operating mode; not supported on 7210 SAS-Sx 10/100GE

Description

This command configures the SAP ingress aggregate policer. The rate of the SAP ingress aggregate policer must be specified by the user. The user can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the ingress traffic on different FCs and determines the final disposition of the packet. The packet is either forwarded to an identified profile or dropped.



Note:

- When operating the 7210 SAS in access-uplink mode, this command is available only for access SAPs; it is not supported for access-uplink SAPs.

- The sum of CIR of the individual FCs configured under the SAP cannot exceed the PIR rate configured for the SAP. Though the 7210 SAS software does not block this configuration, it is not recommended for use.

The following table provides information about the final disposition of the packet based on the operating rate of the per FC policer and the per SAP aggregate policer.

Table 26: Final disposition of the packet based on per-FC and per-SAP policer or meter

Per FC meter operating rate	Per FC assigned color	SAP aggregate meter operating rate	SAP aggregate meter color	Final packet color
Within CIR	Green	Within PIR	Green	Green or In-profile
Within CIR ¹³	Green	Above PIR	Red	Green or In-profile
Above CIR, Within PIR	Yellow	Within PIR	Green	Yellow or Out-of-Profile
Above CIR, Within PIR	Yellow	Above PIR	Red	Red or Dropped
Above PIR	Red	Within PIR	Green	Red or Dropped
Above PIR	Red	Above PIR	Red	Red or Dropped

When the SAP aggregate policer is configured, per FC policer can be only configured in "trtcm2" mode (RFC 4115).

The meter modes "srtcm" and "trtcm1" are used in the absence of an aggregate meter.

The SAP ingress meter counters increment the packet or octet counts based on the final disposition of the packet.

If ingress Frame-based accounting is used, the SAP aggregate meter rate accounts for the Ethernet frame overhead. The system accounts for 12 bytes of IFG and 8 bytes of start delimiter.

The **no** form of this command removes the aggregate policer from use.

Default

no aggregate-meter-rate

¹³ This row is not recommended for use. For more information, see the Note in the "[aggregate-meter-rate](#)" description.

Parameters

rate-in-kbps

Specifies the rate in kilobits per second.

Values 1 to 20000000 | max

Default max

burst burst-in-kilobits

Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 to 2146959

Default 512

filter

Syntax

filter [**ip** *ip-filter-id*]

filter [**ipv6** *ipv6-filter-id*]

filter [**mac** *mac-filter-id*]

no filter [**ip** *ip-filter-id*]

no filter [**ipv6** *ipv6-filter-id*]

no filter [**mac** *mac-filter-id*]

Context

config>service>epipe>sap>egress

config>service>epipe>sap>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command associates an IP 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. Only one filter can be applied to a SAP at a time.

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

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

**Note:**

For filter support available on different 7210 SAS platforms, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.

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.

Special Cases**Epipe**

Both MAC and IP filters are supported on an Epipe service SAP.

Parameters**ip *ip-filter-id***

Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

ipv6 *ipv6-filter-id*

Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

mac *mac-filter-id*

Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 to 65535

meter-override**Syntax**

[no] meter-override

Context

config>service>epipe>sap>ingress

config>service>vppls>sap>ingress

config>service>ies>interface>sap>ingress

config>service>vprn>interface>sap>ingress (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command, within the SAP ingress contexts, is used to create a CLI node for specific overrides to one or more meters created on the SAP through the sap-ingress QoS policies.

The **no** form of this command is used to remove any existing meter overrides.

Default

no meter-overrides

meter

Syntax

meter *meter-id* [**create**]

no meter *meter-id*

Context

config>service>epipe>sap>ingress>meter-override

config>service>vpls>sap>ingress>meter-override

config>service>ies>interface>sap>ingress>meter-override

config>service>vprn>interface>sap>ingress>meter-override (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command, within the SAP ingress contexts, is used to create a CLI node for specific overrides to a specific meter created on the SAP through a sap-ingress QoS policies.

The **no** form of this command is used to remove any existing overrides for the specified meter-id.

Parameters

meter-id

The meter-id parameter is required when executing the meter command within the meter-overrides context. The specified meter-id must exist within the sap-ingress QoS policy applied to the SAP. If the meter is not currently used by any forwarding class or forwarding type mappings, the meter does not exist on the SAP. This does not preclude creating an override context for the meter-id.

create

The create keyword is required when a meter *meter-id* override node is being created and the system is configured to expect explicit confirmation that a new object is being created. When the system is not configured to expect explicit confirmation, the create keyword is not required.

adaptation-rule

Syntax

adaptation-rule [*pir adaptation-rule*] [*cir adaptation-rule*]

no adaptation-rule

Context

config>service>epipe>sap>ingress>meter-override>meter

config>service>vpls>sap>ingress>meter-override>meter

config>service>ies>interface>sap>ingress>meter-override>meter

config>service>vprn>interface>sap>ingress>meter-override>meter (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

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

Default

no adaptation-rule

Parameters

pir

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



Note:

When the meter mode in use is 'trtcm2', this parameter is interpreted as EIR value. For more information, see the description and relevant notes for meter modes in the *7210 SAS-Mxp, R6, R12, S, Sx, T Quality of Service Guide*.

cir

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

adaptation-rule

Specifies the criteria to use to compute the operational CIR and PIR values for this meter, while maintaining a minimum offset.

Values

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

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

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

cbs**Syntax**

cbs *size* [**kbits** | **bytes** | **kbytes**]

no cbs

Context

config>service>epipe>sap>ingress>meter-override>meter

config>service>vpls>sap>ingress>meter-override>meter

config>service>ies>interface>sap>ingress>meter-override>meter

config>service>vprn>interface>sap>ingress>meter-override>meter (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command provides a mechanism to override the default CBS for the meter. The committed burst size parameter specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the CBS value then the packets are marked as in-profile by the meter to indicate that the traffic is complying meter configured parameters.

The **no** form of this command returns the CBS size to the default value.

Default

32 kbits

Parameters

size

Specifies the value in either kbits, kilobytes or bytes.

Values kbits : [4..2146959 | default]
 bytes : [512..274810752]
 kbytes : [1..268369]

mbs

Syntax

mbs *size* [kbits | bytes | kbytes]

no mbs

Context

config>service>epipe>sap>ingress>meter-override>meter

config>service>vpls>sap>ingress>meter-override>meter

config>service>ies>interface>sap>ingress>meter-override>meter

config>service>vprn>interface>sap>ingress>meter-override>meter (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command provides a mechanism to override the default MBS for the meter. The maximum burst size parameter specifies the maximum burst size that can be transmitted by the source while still complying with the CIR. If the transmitted burst is lower than the MBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying meter configured parameters.

The **no** form of this command returns the MBS size to the default value.

Default

512kbits

Parameters

size

Specifies the value in either kbits, kilobytes or bytes.

Values kbits: 4 to 2146959 | default
 bytes: 512 to 274810752
 kbytes: 1to 268369

mode

Syntax

mode *mode*
no mode

Context

config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter
config>service>vprn>interface>sap>ingress>meter-override>meter (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command within the SAP ingress meter-overrides contexts is used to override the sap-ingress QoS policy configured mode parameters for the specified meter-id.
The **no** form of this command is used to restore the policy defined metering and profiling mode to a meter.

Parameters

mode

Specifies the rate mode of the meter-override.

Values trtcm1 | trtcm2 | srctm

rate

Syntax

rate *cir cir-rate* [*pir pir-rate*]
no rate

Context

```
config>service>epipe>sap>ingress>meter-override>meter
config>service>vpls>sap>ingress>meter-override>meter
config>service>ies>interface>sap>ingress>meter-override>meter
config>service>vprn>interface>sap>ingress>meter-override>meter (not supported in access-uplink
operating mode)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command within the SAP ingress meter-overrides contexts is used to override the sap-ingress QoS policy configured rate parameters for the specified meter-id.

The **no** form of this command is used to restore the policy defined metering and profiling rate to a meter.

Default

max

The **max** default specifies the amount of bandwidth in kilobits per second (thousand bits per second). The **max** value is mutually exclusive to the **pir-rate** value.

Parameters

pir-rate

Specifies the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed. Fractional values are not allowed and must be specified as a positive integer.



Note:

When the meter mode is set to 'trtcm2' the PIR value is interpreted as the EIR value. For more information, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Quality of Service Guide*.

The actual PIR rate is dependent on the queue **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values [0..20000000 | max]

Default max

cir-rate

The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed. Fractional values are not allowed and must be specified as a positive integer.

Values [0..20000000 | max]

Default 0

qos

Syntax

qos *policy-id*

qos *policy-id* [**enable-table-classification**] (for 7210 SAS-Mxp Epipe ingress only)

no qos *policy-id*

Context

config>service>epipe>sap>ingress

config>service>epipe>sap>egress (7210 SAS-Mxp with Epipe)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command associates a Quality of Service (QoS) policy with an ingress or egress SAP.

QoS ingress 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 only allows ingress policies to be associated on SAP or IP interface ingress, and only allows egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second policy of same or different type replaces the earlier one with the new policy.



Note:

SAP egress QoS policies are only supported on the 7210 SAS-Mxp.

On the 7210 SAS-Mxp (ingress), using the **enable-table-classification** keyword enables the use of IP DSCP tables to assign FC and profile on a per-SAP ingress basis. The match-criteria configured from the service ingress policy, which require CAM resources, are ignored. Only meters from the service ingress policy are used (and the meters still require CAM resources). The IP DSCP classification policy configured in the SAP ingress policy is used to assign FC and profile. The default FC is assigned from the SAP ingress policy.

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

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

Parameters

policy-id

Specifies the ingress or egress policy ID to associate with SAP on ingress or egress. The policy ID must already exist.

Values 1 to 65535

enable-table-classification

Enables the use of table-based classification instead of CAM-based classification at SAP ingress. The FC and profile are taken from the IP DSCP classification policy configured in the ingress policy, along with the meters from the SAP ingress policy. Match-criteria entries in the SAP ingress policy are ignored. This parameter is only supported on the 7210 SAS-Mxp.

3.9.2.1.5 Epipe Service SAP Statistics Commands

statistics

Syntax

statistics

Context

config>service>epipe>sap

Platforms

7210 SAS-T (both network and access uplink mode), 7210 SAS-Mxp, and 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC)

Description

Commands in this context configure the counters associated with SAP ingress and egress.

ingress

Syntax

ingress

Context

config>service>epipe>sap>statistics

Platforms

7210 SAS-T (both network and access uplink mode), 7210 SAS-Mxp, and 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC)

Description

Commands in this context configure the ingress SAP statistics counters.

counter-mode

Syntax

counter-mode {in-out-profile-count | forward-drop-count}

Context

config>service>epipe>sap>statistics>ingress

Platforms

7210 SAS-T (network and access-uplink mode), 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), and 7210 SAS-Sx 10/100GE (standalone)

Description

This command sets the counter mode for the counters associated with SAP ingress meters (also known as policers). A pair of counters is available with each meter. These counters count different events based on the counter mode value.



Note:

- The counter mode can be changed if an accounting policy is associated with a SAP. If the counter mode is changed, the counters associated with the meter are reset and the counts are cleared. If an accounting policy is in use when the counter mode is changed, a new record is written into the current accounting file.
- The configuration information is not saved across a reboot.

Execute the following sequence of commands on the specified SAP to ensure that the correct statistics are collected when the counter-mode is changed:

1. Execute the command **config>service>epipe>sap>no collect-stats**, to disable writing of accounting records for the SAP.
2. Change the counter-mode to the needed option, execute the command **config>service>epipe>sap>counter-mode {in-out-profile-count | forward-drop-count}**.
3. Execute the command **config>service>epipe>sap>collect-stats**, to enable writing of accounting records for the SAP.

Default

in-out-profile-count

Parameters

forward-drop-count

Counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets

dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

in-out-profile-count

Counts on one counter the total in-profile packets and octets received on ingress of a SAP, and on another, counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

drop-count-extra-vlan-tag-pkts

Syntax

[no] drop-count-extra-vlan-tag-pkts

Context

config>service>epipe>sap>statistics>ingress

config>service>epipe>spoke-sdp>statistics>ingress

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp, and 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC)

Description

This command associates a counter which enables the counting of extra VLAN-tag dropped packets for the SAP or spoke-SDP. A limited number of counters are available for use.

The **no** form of this command removes the associated counter.

3.9.2.1.6 VLL SDP Commands



Note:

VLL SDP commands are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode. Any exceptions are noted explicitly.

spoke-sdp

Syntax

spoke-sdp *sdp-id[:vc-id]* [**no-endpoint**] [**create**]

spoke-sdp *sdp-id[:vc-id]* **endpoint** *endpoint-name*

no spoke-sdp *sdp-id[:vc-id]*

Context

config>service>epipe

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command binds a service to an existing Service Distribution Point (SDP).

The SDP has an operational state that determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service is down.

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

SDPs must be explicitly associated and bound to a service to allow far-end 7210 SAS devices to participate in the service.

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

Default

no sdp-id is bound to a service

Special Cases

Epipe

At most, only one *sdp-id* can be bound to an Epipe service. Since an Epipe is a point-to-point service, it can have, at most, two end points. The two end points can be one SAP and one SDP or two SAPs. VC-switching VLLs are an exception. If the VLL is a "vc-switching" VLL, the two endpoints must both be SDPs.

Parameters

sdp-id

Specifies the SDP identifier.

Values 1 to 17407

vc-id

Specifies the virtual circuit identifier.

Values 1 to 4294967295

no-endpoint

Keyword to remove the association of a spoke-SDP with an explicit endpoint name.

endpoint-name

Specifies the name of the service endpoint.

create

Keyword to create spoke-SDP.

control-word**Syntax**

[no] control-word

Context

config>service>epipe>spoke-sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

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

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

The service only comes up if the same C-bit value is signaled in both directions. If a spoke-SDP is configured to use the control word but the node receives a label mapping message with a C-bit clear, the node releases the label with the an "Illegal C-bit" status code as per Section 6.1 of RFC 4447. As soon as the user also enabled the control the remote peer, the remote peer withdraws its original label and sends a label mapping with the C-bit set to 1 and the VLL service is up in both nodes.

control-channel-status**Syntax**

[no] control-channel-status

Context

config>service>epipe>spoke-sdp

Platforms

7210 SAS-T

Description

This command enables the configuration of static pseudowire status signaling on a spoke-SDP for which signaling for its SDP is set to OFF.

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

- The system is using network chassis mode D
- SDP signaling is off
- The control-word is enabled (control-word by default is disabled)
- The service type is Epipe or VPLS.
- Mate SDP signaling is off (in vc-switched services)
- pw-path-id is configured for this spoke

The **no** form of this command removes control channel status signaling from a spoke-SDP. It can only be removed if control channel status is shutdown.

Default

no control-channel-status

acknowledgment

Syntax

[no] **acknowledgment**

Context

config>service>epipe>spoke-sdp>control-channel-status

Platforms

7210 SAS-T

Description

This command enables the acknowledgment of control channel status messages. By default, no acknowledgment packets are sent.

refresh-timer

Syntax

refresh-timer *value*

no refresh-timer

Context

config>service>epipe>spoke-sdp>control-channel-status

Platforms

7210 SAS-T

Description

This command configures the refresh timer for control channel status signaling packets. By default, no refresh packets are sent.

Default

no refresh-timer

Parameters

value

Specifies the refresh timer value.

Values 10 to 65535 seconds

Default 0 (off)

request-timer

Syntax

request-timer **request-timer** *request-timer-secs* **retry-timer** *retry-timer-secs* **timeout-multiplier** *multiplier*

Context

config>service>epipe>spoke-sdp>control-channel-status

Platforms

7210 SAS-T

Description

This command configures the control channel status request mechanism. When it is configured, control channel status request procedures are used. These augment the procedures for control channel status messaging from RFC 6478. This command is mutually exclusive with a non-zero refresh-timer value.

Parameters

request-timer

Specifies the interval at which pseudowire status messages, including a reliable delivery TLV, with the "request" bit set, are sent.

Values 10 to 65535 seconds

retry-timer

Specifies the timeout interval if no response to a pseudowire status request is received. This parameter must be configured. A value of zero (0) disables retries.

Values 0, 3 to 60 seconds

timeout-multiplier

Specifies the optional timeout multiplier for the retry timer. If a requesting node does not receive a valid response to a pseudowire status request within this multiplier value times the retry timer value, it assumes the pseudowire is down.

Values 3 to 20 seconds

precedence**Syntax**

precedence [*precedence-value* | **primary**]

no precedence

Context

config>service>epipe>spoke-sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding begins to forward traffic.

The **no** form of this command returns the precedence value to the default.

Default

4

Parameters***precedence-value***

Specifies the spoke-SDP precedence.

Values 1 to 4

primary

Specifies to make this the primary spoke-SDP.

pw-path-id**Syntax**

[no] pw-path-id

Context

config>service>epipe>spoke-sdp

Platforms

7210 SAS-T

Description

This command configures an MPLS-TP Pseudowire Path Identifier for a spoke-SDP. All elements of the PW path ID must be configured to enable a spoke-SDP with a PW path ID.

For an IES or VPRN spoke-SDP, the pw-path-id is only valid for Ethernet spoke-SDPs.

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

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

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

Default

no pw-path-id

agi

Syntax

agi *agi*

no agi

Context

config>service>epipe>spoke-sdp>pw-path-id

Platforms

7210 SAS-T

Description

This command configures the attachment group identifier for an MPLS-TP PW.

Parameters

agi

Specifies the attachment group identifier.

Values 0 to 4294967295

saii-type2

Syntax

saii-type2 *global-id:node-id:ac-id*

no saii-type2

Context

config>service>epipe>spoke-sdp>pw-path-id

Platforms

7210 SAS-T

Description

This command configures the source individual attachment identifier (SAII) for an MPLS-TP spoke-SDP. If this is configured on a spoke-SDP for which vc-switching is also configured, that is, if it is at an S-PE, the values must match those of the taii-type2 of the mate spoke-SDP.

Parameters

global-id

Specifies the global ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values 0 to 4294967295

node-id

Specifies the node ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values a.b.c.d or 0 to 4294967295

ac-id

Specifies the attachment circuit ID at the source PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, the AC ID must be set to a locally unique value.

Values 1 to 4294967295

taii-type2

Syntax

taii-type2 *global-id:node-id:ac-id*

no taii-type2

Context

config>service>epipe>spoke-sdp>pw-path-id

Platforms

7210 SAS-T

Description

This command configures the source individual attachment identifier (SAII) for an MPLS-TP spoke-SDP. If this is configured on a spoke-SDP for which vc-switching is also configured (it is at an S-PE), the values must match those of the taii-type2 of the mate spoke-SDP.

Parameters

global-id

Specifies the global ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values 0 to 4294967295

node-id

Specifies the node ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP.

Values a.b.c.d or 0 to 4294967295

ac-id

Specifies the attachment circuit ID at the target PE or T-PE for the MPLS-TP PW for a spoke-SDP. If this node is the source of the PW, the AC ID must be set to a locally unique value.

Values 1 to 4294967295

pw-status-signaling

Syntax

[no] pw-status-signaling

Context

config>service>epipe>spoke-sdp

Platforms

7210 SAS-T

Description

This command enables pseudowire status signaling for this spoke-SDP binding.

The **no** form of this command disables the status signaling.

Default

pw-status-signaling

vc-label

Syntax

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

Context

config>service>epipe>spoke-sdp>egress

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command configures the egress VC label.

Parameters

vc-label

A VC egress value that indicates a specific connection.

Values 16 to 1048575

vc-label

Syntax

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

Context

config>service>epipe>spoke-sdp>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command configures the ingress VC label.

Parameters

vc-label

Specifies a VC ingress label value for a 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

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

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.

spoke-sdp-fec

Syntax

spoke-sdp-fec

spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**]

spoke-sdp-fec *spoke-sdp-fec-id* **no-endpoint**

spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**] **endpoint** *name* [**icb**]

Context

config>service>epipe

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command binds a service to an existing SDP, using a dynamic MS-PW.

A spoke-SDP is treated like the equivalent of a traditional bridge "port" where flooded traffic received on the spoke-SDP is replicated on all other "ports" (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service is down.

When using dynamic MS-PWs, the particular SDP to bind to is automatically selected based on the Target Attachment Individual Identifier (TAII) and the path to use, specified under spoke-SDP FEC. The selected SDP terminates on the first hop S-PE of the MS-PW. Therefore, an SDP must already be defined in the **configure service sdp** context that reaches the first hop 7210 SAS of the MS-PW. The 7210 SAS associates an SDP with a service. If an SDP is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that sdp-id and the service is created.

It differs from the **spoke-sdp** command in that the **spoke-sdp** command creates a spoke-SDP binding that uses a PW with the PW ID FEC. However, the **spoke-sdp-fec** command enables PWs with other FEC types to be used. In Release 9.0, only the Generalized ID FEC (FEC129) may be specified using this command.

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

Parameters

spoke-sdp-fec-id

Specifies an unsigned integer value of the spoke-SDP.

Values 1 to 4294967295

fec fec-type

Specifies an unsigned integer value for the type of the FEC used by the MS-PW.

Values 129 to 130

aii-type aii-type

Specifies an unsigned integer value for the Attachment Individual Identifier (AII) type used to identify the MS-PW endpoints.

Values 1 to 2

endpoint endpoint-name

Specifies the name of the service endpoint.

no endpoint

Adds or removes a spoke-SDP association.

icb

Configures the spoke-SDP as an inter-chassis backup SDP binding.

auto-config**Syntax**

[no] auto-config

Context

config>service>epipe>spoke-sdp-fec

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command enables single-sided automatic endpoint configuration of the spoke-SDP. The 7210 SAS acts as the passive T-PE for signaling this MS-PW.

Automatic Endpoint Configuration allows the configuration of a spoke-SDP endpoint without specifying the TAIL associated with that spoke-SDP. It allows a single-sided provisioning model where an incoming label mapping message with a TAIL that matches the SAIL of that spoke-SDP to be automatically bound to that endpoint. In this mode, the far end T-PE actively initiates MS-PW signaling and sends the initial label mapping message using T-LDP, while the 7210 T-PE for which auto-config is specified, acts as the passive T-PE.

The **auto-config** command is blocked in CLI if signaling active has been enabled for this spoke-SDP. It is only applicable to spoke-SDPs configured under the Epipe, IES and VPRN interface context.

The **no** form of this command means that the 7210 T-PE either acts as the active T-PE (if signaling active is configured) or automatically determines which 7210 initiates MS-PW signaling based on the prefix values configured in the SAIL and TAIL of the spoke-SDP. If the SAIL has the greater prefix value, the 7210 initiates MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAIL has the greater value prefix, then the 7210 assumes that the far end T-PE initiates MS-PW signaling and waits for that label mapping message before responding with a T-LDP label mapping message for the MS-PW in the reverse direction.

Default

no auto-config

path**Syntax**

path *name*

no path

Context

config>service>epipe>spoke-sdp-fec

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command specifies the explicit path, containing a list of S-PE hops, that should be used for this spoke-SDP. The path-name should correspond to the name of an explicit path configured in the **config>service>pw-routing** context.

If no path is configured, each next-hop of the MS-PW used by the spoke-SDP is chosen locally at each T-PE and S-PE.

Default

no path

Parameters

path-name

The name of the explicit path to be used, as configured under config>service>pw-routing.

precedence

Syntax

precedence *precedence-value*

precedence primary

no precedence

Context

config>service>epipe>spoke-sdp-fec

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command specifies the precedence of the SDP binding when there are multiple SDP bindings attached to one service endpoint. The value of zero can only be assigned to one SDP bind making it the primary SDP bind. When an SDP binding goes down, the next highest precedence SDP binding begins to forward traffic.

The **no** form of this command returns the precedence value to the default.

Default

42

Parameters***precedence-value***

Specifies the spoke-SDP precedence.

Values 1 to 4***primary***

Specifies to make this the primary spoke-SDP.

pw-template-bind**Syntax****pw-template-bind** *policy-id***no pw-template-bind****Context**

config>service>epipe>spoke-sdp-fec

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command binds the parameters included in a specific PW Template to a spoke-SDP.

The **no** form of this command removes the values from the configuration.**Parameters*****policy-id***

Specifies the existing policy ID.

Values 1 to 2147483647**retry-count****Syntax****retry-count** *retry-count***no retry-count**

Context

config>service>epipe>spoke-sdp-fec

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This optional command specifies the number of attempts software should make to reestablish the spoke-SDP after it has failed. After each successful attempt, the counter is reset to zero.

When the specified number is reached, no more attempts are made and the spoke-SDP is put into the shutdown state.

Use the **no shutdown** command to bring up the path after the retry limit is exceeded.

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

Default

30

Parameters

retry-count

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

Values 10 to 10000

retry-timer

Syntax

retry-timer *retry-timer*

no retry-timer

Context

config>service>epipe>spoke-sdp-fec

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command specifies a retry-timer for the spoke-SDP. This is a configurable exponential back-off timer that determines the interval between retries to reestablish a spoke-SDP if it fails and a label withdraw message is received with the status code "All unreachable".

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

Default

30

Parameters***retry-timer***

Specifies the initial retry-timer value in seconds.

Values 10 to 480**saii-type2****Syntax****saii-type2** *global-id:prefix:ac-id***no saii-type2****Context**

config>service>epipe>spoke-sdp-fec

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command configures the source attachment individual identifier for the spoke-SDP. This is only applicable to FEC129 All type 2.

Parameters***global-id***Specifies the global ID of this 7210 T-PE. This value must correspond to one of the `global_id` values configured for a local-prefix under **config>service>pw-routing>local-prefix** context.**Values** 1 to 4294967295***prefix***Specifies the prefix on this 7210 T-PE that the spoke-SDP SDP is associated with. This value must correspond to one of the prefixes configured under **config>service>pw-routing>local-prefix** context.**Values** an IPv4-formatted address a.b.c.d or 1 to 4294967295***ac-id***

Specifies an unsigned integer representing a locally unique identifier for the spoke-SDP.

Values 1 to 4294967295

taii-type2

Syntax

taii-type2 *global-id:prefix:ac-id*
no taii-type2

Context

config>service>epipe>spoke-sdp-fec

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command configures the target attachment individual identifier for the spoke-SDP. This is only applicable to FEC129 All type 2.

This command is blocked in CLI if this end of the spoke-SDP is configured for single-sided auto configuration (using the **auto-config** command).

Parameters

global-id

Specifies the Global ID of this 7210 T-PE. This value must correspond to one of the `global_id` values configured for a local-prefix under **config>service>pw-routing>local-prefix** context.

Values 1 to 4294967295

prefix

Specifies the prefix on this 7210 T-PE that the spoke-SDP is associated with. This value must correspond to one of the prefixes configured under **config>service>pw-routing>local-prefix** context.

Values an IPv4-formatted address a.b.c.d or 1 to 4294967295

ac-id

Specifies an unsigned integer representing a locally unique identifier for the spoke-SDP.

Values 1 to 4294967295

3.9.2.1.7 Connection profile commands

connection-profile

Syntax

connection-profile *conn-prof-id* [create]

no connection-profile *conn-prof-id*

Context

config

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates a list of VLAN values to be assigned to a Dot1q SAP in an Epipe service.

A connection profile can only be assigned to a Dot1q SAP which is part of an Epipe service.

The **no** form of this command deletes the profile from the configuration.

Parameters

conn-prof-id

Specifies the profile number.

Values 1 to 8000

ethernet

Syntax

ethernet

Context

config>connprof

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command provides the context to configure the VLAN ranges values.

ranges

Syntax

no ranges
ranges *vlan-ranges* [*vlan-ranges...*(up to 32 max)]

Context

config>connprof>ethernet

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Specifies the list of VLAN ranges or individual VLAN ID to be used for mapping the specific VLANs to the Epipe SAP.

The system validates that the values specified are valid VLAN ID in the range 0-4094 (VLAN ID 4095 is reserved). Ranges are specified in the format 'a-b ', the expression (a < b) should be true. Up to about 32 individual VLAN values or VLAN ranges can be specified. A maximum of up to 8 VLAN ranges are allowed per connection profile.

Parameters

vlan-ranges

Specifies the list of VLAN ranges or individual VLAN ID to be used for mapping the specific VLANs to the Epipe SAP.

A list of space separated values specified as either a-b or individual VLAN IDs. Both the VLAN IDs and the value used for 'a' and 'b' must be in the range of 0-4094. Additionally, value 'a' must be less than value 'b'.

For example:

ranges	100-200 5 6 4000-4020
ranges	4 5 6 10 11 12
ranges	250-350 500-600 1000-1023

3.9.2.2 Show commands

```
id
```

Syntax

```
id service-id {all | arp | base | endpoint | fdb | interface | label | labels | sap | split-horizon-group | stp|
interface | mstp-configuration}
```

Context

```
show>service
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays information for a particular service-id.

Parameters

service-id
Display the service identification number that identifies the service in the domain.

Values service-id: 1 to 214748364
 svc-name: A string up to 64 characters.

all
Displays more information about the service.

arp
Displays ARP entries for the service.

base
Displays basic service information.

endpoint
Displays service endpoint information.

fdb
Displays FDB information.

interface
Displays service interfaces.

labels
Displays labels being used by this service.

mstp-configuration
Displays MSTP information.

sap

Displays SAPs associated with the service.

sdp

Displays SDPs associated with the service. This keyword is not supported on 7210 SAS platforms operating in access-uplink mode.

split-horizon-group

Displays split horizon group information.

stp

Displays STP information.

Output

The following output is an example of information for a specific service ID.

Sample output

```
*A:ces-A# show service id 1 sap
=====
SAP(Summary), Service 1
=====
PortId                SvcId      Ing.   Ing.   Egr.   Adm   Opr
                    QoS      Fltr   Fltr
-----
1/2/1.1                1          1    none  none   Up    Up
-----
Number of SAPs : 1
=====

*A:Dut-A# /show service id 1 base
=====
Service Basic Information
=====
Service Id       : 1          Vpn Id       : 0
Service Type    : VPLS
Name             : (Not Specified)
Description      : (Not Specified)
Customer Id     : 1          Creation Origin : manual
Last Status Change: 03/20/2023 12:49:45
Last Mgmt Change : 03/20/2023 12:49:45
Admin State     : Up        Oper State    : Up
MTU              : 1514     Def. Mesh VC Id : 1
MTU Check       : Enabled
SAP Count       : 1        SDP Bind Count : 1
Snd Flush on Fail : Disabled  Host Conn Verify : Disabled
SAP Type        : Any
Propagate MacFlush: Disabled  Per Svc Hashing : Disabled
Allow IP Intf Bind: Disabled  Fwd-IPv4-Mcast-To*: Disabled
VSD Domain      : <none>
SPI load-balance : Disabled
TEID load-balance : Disabled

-----
Service Access & Destination Points
-----
Identifier                Type      AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/21                null      1514    1514    Up    Up
```

```
sdp:1:1 S(2.2.2.2)           Spok           0           9190      Up      Up
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-A#
```

all

Syntax

all

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays detailed information for all aspects of the service.

Output

The following outputs are examples of detailed service information, and the associated tables describe the output fields.

- [Sample output — Standard](#), [Sample output \(PW-Entropy/Hash-label\)](#), [Table 27: Output fields: service ID all](#)
- [Sample output \(Meter-override\)](#)

Sample output — Standard

```
*A:Dut-A>show>service>id# all
=====
Service Detailed Information
=====
Service Id       : 1501           Vpn Id           : 1501
Service Type     : Epipe
Description      : Default epipe description for service id 1501
Customer Id      : 1
Last Status Change: 02/21/2011 13:07:03
Last Mgmt Change  : 02/21/2011 13:03:58
Admin State      : Up             Oper State        : Up
MTU              : 1514
MTU Check        : Enabled
Vc Switching     : False
SAP Count        : 1             SDP Bind Count    : 2
-----
Service Destination Points(SDPs)
-----
Sdp Id 1413:1501  -(10.20.1.4)
-----
```

```

Description      : Default sdp description
SDP Id          : 1413:1501                Type           : Spoke
VC Type         : Ether                    VC Tag          : n/a
Admin Path MTU  : 0                       Oper Path MTU    : 9182
Far End         : 10.20.1.4                Delivery        : MPLS

Admin State     : Up                      Oper State       : Up
Acct. Pol      : 14                      Collect Stats    : Enabled
Ingress Label   : 130948                  Egress Label     : 130483
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 : 02/21/2011 13:07:12 Signaling        : TLDP
Last Mgmt Change : 02/21/2011 13:03:58    Force Vlan-Vc    : Disabled
Endpoint        : coreSide                Precedence       : 1
Class Fwding State : Down
Flags           : None
Peer Pw Bits    : None
Peer Fault Ip   : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : pwe3ControlWord mplsRouterAlertLabel

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

Statistics      :
I. Fwd. Pkts.   : 48319                    I. Fwd. Octs.    : 5690869
E. Fwd. Pkts.   : 34747                    E. Fwd. Octets   : 4013709
-----
Eth-Cfm Configuration Information
-----
Md-index        : 1000                     Direction        : Down
Ma-index        : 1150114                  Admin            : Enabled
MepId           : 1                       CCM-Enable       : Enabled
LowestDefectPri : macRemErrXcon            HighestDefect     : none
Defect Flags    : None
Mac Address     : 7c:20:64:ad:04:07        ControlMep       : False
CcmLtmPriority  : 7
CcmTx           : 11385                    CcmSequenceErr   : 0
Eth-1Dm Threshold : 3(sec)
Eth-Ais        : Disabled
Eth-Tst        : Disabled
LbRxReply      : 0                         LbRxBadOrder     : 0
LbRxBadMsdu    : 0                         LbTxReply        : 0
LbNextSequence : 1                         LtNextSequence    : 1
LtRxUnexplained : 0

Associated LSP LIST :
Lsp Name        : A_D_21
Admin State     : Up                      Oper State       : Up
Time Since Last Tr*: 03h49m30s

-----
Sdp Id 1613:1501  -(10.20.1.6)
-----
Description      : Default sdp description
SDP Id          : 1613:1501                Type           : Spoke
VC Type         : Ether                    VC Tag          : n/a
Admin Path MTU  : 0                       Oper Path MTU    : 9182
Far End         : 10.20.1.6                Delivery        : MPLS

```

```

Admin State      : Up
Acct. Pol       : 14
Ingress Label    : 130526
Ing mac Fltr     : n/a
Ing ip Fltr      : n/a
Admin ControlWord : Not Preferred
Admin BW(Kbps)   : 0
Last Status Change : 02/21/2011 13:07:03
Last Mgmt Change  : 02/21/2011 13:03:58
Endpoint         : coreSide
Class Fwding State : Down
Flags           : None
Peer Pw Bits     : pwFwdingStandby
Peer Fault Ip    : None
Peer Vccv CV Bits : lspPing
Peer Vccv CC Bits : mplsRouterAlertLabel

Oper State      : Up
Collect Stats   : Enabled
Egress Label    : 130424
Egr mac Fltr    : n/a
Egr ip Fltr     : n/a
Oper ControlWord : False
Oper BW(Kbps)   : 0
Signaling       : TLDP
Force Vlan-Vc   : Disabled
Precedence      : 2

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

Statistics       :
I. Fwd. Pkts.    : 25
E. Fwd. Pkts.    : 23
I. Fwd. Octs.    : 2776
E. Fwd. Octets   : 2557

-----
Eth-Cfm Configuration Information
-----
Md-index        : 1000
Ma-index        : 1150116
MepId           : 1
LowestDefectPri : macRemErrXcon
Defect Flags    : None
Mac Address     : 7c:20:64:ad:04:07
CcmLtmPriority   : 7
CcmTx           : 11414
Eth-1Dm Threshold : 3(sec)
Eth-Ais         : Disabled
Eth-Tst         : Disabled
LbRxReply       : 0
LbRxBadMsdu     : 0
LbNextSequence  : 1
LtRxUnexplained : 0

Direction       : Down
Admin           : Enabled
CCM-Enable      : Enabled
HighestDefect    : none
ControlMep      : False
CcmSequenceErr  : 0

LbRxBadOrder    : 0
LbTxReply       : 0
LtNextSequence  : 1

Associated LSP LIST :
Lsp Name        : A_F_21
Admin State     : Up
Time Since Last Tr*: 03h48m45s
Oper State      : Up

-----
Number of SDPs : 2
-----

Service Access Points
-----

SAP lag-3:1501.1501
-----
Service Id      : 1501
SAP             : lag-3:1501.1501
QinQ Dot1p     : Default
Description     : (Not Specified)
Admin State     : Up
Encap           : qinq
Oper State      : Up

```

```

Flags          : None
Last Status Change : 02/21/2011 13:06:45
Last Mgmt Change  : 02/21/2011 13:03:58

Admin MTU       : 9212                Oper MTU       : 9212
Ingr IP Fltr-Id : n/a                 Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : 1501               Egr Mac Fltr-Id : n/a
tod-suite       : None
Egr Agg Rate Limit : max
Endpoint        : accessSide

Acct. Pol       : Default              Collect Stats    : Enabled

```

QOS

```

Ingress qos-policy : 1500                Egress qos-policy : 1500

```

Sap Egress Policy (1500)

```

Scope          : Template
Remark         : False                Remark Pol Id    : 2
Accounting      : frame-based
Description     : Sap Egress Policy for svcList 1500

```

Queue Rates and Rules

QueueId	CIR	CIR Adpt Rule	PIR	PIR Adpt Rule
Queue1	10000	max	10000	max
Queue2	10000	max	10000	max
Queue3	10000	max	10000	max
Queue4	10000	max	10000	max
Queue5	10000	max	10000	max
Queue6	10000	max	10000	max
Queue7	10000	max	10000	max
Queue8	10000	max	10000	max

Parent Details

QueueId	Port	CIR Level	PIR Weight
Queue1	True	1	1
Queue2	True	2	2
Queue3	True	3	3
Queue4	True	4	4
Queue5	True	5	5
Queue6	True	6	6
Queue7	True	7	7
Queue8	True	8	8

High Slope

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)
Queue1	Up	50	100	50
Queue2	Up	50	100	50
Queue3	Up	50	100	50
Queue4	Up	50	100	50
Queue5	Up	50	100	50

Queue6	Up	50	100	50
Queue7	Up	50	100	50
Queue8	Up	50	100	50

Low Slope				

QueueId	State	Start-Avg(%)	Max-Avg(%)	Max-Prob(%)

Queue1	Up	10	50	50
Queue2	Up	10	50	50
Queue3	Up	10	50	50
Queue4	Up	10	50	50
Queue5	Up	10	50	50
Queue6	Up	10	50	50
Queue7	Up	10	50	50
Queue8	Up	10	50	50

Burst Sizes and Time Average Factor				

QueueId	CBS	MBS	Time Average Factor	Queue-Mgmt

Queue1	200	400	10	qM_1500
Queue2	200	400	10	qM_1500
Queue3	200	400	10	qM_1500
Queue4	200	400	10	qM_1500
Queue5	200	400	10	qM_1500
Queue6	200	400	10	qM_1500
Queue7	200	400	10	qM_1500
Queue8	200	400	10	qM_1500

Aggregate Policer (Available)				

rate	: n/a		burst	: n/a

Ingress QoS Classifier Usage				

Classifiers Allocated: 32		Meters Allocated : 16		
Classifiers Used : 8		Meters Used : 5		

Sap Statistics				

Ingress Stats:	Packets		Octets	
Egress Stats:	34659		3241035	
Extra-Tag Drop Stats:	48099		5291928	
	n/a		n/a	

Sap per Meter stats				

	Packets		Octets	

Ingress Meter 1 (Unicast)				
For. InProf	: 7209		468585	
For. OutProf	: 0		0	

Ingress Meter 2 (Unicast)				
For. InProf	: 0		0	
For. OutProf	: 0		0	

Ingress Meter 3 (Unicast)				
For. InProf	: 0		0	
For. OutProf	: 0		0	


```

Ingress Meter 4 (Unicast)
For. InProf      : 0           0
For. OutProf     : 0           0

```

```

Ingress Meter 5 (Unicast)
For. InProf      : 27454       2772854
For. OutProf     : 0           0

```

Sap per Queue stats

	Packets	Octets
Egress Queue 1 (be)		
Fwd Stats	: 0	0
Drop InProf	: 0	0
Drop OutProf	: 0	0
Egress Queue 2 (l2)		
Fwd Stats	: 3	180
Drop InProf	: 0	0
Drop OutProf	: 0	0
Egress Queue 3 (af)		
Fwd Stats	: 0	0
Drop InProf	: 0	0
Drop OutProf	: 0	0
Egress Queue 4 (l1)		
Fwd Stats	: 0	0
Drop InProf	: 0	0
Drop OutProf	: 0	0
Egress Queue 5 (h2)		
Fwd Stats	: 0	0
Drop InProf	: 0	0
Drop OutProf	: 0	0
Egress Queue 6 (ef)		
Fwd Stats	: 0	0
Drop InProf	: 0	0
Drop OutProf	: 0	0
Egress Queue 7 (h1)		
Fwd Stats	: 0	0
Drop InProf	: 0	0
Drop OutProf	: 0	0
Egress Queue 8 (nc)		
Fwd Stats	: 20842	1938306
Drop InProf	: 0	0
Drop OutProf	: 0	0

Service Endpoints

```

Endpoint name      : coreSide
Description        : (Not Specified)
Revert time        : 0
Act Hold Delay     : 0
Standby Signaling Master : true
Tx Active          : 1413:1501
Tx Active Up Time  : 0d 03:48:41
Revert Time Count Down : N/A
Tx Active Change Count : 2

```

```

Last Tx Active Change      : 02/21/2011 13:07:12
-----
Members
-----
Spoke-sdp: 1413:1501 Prec:1                      Oper Status: Up
Spoke-sdp: 1613:1501 Prec:2                      Oper Status: Up
=====
Endpoint name              : accessSide
Description                 : (Not Specified)
Revert time                 : 0
Act Hold Delay              : 0
Standby Signaling Master   : false
Tx Active                   : lag-3:1501.1501
Tx Active Up Time           : 0d 03:49:08
Revert Time Count Down     : N/A
Tx Active Change Count     : 1
Last Tx Active Change      : 02/21/2011 13:06:45
-----
Members
-----
SAP      : lag-3:1501.1501                      Oper Status: Up
=====
=====

```

Sample output (PW-Entropy/Hash-label)

```

*A:7210SAS>config>service# /show service id 1 all

=====
Service Detailed Information
=====
Service Id      : 1                      Vpn Id      : 0
Service Type    : VPLS
Name            : (Not Specified)
Description     : (Not Specified)
Customer Id     : 1
Last Status Change: 01/07/2000 21:19:14
Last Mgmt Change : 01/07/2000 21:15:25
Admin State     : Up                    Oper State   : Up
MTU             : 1514                  Def. Mesh VC Id : 1
MTU Check       : Enabled
SAP Count       : 0                    SDP Bind Count : 1
Snd Flush on Fail : Disabled            Host Conn Verify : Disabled
SAP Type        : Any
Propagate MacFlush: Disabled            Per Svc Hashing : Disabled
Allow IP Intf Bind: Disabled

-----
Split Horizon Group specifics
-----

-----
ETH-CFM service specifics
-----
Tunnel Faults    : ignore                V-Mep Extensions : Enabled

-----
Service Destination Points(SDPs)
-----
Sdp Id 1:1      : (2.2.2.2)
-----

```

```

Description      : (Not Specified)
SDP Id           : 1:1
Spoke Descr      : (Not Specified)
Split Horiz Grp  : (Not Specified)
VC Type          : Ether
Admin Path MTU   : 0
Far End          : 2.2.2.2
Tunnel Far End   : 2.2.2.2
Hash Label       : Enabled
Oper Hash Label  : Enabled

Type             : Spoke
VC Tag           : n/a
Oper Path MTU    : 9190
Delivery         : MPLS
LSP Types        : LDP
Hash Lbl Sig Cap : Disabled

Admin State      : Up
Acct. Pol        : None
Ingress Label    : 131069
Ingr Mac Fltr-Id : n/a
Ingr IP Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a
Admin ControlWord : Not Preferred
Last Status Change : 01/07/2000 21:19:14
Last Mgmt Change  : 01/07/2000 21:15:25
Endpoint         : N/A
PW Status Sig     : Enabled
Class Fwding State : Down
Flags            : None
Local Pw Bits     : None
Peer Pw Bits      : None
Peer Fault Ip     : None

Oper State       : Up
Collect Stats    : Disabled
Egress Label     : 131069
Egr Mac Fltr-Id  : n/a
Egr IP Fltr-Id   : n/a
Egr IPv6 Fltr-Id : n/a
Oper ControlWord : False
Signaling        : TLDP
Force Vlan-Vc    : Disabled
Precedence       : 4

Application Profile: None
Transit Policy    : None
Max Nbr of MAC Addr: No Limit
Learned MAC Addr  : 0

Total MAC Addr    : 0
Static MAC Addr   : 0

MAC Learning      : Enabled
MAC Aging         : Enabled
BPDU Translation  : Disabled
L2PT Termination  : Disabled
MAC Pinning       : Disabled
Ignore Standby Sig : False
Oper Group        : (none)
Rest Prot Src Mac : Disabled
Auto Learn Mac Prot: Disabled

Discard Unkwn Srce: Disabled

Block On Mesh Fail: False
Monitor Oper Grp   : (none)

RestProtSrcMacAct : Disable

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

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

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

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

Statistics          :
I. Fwd. Pkts.       : 0
E. Fwd. Pkts.       : 0
Extra-Tag-Drop-Pkts: n/a

I. Fwd. Octs.       : 0
E. Fwd. Octets      : 0
Extra-Tag-Drop-Oc*: n/a

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

Refresh Timer       : <none>

```

```

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

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

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

Stp Service Destination Point specifics
-----
Stp Admin State      : Up                Stp Oper State      : Down
Core Connectivity    : Down
Port Role            : N/A                Port State          : Forwarding
Port Number          : 0                  Port Priority        : 128
Port Path Cost       : 10                 Auto Edge           : Enabled
Admin Edge           : Disabled            Oper Edge           : N/A
Link Type            : Pt-pt              BPDU Encap          : Dot1d
Root Guard           : Disabled            Active Protocol      : N/A
Last BPDU from       : N/A                Designated 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
TC bit BPDUs rcvd    : 0                  TC bit BPDUs tx      : 0
RST BPDUs rcvd       : 0                  RST BPDUs tx         : 0
-----

Number of SDPs : 1
-----
* indicates that the corresponding row element may have been truncated.
-----

Service Access Points
-----
No Sap Associations
-----

VPLS Spanning Tree Information
-----
VPLS oper state      : Up                Core Connectivity    : Down
Stp Admin State      : Down                Stp Oper State      : Down
Mode                 : Rstp                Vcp Active Prot.    : N/A

Bridge Id            : 80:00:c4:08:4a:59:b2:61 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

Root Bridge          : N/A
Primary Bridge       : N/A

Root Path Cost       : 0                    Root Forward Delay: 0
Rcvd Hello Time     : 0                    Root Max Age         : 0

```

```

Root Priority      : 0                               Root Port      : N/A
-----
Forwarding Database specifics
-----
Service Id       : 1                               Mac Move       : Disabled
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
Remote Age       : 900                             Local Age      : 300
High Watermark   : 95%                             Low Watermark  : 90%
Mac Learning     : Enabled                         Discard Unknown : Disabled
Mac Aging        : Enabled                         Relearn Only   : False
-----
IGMP Snooping Base info
-----
Admin State : Down
Querier     : No querier found
-----
Sap/Sdp      Oper MRtr Send Max MVR      Num
Id           State Port Qries Grps From-VPLS Grps
-----
sdp:1:1      Up    No   No   None N/A      0
-----
Service Endpoints
-----
No Endpoints found.
-----
=====
VPLS Sites
=====
Site           Site-Id  Dest           Mesh-SDP  Admin  Oper  Fwdr
-----

```

Table 27: Output fields: service ID all

Label	Description
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.
Service Type	The type of service.
VLL Type	The VLL type.
SDP Id	The SDP identifier.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change.

Label	Description
Endpoint	The name of the service endpoint.
Flags	The conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, L2OperDown, RelearnLimit Exceeded, RxProtSrcMac, ParentIfAdminDown, TodResource Unavail, TodMssResourceUnavail, SapParamMismatch, Sap IngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
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.
Admin Path MTU	The needed largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	The type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Jitter Buffer (packets)	The jitter buffer length in number of packet buffers.
Playout Threshold (packets)	The playout buffer packets threshold in number of packet buffers.
Playout Threshold (packets)	The current packet depth of the jitter buffer.
Peer Pw Bits	The bits set by the LDP peer when there is a fault on its side of the pseudowire. LAC failures occur on the SAP that has been configured on the pipe service, PSN bits are set by SDP-binding failures on the pipe service. The pwNotForwarding bit is set when none of the preceding failures apply, such as an MTU mismatch failure. This value is only applicable if the peer is using the pseudowire status signalling method to indicate faults. pwNotForwarding — Pseudowire not forwarding lacIngressFault Local — Attachment circuit RX fault

Label	Description
	lacEgressFault Local — Attachment circuit TX fault psnIngressFault Local — PSN-facing PW RX fault psnEgressFault Local — PSN-facing PW TX fault pwFwdingStandby — Pseudowire in standby mode
LLF Admin State	The Link Loss Forwarding administrative state.
LLF Oper State	The Link Loss Forwarding operational state.
Standby Signaling Master	If the parameter standby signaling master is enabled.
Hash Label	If use of PW hash label is enabled or not.
Oper Hash Label	If MPLS packet originated by the node, is using PW Hash label if the value displayed is "Enabled". If the value displayed is "Disabled", the MPLS packets originated by the node is not using Pseudowire Hash label.
Hash Lbl Sig Cap	If PW hash label signaling is enabled or not.

Sample output (Meter-override)

```

A:7210SAS>show>service# id 1101 sap 1/2/1:1 detail
Ingress Meter Override
-----
Meter Id           : 1
Admin PIR          : 12000                Admin CIR          : 10000
Oper PIR           : 12000                Oper CIR           : 10000
PIR Rule           : closest*             CIR Rule           : closest*
MBS                : 20 KBytes CBS : 15 Kbytes
Mode               : TrtcM2*

* means the value is inherited
-----
A:7210SAS>show>service#

```

base

Syntax

base

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Displays basic information about the service ID including service type, description, SAPs.

Output

The following output is an example of basic service information, and [Table 28: Output fields: base](#) describes the output fields.

Sample output

```
A:Dut-A# show service id 1101 base
=====
Service Basic Information
=====
Service Id       : 1101           Vpn Id           : 1101
Service Type     : Epipe
Description      : Default epipe description for service id 1101
Customer Id      : 1
Last Status Change: 07/07/2009 18:13:43
Last Mgmt Change  : 07/07/2009 14:39:14
Admin State      : Up             Oper State        : Up
MTU              : 1514
Vc Switching     : False
SAP Count        : 1             SDP Bind Count    : 1
-----
Service Access & Destination Points
-----
Identifier                Type      AdmMTU  OprMTU  Adm  Opr
-----
sap:lag-4:1101             q-tag    9212   9212    Up   Up
sdp:1409:1101 S(10.20.1.4) n/a       0      9186    Up   Up
=====
A:Dut-A#
```

Table 28: Output fields: base

Label	Description
Service Id	The service identifier.
Vpn Id	Specifies the VPN ID assigned to the service.
Service Type	The type of service: Epipe, VPLS
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Adm	The desired state of the service.

Label	Description
Oper	The operating state of the service.
Mtu	The largest frame size (in octets) that the service can handle.
Def. Mesh VC Id	This object is only valid in services that accept mesh SDP bindings. It is used to validate the VC ID portion of each mesh SDP binding defined in the service.
SAP Count	The number of SAPs defined on the service.
SDP Bind Count	The number of SDPs bound to the service.
Identifier	Specifies the service access (SAP) points.
Type	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received.
AdmMTU	Specifies the desired largest service frame size (in octets) that can be transmitted through this SAP, without requiring the packet to be fragmented.
PBB Tunnel Point	Specifies the endpoint in the B-VPLS environment where the Epipe terminates.
Admin MTU	Specifies the B-VPLS admin MTU.
Backbone-Flooding	Specifies whether or not the traffic is flooded in the B-VPLS for the destination instead of unicast. If the backbone destination MAC is in the B-VPLS FDB, it is unicast.
ISID	The 24 bit field carrying the service instance identifier associated with the frame. It is used at the destination PE as a demultiplexor field.

endpoint

Syntax

endpoint [*endpoint-name*]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays service endpoint information.

Parameters

endpoint-name

Specifies the name of an existing endpoint for the service.

Output

The following output is an example of service endpoint information, and [Table 29: Output fields: service ID endpoint](#) describes the output fields.

Sample output

```
*A:Dut-A>show>service>id# endpoint

=====
Service 1501 endpoints
=====
Endpoint name      : coreSide
Description        : (Not Specified)
Revert time       : 0
Act Hold Delay    : 0
Standby Signaling Master : true
Tx Active         : 1413:1501
Tx Active Up Time : 0d 03:46:25
Revert Time Count Down : N/A
Tx Active Change Count : 2
Last Tx Active Change : 02/21/2011 13:07:12
-----
Members
-----
Spoke-sdp: 1413:1501 Prec:1           Oper Status: Up
Spoke-sdp: 1613:1501 Prec:2           Oper Status: Up
=====
Endpoint name      : accessSide
Description        : (Not Specified)
Revert time       : 0
Act Hold Delay    : 0
Standby Signaling Master : false
Tx Active         : lag-3:1501.1501
Tx Active Up Time : 0d 03:46:52
Revert Time Count Down : N/A
Tx Active Change Count : 1
Last Tx Active Change : 02/21/2011 13:06:45
-----
Members
-----
SAP      : lag-3:1501.1501           Oper Status: Up
=====
```

Table 29: Output fields: service ID endpoint

Label	Description
Service endpoints	

Label	Description
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

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays the labels being used by the service.

Output

The following output is an example of service label information, and [Table 30: Output fields: labels](#) describes the output fields.

Sample output

```
*A:ALA-12# show service id 1 labels
=====
Martini Service Labels
=====
Svc Id      Sdp Id      Type I.Lbl      E.Lbl
-----
1           10:1        Mesh 0         0
1           20:1        Mesh 0         0
1           30:1        Mesh 0         0
1           40:1        Mesh 130081    131061
1           60:1        Mesh 131019    131016
1           100:1       Mesh 0         0
-----
Number of Bound SDPs : 6
-----
*A:ALA-12#
```

Table 30: Output fields: labels

Label	Description
Svc Id	The service identifier.
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP.

sap

Syntax

sap sap-id [detail]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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

Displays SAPs for the service in the form *slot/mdalport[channel]*. See [Common CLI command descriptions](#) for command syntax.

interface interface-name

Displays information for the specified IP interface.

ip-address ip-address

Displays information associated with the specified IP address.

detail

Displays detailed information for the SAP.

Output

The following outputs are examples of SAP information, and [Table 31: Output fields: service ID SAP](#) describes the output fields.

- [Sample output](#)
- [Sample output for 7210 SAS-Mxp](#)

Sample output

```
A:Dut-A>config>service>epipe# show service id 2011 sap 1/1/18
=====
Service Access Points(SAP)
=====
Service Id       : 2011
SAP              : 1/1/18                      Encap           : null
Dot1Q Ethertype : 0x8100                      QinQ Ethertype  : 0x8100
Description      : Default sap description for service id 2011

Admin State      : Up                         Oper State      : Up
Flags            : None
Last Status Change : 07/07/2009 14:39:57
Last Mgmt Change  : 07/07/2009 14:39:14
Admin MTU        : 1514                      Oper MTU        : 1514
LLF Admin State  : Up LLF Oper State : Clear
Ingress qos-policy : 10
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
Egr Agg Rate Limit : max                     Endpoint        : N/A

Acct. Pol        : None                      Collect Stats    : Disabled
Ignore Oper Down : Disabled
=====
A:Dut-A>config>service>epipe#
```

```

A:Dut-A>config>service>epipe# show service id 2011 sap 1/1/18 detail
=====
Service Access Points(SAP)
=====
Service Id      : 2011
SAP             : 1/1/18
Dot1Q Ethertype : 0x8100
Encap           : null
QinQ Ethertype  : 0x8100
Description     : Default sap description for service id 2011

Admin State     : Up
Flags           : None
Oper State      : Up
Last Status Change : 07/07/2009 14:39:57
Last Mgmt Change  : 07/07/2009 14:39:14
Admin MTU       : 1514
Oper MTU        : 1514
LLF Admin State : Up LLF Oper State : Clear
Ingress qos-policy : 10
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
Egr Agg Rate Limit : max
Endpoint        : N/A

Acct. Pol       : None
Collect Stats   : Disabled
Ignore Oper Down : Disabled

-----
Sap Statistics
-----
Ingress Stats:      Packets      Octets
Egress Stats:       0            0

-----
Sap per Meter stats
-----
Ingress Meter 1 (Unicast)
For. InProf         : 0
For. OutProf        : 0

Ingress Meter 2 (Unicast)
For. InProf         : 0
For. OutProf        : 0

Ingress Meter 3 (Unicast)
For. InProf         : 0
For. OutProf        : 0

Ingress Meter 4 (Unicast)
For. InProf         : 0
For. OutProf        : 0

=====
A:Dut-A>config>service>epipe#

*A:ces-A# show service id 1 sap 1/2/1.1 detail
=====
Service Access Points(SAP)
=====
Service Id      : 1
SAP             : 1/2/1.1
Encap           : cem
Description     : (Not Specified)
Admin State     : Up
Flags           : None
Oper State      : Up
Last Status Change : 07/06/2010 14:16:41

```

```

Last Mgmt Change   : 07/06/2010 11:31:34

Admin MTU          : 1514                Oper MTU          : 1514
Endpoint           : N/A

Acct. Pol          : None                Collect Stats       : Disabled
Ignore Oper Down   : Disabled
-----
QOS
-----
Ingress qos-policy : 1
-----
Sap Statistics
-----
Ingress Stats:      Packets      Octets
Egress Stats:      2815         613670
                   2815         613670
-----
CEM SAP Configuration Information
-----
Endpoint Type       : Unstruct. T1      Bit-rate            : 24
Payload Size        : 192               Jitter Buffer (ms)   : 5
Jitter Buffer (packets): 6               Playout Threshold (packets): 4
Use RTP Header      : No                Differential         : No
Timestamp Freq      : 0                 CAS Framing          : No CAS
Effective PDVT       : +/-2.984 ms

Cfg Alarm           : stray malformed pktloss overrun underrun
Alarm Status        :
-----
CEM SAP Statistics
-----
Egress Stats      Packets      Seconds      Events
Forwarded         : 2915
Dropped           : 0
Missing           : 0
Reordered Forwarded : 0
Underrun          : 0                0
Overrun           : 0                0
Misordered Dropped : 0
Malformed Dropped : 0
LBit Dropped      : 0
Multiple Dropped  : 0
Error             : 0
Severely Error    : 0
Unavailable       : 0
Failure Count     : 0                0
Jitter Buffer Depth : 3

Ingress Stats
Forwarded         : 2915
Dropped           : 0
=====

*A:Dut-A>show# service id 104 sap 1/2/1.2 detail

=====
Service Access Points(SAP)
=====
Service Id        : 104
SAP               : 1/2/1.2          Encap             : cem
Description       : (Not Specified)

```

```

Admin State      : Up                Oper State      : Up
Flags           : None
Last Status Change : 12/15/2010 07:39:05
Last Mgmt Change  : 12/15/2010 07:25:37

Admin MTU        : 1514              Oper MTU        : 1514
Endpoint         : N/A

Acct. Pol        : None              Collect Stats    : Disabled
Ignore Oper Down : Disabled

-----
QOS
-----
Ingress qos-policy : 1                Egress qos-policy : 1
-----
Aggregate Policer
-----
rate             : n/a                burst           : n/a
-----
Sap Statistics
-----
Ingress Stats:      Packets           Octets
                   786701             70803090
Egress Stats:       784531             70607790
Extra-Tag Drop Stats: n/a             n/a
-----
CEM SAP Configuration Information
-----
Endpoint Type      : NxDS0           Bit-rate         : 1
Payload Size       : 64               Jitter Buffer (ms) : 32
Jitter Buffer (packets): 4           Playout Threshold (packets): 3
Use RTP Header     : No              Differential      : No
Timestamp Freq     : 0               CAS Framing       : No CAS
Effective PDVT     : +/-16.0 ms

Cfg Alarm          : stray malformed pktloss overrun underrun
Alarm Status       :

-----
CEM SAP Statistics
-----
Packets           Seconds           Events
Egress Stats
Forwarded         : 784407
Dropped           : 132
Missing           : 0
Reordered Forwarded : 0
Underrun          : 2355              1
Overrun           : 0                  0
Misordered Dropped : 0
Malformed Dropped : 0
LBit Dropped      : 132
Multiple Dropped  : 0
Error              : 1
Severely Error    : 0
Unavailable       : 18
Failure Count     : 1
Jitter Buffer Depth : 2

Ingress Stats
Forwarded         : 786762
Dropped           : 0
=====
*A:Dut-A>show#

```


CLI output for 7210 SAS-T configured in access uplink mode:
 A:SAS-T-A0-2>show>service>id# sap 1/1/1:10. detail

=====

Service Access Points(SAP)

=====

Service Id	: 1		
SAP	: 1/1/1:10.*	Encap	: qinq
QinQ Dot1p	: Default		
Description	: (Not Specified)		
Admin State	: Up	Oper State	: Up
Flags	: None		
Last Status Change	: 04/29/2001 06:59:15		
Last Mgmt Change	: 04/28/2001 03:09:30		
Dot1Q Ethertype	: 0x8100	QinQ Ethertype	: 0x8100
Max Nbr of MAC Addr	: No Limit	Total MAC Addr	: 0
Learned MAC Addr	: 0	Static MAC Addr	: 0
Admin MTU	: 1522	Oper MTU	: 1522
Ingr IP Fltr-Id	: n/a	Egr IP Fltr-Id	: n/a
Ingr Mac Fltr-Id	: 1	Egr Mac Fltr-Id	: n/a
tod-suite	: None		
Mac Learning	: Enabled	Discard Unkwn Srce	: Disabled
Mac Aging	: Enabled	Mac Pinning	: Disabled
BPDU Translation	: Disabled		
L2PT Termination	: Disabled		
Acct. Pol	: None	Collect Stats	: Disabled
Ignore Oper Down	: Disabled		

Stp Service Access Point specifics

Stp Admin State	: Up	Stp Oper State	: Down
Core Connectivity	: Down		
Port Role	: N/A	Port State	: Forwarding
Port Number	: 2048	Port Priority	: 128
Port Path Cost	: 10	Auto Edge	: Enabled
Admin Edge	: Disabled	Oper Edge	: N/A
Link Type	: Pt-pt	BPDU Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: N/A
Last BPDU from	: N/A		
CIST Desig Bridge	: N/A	Designated Port	: N/A
Forward transitions	: 0	Bad BPDUs rcvd	: 0
Cfg BPDUs rcvd	: 0	Cfg BPDUs tx	: 0
TCN BPDUs rcvd	: 0	TCN BPDUs tx	: 0
RST BPDUs rcvd	: 0	RST BPDUs tx	: 0
MST BPDUs rcvd	: 0	MST BPDUs tx	: 0

ARP host

Admin State	: outOfService		
Host Limit	: 1	Min Auth Interval	: 15 minutes

QoS

Ingress qos-policy : 1

Aggregate Policer

```

rate           : n/a                burst           : n/a
-----
Ingress QoS Classifier Usage
-----
Classifiers Allocated: 4            Meters Allocated : 2
Classifiers Used      : 2            Meters Used       : 2
-----
Sap Statistics
-----
Ingress Stats:      Packets      Octets
                   142761481188  9707780720784
Egress Stats:       0            0
Extra-Tag Drop Stats: n/a        n/a
-----
Sap per Meter stats
-----
                   Packets      Octets

Ingress Meter 1 (Unicast)
For. InProf         : 17        1162
For. OutProf        : 0         0

Ingress Meter 11 (Multipoint)
For. InProf         : 61        4148
For. OutProf        : 142761547917 9707785259394
=====

```

Sample output for 7210 SAS-Mxp

```

*A:Dut-A# show service id 10 sap 5/1/1:800 detail
=====
Service Access Points(SAP)
=====
Service Id       : 10
SAP              : 5/1/1:800          Encap           : q-tag
Description      : (Not Specified)
Admin State      : Up                 Oper State      : Down
Flags           : PortOperDown
Last Status Change : 11/07/2017 04:48:25
Last Mgmt Change  : 11/07/2017 05:02:47
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
BGP IPv4 FlowSpec : Disabled
BGP IPv6 FlowSpec : Disabled
tod-suite        : None
Egr Agg Rate CIR  : 0                Egr Agg Rate PIR : max
Limit Unused BW   : Disabled
Collect Stats     : Disabled
Dynamic Hosts     : Enabled
Monitor Oper Grp  : (none)
Acct. Pol         : None
Anti Spoofing     : None
Oper Group        : (none)
Host Lockout Plcy : n/a
Lag Link Map Prof : (none)
-----
QoS
-----
Ingress qos-policy : 1              Egress qos-policy : 1
Table-based        : enabled
-----
Aggregate Policer

```

```

-----
Rate                : n/a                Burst                : n/a
-----
Egress Aggregate Meter
-----
Rate                : n/a                Burst                : n/a
-----
Ingress QoS Classifier Usage
-----
Classifiers Allocated: 60                Meters Allocated   : 30
Classifiers Used      : 9                 Meters Used        : 8
-----
Sap Statistics
-----
                Packets                Octets
Ingress Stats:      0                   0
Egress Stats:       0                   0
Ingress Drop Stats: 0                   0

Extra-Tag Drop Stats:  n/a                n/a
-----
Sap per Meter stats (in/out counter mode)
-----
                Packets                Octets
Ingress Meter 1
For. InProf         : 0                   0
For. OutProf        : 0                   0

Ingress Meter 2
For. InProf         : 0                   0
For. OutProf        : 0                   0

Ingress Meter 3
For. InProf         : 0                   0
For. OutProf        : 0                   0

Ingress Meter 4
For. InProf         : 0                   0
For. OutProf        : 0                   0

Ingress Meter 5
For. InProf         : 0                   0
For. OutProf        : 0                   0

Ingress Meter 6
For. InProf         : 0                   0
For. OutProf        : 0                   0

Ingress Meter 7
For. InProf         : 0                   0
For. OutProf        : 0                   0

Ingress Meter 8
For. InProf         : 0                   0
For. OutProf        : 0                   0
=====

```

Table 31: Output fields: service ID SAP

Label	Description
Service Id	The service identifier.
SAP	The SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, TodResourceUnavail, Tod MssResourceUnavail, SapParamMismatch, ServiceMTUToo Small, SapIngressNamedPoolMismatch, SapEgressNamedPool Mismatch, NoSapEpipeRingNode.
Last Status Change	The time of the most recent operating status change to this SAP.
Last Mgmt Change	The time of the most recent management-initiated change to this SAP.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the port 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 port to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Egress qos-policy	The egress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Indicates whether collect stats is enabled.

Label	Description
Ignore Oper Down	Whether the user has enabled or disabled ignore-oper-down parameter.
LLF Admin State	The Link Loss Forwarding administrative state.
LLF Oper State	The Link Loss Forwarding operational state.
Loopback Mode	The Ethernet port loopback mode.
Loopback Src Addr	The configured loopback source address
Loopback Dst Addr	The configured loopback destination address.
No-svc-port used	The port ID of the port on which no service is configured. This port is used for the port loop back with MAC swap functionality.
Table-based	Indicates the use of table-based resource classification: Enabled (table-based) or Disabled (CAM-based)

sdp

Syntax

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

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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

Default All SDPs.

far-end ip-addr

Displays only SDPs matching the specified far-end IP address.

Default SDPs with any far-end IP address.

detail

Displays detailed SDP information.

Output

The following outputs are examples of service SDP information, and [Table 32: Output fields: service ID SDP](#) describes the output fields.

- [Sample output 1](#)
- [Sample output 2](#)
- [Sample output 3](#)
- [Sample output 4](#)
- [Sample output 5](#)

Sample output 1

```
A:Dut-A# show service id 1 sdp detail
=====
Services: Service Destination Points Details
=====
Sdp Id 1:1 -(10.20.1.2)
-----
Description      : Default sdp description
SDP Id           : 1:1                               Type           : Spoke
VC Type          : Ether                               VC Tag          : n/a
Admin Path MTU   : 0                                  Oper Path MTU   : 9186
Far End          : 10.20.1.2                           Delivery        : MPLS

Admin State      : Up                                Oper State      : Up
Acct. Pol        : None                              Collect Stats   : Disabled
Ingress Label    : 2048                              Egress Label    : 2048
Ing mac Fltr     : n/a                               Egr mac Fltr    : n/a
Ing ip Fltr      : n/a                               Egr ip Fltr     : n/a
Ing ipv6 Fltr    : n/a                               Egr ipv6 Fltr   : n/a
Admin ControlWord : Not Preferred                     Oper ControlWord : False
Last Status Change : 05/31/2007 00:45:43             Signaling       : None
Last Mgmt Change  : 05/31/2007 00:45:43
Class Fwding State : Up
Flags            : None
Peer Pw Bits     : None
Peer Fault Ip    : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None
Max Nbr of MAC Addr: No Limit                         Total MAC Addr  : 0
Learned MAC Addr  : 0                                Static MAC Addr  : 0

MAC Learning     : Enabled                           Discard Unkwn Srce: Disabled
MAC Aging        : Enabled
L2PT Termination : Disabled                           BPDU Translation : Disabled
MAC Pinning      : Disabled

KeepAlive Information :
Admin State       : Disabled                           Oper State      : Disabled
Hello Time       : 10                                Hello Msg Len    : 0
```

```

Max Drop Count      : 3                      Hold Down Time    : 10

Statistics          :
I. Fwd. Pkts.       : 0                      I. Dro. Pkts.       : 0
I. Fwd. Octs.       : 0                      I. Dro. Octs.       : 0
E. Fwd. Pkts.       : 0                      E. Fwd. Octets      : 0
MCAC Policy Name    :
MCAC Max Unconst BW: no limit                MCAC Max Mand BW    : no limit
MCAC In use Mand BW: 0                      MCAC Avail Mand BW  : unlimited
MCAC In use Opnl BW: 0                      MCAC Avail Opnl BW  : unlimited
Associated LSP LIST :
Lsp Name            : A_B_1                  Oper State          : Up
Admin State         : Up
Time Since Last Tr* : 00h26m35s

Lsp Name            : A_B_2                  Oper State          : Up
Admin State         : Up
Time Since Last Tr* : 00h26m35s

Lsp Name            : A_B_3                  Oper State          : Up
Admin State         : Up
Time Since Last Tr* : 00h26m34s

Lsp Name            : A_B_4                  Oper State          : Up
Admin State         : Up
Time Since Last Tr* : 00h26m34s

Lsp Name            : A_B_5                  Oper State          : Up
Admin State         : Up
Time Since Last Tr* : 00h26m34s

Lsp Name            : A_B_6                  Oper State          : Up
Admin State         : Up
Time Since Last Tr* : 00h26m34s

Lsp Name            : A_B_7                  Oper State          : Up
Admin State         : Up
Time Since Last Tr* : 00h26m34s

Lsp Name            : A_B_8                  Oper State          : Up
Admin State         : Up
Time Since Last Tr* : 00h26m35s

Lsp Name            : A_B_9                  Oper State          : Up
Admin State         : Up
Time Since Last Tr* : 00h26m34s

Lsp Name            : A_B_10                 Oper State          : Up
Admin State         : Up
Time Since Last Tr* : 00h26m34s
-----
Class-based forwarding :
-----
Class forwarding      : enabled
Default LSP           : A_B_10                Multicast LSP        : A_B_9
=====
FC Mapping Table
=====
FC Name               LSP Name
-----
af                    A_B_3
be                    A_B_1
ef                    A_B_6
hl                    A_B_7

```

```

h2          A_B_5
l1          A_B_4
l2          A_B_2
nc          A_B_8
=====
Stp Service Destination Point specifics
-----
Mac Move      : Blockable
Stp Admin State : Up                      Stp Oper State : Down
Core Connectivity : Down
Port Role      : N/A                      Port State      : Forwarding
Port Number    : 2049                     Port Priority    : 128
Port Path Cost : 10                       Auto Edge       : Enabled
Admin Edge     : Disabled                  Oper Edge       : N/A
Link Type      : Pt-pt                     BPDU Encap      : Dot1d
Root Guard     : Disabled                  Active Protocol  : N/A
Last BPDU from : N/A
Designated Bridge : N/A                   Designated Port Id: 0
Fwd Transitions : 0                       Bad BPDUs rcvd  : 0
Cfg BPDUs rcvd  : 0                       Cfg BPDUs tx    : 0
TCN BPDUs rcvd  : 0                       TCN BPDUs tx    : 0
RST BPDUs rcvd  : 0                       RST BPDUs tx    : 0
-----
Number of SDPs : 1
-----
* indicates that the corresponding row element may have been truncated.
-----
A:Dut-A#

```

Table 32: Output fields: service ID SDP

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
Split Horizon Group	Name of the split horizon group that the SDP belongs to.
VC Type	The VC type: ether, vlan, or vpls.
VC Tag	The explicit dot1Q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case).
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.

Label	Description
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The current state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the keepalive process.
Oper State	The operational state of the keepalive process.
Hello Time	Transmission frequency of the SDP echo request messages.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	The length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field.

Sample output 2

The following output is an example of when both sides have the control word disabled.

```
*A:ALA-Dut-B>config>service>epipe# show service id 2100 sdp detail
=====
Services: Service Destination Points Details
=====
Sdp Id 1:2001  -(1.1.1.1)
-----
Description    : Default sdp description
```

```

SDP Id      : 1:2001
VC Type     : Ether
Admin Path MTU : 1600
Far End     : 1.1.1.1
Admin State  : Up
Acct. Pol   : None
Ingress Label : 115066
Ing mac Fltr : n/a
Ing ip Fltr  : n/a
Ing ipv6 Fltr : n/a
Admin ControlWord : Not Preferred
Last Status Change : 02/05/2007 16:49:05
Last Mgmt Change  : 02/05/2007 16:47:54
Flags         : None
Peer Pw Bits   : None
Peer Fault Ip  : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None
Max Nbr of MAC Addr: No Limit
Learned MAC Addr : 0
MAC Learning   : Enabled
MAC Aging      : Enabled
L2PT Termination : Disabled
MAC Pinning    : Disabled
KeepAlive Information :
Admin State    : Disabled
Hello Time     : 10
Max Drop Count : 3
Statistics     :
I. Fwd. Pkts.  : 0
E. Fwd. Pkts.  : 0
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-----
Number of SDPs : 1
=====
*A:ALA-Dut-B>config>service>epipe#

```

Sample output 3

The following examples show both sides (PE nodes) when the control word is enabled.

```

*A:ALA-Dut-B>config>service>epipe# show service id 2100 sdp detail
=====
Services: Service Destination Points Details
-----
Sdp Id 1:2001  -(1.1.1.1)
-----
Description      : Default sdp description
SDP Id          : 1:2001
VC Type         : Ether
Admin Path MTU  : 1600
Far End         : 1.1.1.1
Type            : Spoke
VC Tag          : n/a
Oper Path MTU   : 1600
Delivery        : MPLS

Admin State     : Up
Acct. Pol      : None
Ingress Label   : 115066
Ing mac Fltr    : n/a
Ing ip Fltr     : n/a
Ing ipv6 Fltr   : n/a
Admin ControlWord : Preferred Oper ControlWord : True
Last Status Change : 02/05/2007 16:39:22
Last Mgmt Change  : 02/05/2007 16:39:22
Class Fwding State : Up
Oper State      : Up
Collect Stats   : Disabled
Egress Label    : 119068
Egr mac Fltr    : n/a
Egr ip Fltr     : n/a
Egr ipv6 Fltr   : n/a
Signaling       : TLDP

```

```

Endpoint      : N/A                      Precedence      : 4
Flags         : None
Peer Pw Bits  : None
Peer Fault Ip : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None
Max Nbr of MAC Addr: No Limit
Learned MAC Addr : 0
Total MAC Addr : 0
Static MAC Addr : 0

MAC Learning   : Enabled
MAC Aging      : Enabled
L2PT Termination : Disabled
MAC Pinning    : Disabled
Discard Unkwn Srce: Disabled
BPDU Translation : Disabled

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

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

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-----
Number of SDPs : 1
=====
*A:ALA-Dut-B>config>service>epipe#

```

Sample output 4

The following output is an example of when one side (PE) has the control word enabled (the pipe is down). The following is the side with the control word disabled.

```

*A:ALA-Dut-B>config>service>epipe# show service id 2100 sdp detail
=====
Services: Service Destination Points Details
=====
Sdp Id 1:2001  -(1.1.1.1)
-----
Description      : Default sdp description
SDP Id           : 1:2001
VC Type          : Ether
Admin Path MTU   : 1600
Far End          : 1.1.1.1
Type             : Spoke
VC Tag           : n/a
Oper Path MTU    : 1600
Delivery         : MPLS

Admin State      : Up
Acct. Pol        : None
Ingress Label    : 115066
Ing mac Fltr     : n/a
Ing ip Fltr      : n/a
Ing ipv6 Fltr    : n/a
Admin ControlWord : Not Preferred
Last Status Change : 02/05/2007 16:47:54
Last Mgmt Change  : 02/05/2007 16:47:54
Flags           : None
Peer Pw Bits     : None
Peer Fault Ip    : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None
Max Nbr of MAC Addr: No Limit
Learned MAC Addr : 0
MAC Learning     : Enabled
Oper State       : Down
Collect Stats    : Disabled
Egress Label     : 119068
Egr mac Fltr     : n/a
Egr ip Fltr      : n/a
Egr ipv6 Fltr    : n/a
Oper ControlWord : False
Signaling        : TLDP
Total MAC Addr   : 0
Static MAC Addr  : 0
Discard Unkwn Srce: Disabled

```

```

MAC Aging           : Enabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
KeepAlive Information :
Admin State        : Disabled
Hello Time         : 10
Max Drop Count     : 3
Statistics         :
I. Fwd. Pkts.      : 0
E. Fwd. Pkts.      : 0
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS

```

```
-----
Number of SDPs : 1
=====
```

```
*A:ALA-Dut-B>config>service>epipe#
```

Sample output 5

The following is the side with the control word enabled.

```

*A:ALA-B# show service id 2100 sdp detail
=====
Services: Service Destination Points Details
=====
Sdp Id 1:12000 - (3.3.3.3)
-----
Description      : Default sdp description
SDP Id          : 1:12000
VC Type         : Ether
Admin Path MTU   : 1600
Far End         : 3.3.3.3
Admin State     : Up
Acct. Pol       : None
Ingress Label   : 119066
Ing mac Fltr    : n/a
Ing ip Fltr     : n/a
Ing ipv6 Fltr   : n/a
Admin ControlWord : Preferred
Last Status Change : 02/04/2007 22:52:43
Last Mgmt Change  : 02/04/2007 02:06:08
Flags           : None
Peer Pw Bits     : None
Peer Fault Ip    : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None
Max Nbr of MAC Addr: No Limit
Learned MAC Addr : 0
MAC Learning     : Enabled
MAC Aging       : Enabled
L2PT Termination : Disabled
MAC Pinning     : Disabled
KeepAlive Information :
Admin State     : Disabled
Hello Time      : 10
Max Drop Count  : 3
Statistics      :
I. Fwd. Pkts.   : 0
E. Fwd. Pkts.   : 0
Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-----
Number of SDPs : 1
=====
Type           : Spoke
VC Tag         : n/a
Oper Path MTU  : 1600
Delivery       : MPLS
Oper State     : Down
Collect Stats  : Disabled
Egress Label   : 0
Egr mac Fltr   : n/a
Egr ip Fltr    : n/a
Egr ipv6 Fltr  : n/a
Oper ControlWord : True
Signaling      : TLDP
Total MAC Addr : 0
Static MAC Addr : 0
Discard Unkwn Srce: Disabled
BPDU Translation : Disabled
Oper State     : Disabled
Hello Msg Len  : 0
Hold Down Time : 10
I. Dro. Pkts.  : 0
E. Fwd. Octets : 0

```

```
=====
*A:ALA-B#
```

split-horizon-group

Syntax

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

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays service split horizon groups.

Output

The following output is an example of split horizon group information, and [Table 33: Output fields: Split Horizon group](#) describes the output fields.

Sample output

```
*A:7210-SAS>show>service# id 1 split-horizon-group

=====
Service: Split Horizon Group
=====
Name                               Description
-----
    access
-----
R = Residential Split Horizon Group
A = Auto Created Split Horizon Group
No. of Split Horizon Groups: 1
=====
*A:7210-SAS>show>service# id 1 split-horizon-group access

=====
Service: Split Horizon Group
=====
Name                               Description
-----
    access
-----
Associations
-----
R = Residential Split Horizon Group
SAPs Associated : 0                SDPs Associated : 0
*A:7210-SAS>show>service#
```

Table 33: Output fields: Split Horizon group

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

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays information for the spanning tree protocol (STP) instance for the service.

Parameters

detail

Displays detailed information.

Output

The following output is an example of STP information, and [Table 34: Output fields: STP](#) describes the output fields.

Sample output

```
A:Dut-A>show>service>id# stp
=====
Stp info, Service 305
=====
Bridge Id       : 00:0d.00:20:ab:cd:00:01  Top. Change Count : 5
Root Bridge    : This Bridge               Stp Oper State    : Up
Primary Bridge : N/A                       Topology Change   : Inactive
Mode           : Rstp                      Last Top. Change  : 0d 08:35:16
Vcp Active Prot. : N/A
Root Port      : N/A                       External RPC      : 0
=====
Stp port info
=====
```

Sap/Sdp Id	Oper- State	Port- Role	Port- State	Port- Num	Oper- Edge	Link- Type	Active Prot.
1/1/16:305	Up	Designated	Forward	2048	False	Pt-pt	Rstp
lag-4:305	Up	Designated	Forward	2000	False	Pt-pt	Rstp
1217:305	Up	N/A	Forward	2049	N/A	Pt-pt	N/A
1317:305	Up	N/A	Forward	2050	N/A	Pt-pt	N/A
1417:305	Up	N/A	Forward	2051	N/A	Pt-pt	N/A
1617:305	Pruned	N/A	Discard	2052	N/A	Pt-pt	N/A

```

=====
A:Dut-A>show>service>id#

A:Dut-A>show>service>id# stp detail
=====
Spanning Tree Information
=====
VPLS Spanning Tree Information
-----
VPLS oper state      : Up
Stp Admin State      : Up
Mode                  : Rstp
Core Connectivity     : Down
Stp Oper State        : Up
Vcp Active Prot.     : N/A

Bridge Id             : 00:0d.00:20:ab:cd:00:01
Bridge Priority        : 0
Topology Change        : Inactive
Last Top. Change      : 0d 08:35:29
Top. Change Count      : 5
MST region revision    : 0
MST region name       :

Bridge Instance Id: 13
Tx Hold Count         : 6
Bridge Hello Time      : 2
Bridge Max Age         : 20
Bridge Fwd Delay       : 15
Bridge max hops        : 20

Root Bridge           : This Bridge
Primary Bridge         : N/A

Root Path Cost         : 0
Rcvd Hello Time        : 2
Root Priority           : 13
Root Forward Delay     : 15
Root Max Age           : 20
Root Port              : N/A
-----
Spanning Tree Sap/Spoke SDP Specifics
-----
SAP Identifier        : 1/1/16:305
Port Role              : Designated
Port Number           : 2048
Port Path Cost         : 10
Admin Edge             : Disabled
Link Type              : Pt-pt
Root Guard             : Disabled
Last BPDUs from        : 80:04.00:0a:1b:2c:3d:4e
CIST Desig Bridge      : This Bridge
Forward transitions     : 5
Cfg BPDUs rcvd         : 0
TCN BPDUs rcvd         : 0
RST BPDUs rcvd         : 29
MST BPDUs rcvd         : 0
Stp Admin State        : Up
Port State             : Forwarding
Port Priority           : 128
Auto Edge              : Enabled
Oper Edge              : False
BPDU Encap             : PVST
Active Protocol        : Rstp
Designated Port        : 34816
Bad BPDUs rcvd         : 0
Cfg BPDUs tx           : 0
TCN BPDUs tx           : 0
RST BPDUs tx           : 23488
MST BPDUs tx           : 0

SAP Identifier        : lag-4:305
Port Role              : Designated
Port Number           : 2000
Port Path Cost         : 10
Admin Edge             : Disabled
Link Type              : Pt-pt
Root Guard             : Disabled
Last BPDUs from        : 80:04.00:0a:1b:2c:3d:4e
CIST Desig Bridge      : This Bridge
Stp Admin State        : Up
Port State             : Forwarding
Port Priority           : 128
Auto Edge              : Enabled
Oper Edge              : False
BPDU Encap             : Dot1d
Active Protocol        : Rstp
Designated Port        : 34768

```

```

Forward transitions: 4
Cfg BPDUs rcvd : 0
TCN BPDUs rcvd : 0
RST BPDUs rcvd : 23
MST BPDUs rcvd : 0

SDP Identifier : 1217:305
Port Role : N/A
Port Number : 2049
Port Path Cost : 10
Admin Edge : Disabled
Link Type : Pt-pt
Root Guard : Disabled
Last BPDUs from : N/A
Designated Bridge : N/A
Fwd Transitions : 0
Cfg BPDUs rcvd : 0
TCN BPDUs rcvd : 0
RST BPDUs rcvd : 0

SDP Identifier : 1317:305
Port Role : N/A
Port Number : 2050
Port Path Cost : 10
Admin Edge : Disabled
Link Type : Pt-pt
Root Guard : Disabled
Last BPDUs from : N/A
Designated Bridge : N/A
Fwd Transitions : 0
Cfg BPDUs rcvd : 0
TCN BPDUs rcvd : 0
RST BPDUs rcvd : 0

SDP Identifier : 1417:305
Port Role : N/A
Port Number : 2051
Port Path Cost : 10
Admin Edge : Disabled
Link Type : Pt-pt
Root Guard : Disabled
Last BPDUs from : N/A
Designated Bridge : N/A
Fwd Transitions : 1
Cfg BPDUs rcvd : 0
TCN BPDUs rcvd : 0
RST BPDUs rcvd : 0

SDP Identifier : 1617:305
Port Role : N/A
Port Number : 2052
Port Path Cost : 10
Admin Edge : Disabled
Link Type : Pt-pt
Root Guard : Disabled
Last BPDUs from : N/A
Designated Bridge : N/A
Fwd Transitions : 0
Cfg BPDUs rcvd : 0
TCN BPDUs rcvd : 0
RST BPDUs rcvd : 0

Bad BPDUs rcvd : 0
Cfg BPDUs tx : 0
TCN BPDUs tx : 0
RST BPDUs tx : 23454
MST BPDUs tx : 0

Stp Admin State : Down
Port State : Forwarding
Port Priority : 128
Auto Edge : Enabled
Oper Edge : N/A
BPDU Encap : Dot1d
Active Protocol : N/A

Designated Port Id: 0
Bad BPDUs rcvd : 0
Cfg BPDUs tx : 0
TCN BPDUs tx : 0
RST BPDUs tx : 0

Stp Admin State : Down
Port State : Forwarding
Port Priority : 128
Auto Edge : Enabled
Oper Edge : N/A
BPDU Encap : Dot1d
Active Protocol : N/A

Designated Port Id: 0
Bad BPDUs rcvd : 0
Cfg BPDUs tx : 0
TCN BPDUs tx : 0
RST BPDUs tx : 0

Stp Admin State : Down
Port State : Forwarding
Port Priority : 128
Auto Edge : Enabled
Oper Edge : N/A
BPDU Encap : Dot1d
Active Protocol : N/A

Designated Port Id: 0
Bad BPDUs rcvd : 0
Cfg BPDUs tx : 0
TCN BPDUs tx : 0
RST BPDUs tx : 0

Stp Admin State : Down
Port State : Discarding
Port Priority : 128
Auto Edge : Enabled
Oper Edge : N/A
BPDU Encap : Dot1d
Active Protocol : N/A

Designated Port Id: 0
Bad BPDUs rcvd : 0
Cfg BPDUs tx : 0
TCN BPDUs tx : 0
RST BPDUs tx : 0
=====
A:Dut-A>show>service>id#

```



```
*7210-SAS>show>service>id# stp detail
```

```
=====
Spanning Tree Information
=====
```

```
-----
VPLS Spanning Tree Information
-----
```

VPLS oper state	: Up	Core Connectivity	: Down
Stp Admin State	: Up	Stp Oper State	: Up
Mode	: Mstp	Vcp Active Prot.	: N/A
Bridge Id	: 80:00.00:25:ba:04:66:a0	Bridge Instance Id	: 0
Bridge Priority	: 32768	Tx Hold Count	: 6
Topology Change	: Inactive	Bridge Hello Time	: 2
Last Top. Change	: 0d 02:54:16	Bridge Max Age	: 20
Top. Change Count	: 27	Bridge Fwd Delay	: 15
Root Bridge	: 40:00.7c:20:64:ac:ff:63		
Primary Bridge	: N/A		
Root Path Cost	: 10	Root Forward Delay	: 15
Rcvd Hello Time	: 2	Root Max Age	: 20
Root Priority	: 16384	Root Port	: 2048
MSTP info for CIST	:		
Regional Root	: 80:00.7c:20:64:ad:04:5f	Root Port	: 2048
Internal RPC	: 10	Remaining Hopcount	: 19
MSTP info for MSTI 1	:		
Regional Root	: This Bridge	Root Port	: N/A
Internal RPC	: 0	Remaining Hopcount	: 20
MSTP info for MSTI 2	:		
Regional Root	: 00:02.7c:20:64:ad:04:5f	Root Port	: 2048
Internal RPC	: 10	Remaining Hopcount	: 19

```
-----
Spanning Tree Sap Specifics
-----
```

SAP Identifier	: 1/1/7:0	Stp Admin State	: Up
Port Role	: Root	Port State	: Forwarding
Port Number	: 2048	Port Priority	: 128
Port Path Cost	: 10	Auto Edge	: Enabled
Admin Edge	: Disabled	Oper Edge	: False
Link Type	: Pt-pt	BPDUs Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: Mstp
Last BPDUs from	: 80:00.7c:20:64:ad:04:5f	Inside Mst Region	: True
CIST Desig Bridge	: 80:00.7c:20:64:ad:04:5f	Designated Port	: 34816
MSTI 1 Port Prio	: 128	Port Path Cost	: 10
MSTI 1 Desig Brid	: This Bridge	Designated Port	: 34816
MSTI 2 Port Prio	: 128	Port Path Cost	: 10
MSTI 2 Desig Brid	: 00:02.7c:20:64:ad:04:5f	Designated Port	: 34816
Forward transitions	: 17	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	: 7310	MST BPDUs tx	: 7277
SAP Identifier	: 1/1/8:0	Stp Admin State	: Up
Port Role	: Alternate	Port State	: Discarding
Port Number	: 2049	Port Priority	: 128
Port Path Cost	: 10	Auto Edge	: Enabled
Admin Edge	: Disabled	Oper Edge	: False

Link Type	: Pt-pt	BPDUs Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: Mstp
Last BPDUs from	: 80:00.7c:20:64:ad:04:5f	Inside Mst Region	: True
CIST Desig Bridge	: 80:00.7c:20:64:ad:04:5f	Designated Port	: 34817
MSTI 1 Port Prio	: 128	Port Path Cost	: 10
MSTI 1 Desig Brid	: This Bridge	Designated Port	: 34817
MSTI 2 Port Prio	: 128	Port Path Cost	: 10
MSTI 2 Desig Brid	: 00:02.7c:20:64:ad:04:5f	Designated Port	: 34817
Forward transitions:	: 14	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	: 7326	MST BPDUs tx	: 7307
SAP Identifier	: 1/1/9:0	Stp Admin State	: Up
Port Role	: Designated	Port State	: Forwarding
Port Number	: 2050	Port Priority	: 128
Port Path Cost	: 10	Auto Edge	: Enabled
Admin Edge	: Disabled	Oper Edge	: True
Link Type	: Pt-pt	BPDUs Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: Mstp
Last BPDUs from	: N/A	Inside Mst Region	: True
CIST Desig Bridge	: This Bridge	Designated Port	: 34818
MSTI 1 Port Prio	: 128	Port Path Cost	: 10
MSTI 1 Desig Brid	: This Bridge	Designated Port	: 34818
MSTI 2 Port Prio	: 128	Port Path Cost	: 10
MSTI 2 Desig Brid	: This Bridge	Designated Port	: 34818
Forward transitions:	: 2	Bad BPDUs rcvd	: 0
Cfg BPDUs rcvd	: 0	Cfg BPDUs tx	: 0
TCN BPDUs rcvd	: 0	TCN BPDUs tx	: 0
RST BPDUs rcvd	: 0	RST BPDUs tx	: 0
MST BPDUs rcvd	: 0	MST BPDUs tx	: 7415
SAP Identifier	: 1/1/25:0	Stp Admin State	: Up
Port Role	: Alternate	Port State	: Discarding
Port Number	: 2051	Port Priority	: 128
Port Path Cost	: 10	Auto Edge	: Enabled
Admin Edge	: Disabled	Oper Edge	: False
Link Type	: Pt-pt	BPDUs Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: Mstp
Last BPDUs from	: 80:00.7c:20:64:ad:04:5f	Inside Mst Region	: True
CIST Desig Bridge	: 80:00.7c:20:64:ad:04:5f	Designated Port	: 34820
MSTI 1 Port Prio	: 128	Port Path Cost	: 10
MSTI 1 Desig Brid	: This Bridge	Designated Port	: 34819
MSTI 2 Port Prio	: 128	Port Path Cost	: 10
MSTI 2 Desig Brid	: 00:02.7c:20:64:ad:04:5f	Designated Port	: 34820
Forward transitions:	: 10	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	: 7329	MST BPDUs tx	: 7303
SAP Identifier	: lag-1:0	Stp Admin State	: Up
Port Role	: Alternate	Port State	: Discarding
Port Number	: 2052	Port Priority	: 128
Port Path Cost	: 10	Auto Edge	: Enabled
Admin Edge	: Disabled	Oper Edge	: False
Link Type	: Pt-pt	BPDUs Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: Mstp
Last BPDUs from	: 80:00.7c:20:64:ad:04:5f	Inside Mst Region	: True
CIST Desig Bridge	: 80:00.7c:20:64:ad:04:5f	Designated Port	: 34822
MSTI 1 Port Prio	: 128	Port Path Cost	: 10
MSTI 1 Desig Brid	: This Bridge	Designated Port	: 34820
MSTI 2 Port Prio	: 128	Port Path Cost	: 10

```

MSTI 2 Desig Brid : 00:02.7c:20:64:ad:04:5f Designated Port : 34822
Forward transitions: 11 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 : 7322 MST BPDUs tx : 7299
=====

```

Table 34: Output fields: STP

Label	Description
RSTP Admin State	Indicates the administrative state of the Rapid Spanning Tree Protocol instance associated with this service.
Core Connectivity	Indicates the connectivity status to the core.
RSTP Oper State	Indicates the operational state of the Rapid Spanning Tree Protocol instance associated with this service. This field is applicable only when STP is enabled on the router.
Bridge-id	Specifies the MAC address used to identify this bridge in the network.
Hold Time	Specifies the interval length during which no more than two Configuration BPDUs shall be transmitted by this bridge.
Bridge fwd delay	Specifies how fast a bridge changes its state when moving toward the forwarding state.
Bridge Hello time	Specifies the amount of time between the transmission of Configuration BPDUs.
Bridge max age	Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded. This is the actual value that this bridge is currently using.
Bridge priority	Defines the priority of the Spanning Tree Protocol instance associated with this service.
Topology change	Specifies whether a topology change is currently in progress.
Last Top. change	Specifies the time (in hundredths of a second) since the last time a topology change was detected by the Spanning Tree Protocol instance associated with this service.
Top. change count	Specifies the total number of topology changes detected by the Spanning Tree Protocol instance associated with this service since the management entity was last reset or initialized.
Root bridge-id	Specifies the bridge identifier of the root of the spanning tree as determined by the Spanning Tree Protocol instance associated

Label	Description
	with this service. This value is used as the Root Identifier parameter in all Configuration BPDUs originated by this node.
Root path cost	Specifies the cost of the path to the root bridge as seen from this bridge.
Root forward delay	Specifies how fast the root changes its state when moving toward the forwarding state.
Root hello time	Specifies the amount of time between the transmission of configuration BPDUs.
Root max age	Specifies the maximum age of Spanning Tree Protocol information learned from the network on any port before it is discarded.
Root priority	This object specifies the priority of the bridge that is currently selected as root-bridge for the network.
Root port	Specifies the port number of the port which provides the lowest cost path from this bridge to the root bridge.
SAP Identifier	The ID of the access port where this SAP is defined.
RSTP State	The operational state of RSTP.
STP Port State	Specifies the port identifier of the port on the designated bridge for this port's segment.
BPDU encap	Specifies the type of encapsulation used on BPDUs sent out and received on this SAP.
Port Number	Specifies the value of the port number field which is contained in the least significant 12 bits of the 16-bit port ID associated with this SAP.
Priority	Specifies the value of the port priority field which is contained in the most significant 4 bits of the 16-bit port ID associated with this SAP.
Cost	Specifies the contribution of this port to the path cost of paths toward the spanning tree root which include this port.
Fast Start	Specifies whether Fast Start is enabled on this SAP.
Designated Port	Specifies the port identifier of the port on the designated bridge for this port's segment.
Designated Bridge	Specifies the bridge identifier of the bridge which this port considers to be the designated bridge for this port's segment.

sap-using

Syntax

```
sap-using [sap sap-id]  
sap-using interface [ip-address | ip-int-name]  
sap-using [ingress | egress] filter filter-id  
sap-using [ingress | egress] qos-policy qos-policy-id  
sap-using encap-type encap-type
```

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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

ip-addr

The IP address of the interface for which to display matching SAPs.

Values a.b.c.d

ip-int-name

Specifies the IP interface name for which to display matching SAPs.

ingress

Specifies matching an ingress policy.

egress

Specifies matching an egress policy.

qos-policy *qos-policy-id*

Specifies the ingress QoS Policy ID for which to display matching SAPs.

Values 1 to 65535

filter *filter-id*

Specifies the ingress or egress filter policy ID for which to display matching SAPs.

Values 1 to 65535

sap sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

encap-type encap-type

Displays the CEM encapsulation type.

Values cem

Output

The following output is an example of service SAP information, and [Table 35: Output fields: SAP-using](#) describes the output fields.

Sample output

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

```
=====
Service Access Points
=====
```

PortId	SvcId	Ing. QoS	Ing. Fltr	Egr. QoS	Egr. Fltr	Adm	Opr
1/1/1:1	1	1	none	1	none	Up	Up
2/1/2:10/11	1	1	none	1	none	Up	Up
2/1/2:10/12	1	1	none	1	none	Up	Up
2/1/2:20/11	1	1	none	1	none	Up	Up
2/1/2:20/12	1	1	none	1	none	Up	Up
2/1/4:cp.10	10	1	none	1	none	Up	Up
2/1/4:cp.20	20	1	none	1	none	Up	Up

```
-----
Number of SAPs : 7
-----
=====
```

```
A:Dut-A>config>service>vpls# show service sap-using
```

```
=====
Service Access Points
=====
```

PortId	SvcId	Ing. QoS	Ing. Fltr	Egr. Fltr	Adm	Opr
lag-3:100	100	1	none	none	Up	Up
1/1/3	101	10	mac	none	Up	Up
lag-3:101	101	10	mac	none	Up	Up
lag-3:102	102	10	mac	none	Up	Up
lag-3:103	103	10	mac	none	Up	Up
lag-3:104	104	10	mac	none	Up	Up
lag-3:105	105	10	mac	none	Up	Up
lag-3:201	201	10	mac	none	Up	Up
lag-3:202	202	10	mac	none	Up	Up
lag-3:203	203	10	mac	none	Up	Up
lag-3:204	204	10	mac	none	Up	Up
lag-3:205	205	10	mac	none	Up	Up
1/1/16:301	301	10	mac	none	Up	Up
lag-4:301	301	10	mac	none	Up	Up
1/1/16:302	302	10	mac	none	Up	Up
lag-4:302	302	10	mac	none	Up	Up
1/1/16:303	303	10	mac	none	Up	Up
lag-4:303	303	10	mac	none	Up	Up

```

1/1/16:304          304      10    mac    none    Up    Up
lag-4:304           304      10    mac    none    Up    Up
1/1/16:305          305      10    mac    none    Up    Up
lag-4:305           305      10    mac    none    Up    Up
...
=====
A:Dut-A>config>service>vpls#

A:Dut-A>config>service# show service sap-using sap 1/1/16:305
=====
Service Access Points Using Port 1/1/16:305
=====
PortId              SvcId      Ing.   Ing.   Egr.   Adm   Opr
                   QoS      Fltr   Fltr
-----
1/1/16:305          305        10    mac    none    Up    Up
-----
Number of SAPs : 1
=====
A:Dut-A>config>service#

A:ces-A# show service sap-using sap 1/2/1.1
=====
Service Access Points
=====
PortId              SvcId      Ing.   Ing.   Egr.   Adm   Opr
                   QoS      Fltr   Fltr
-----
1/2/1.1             1          12    none    none    Up    Up
-----
Number of SAPs : 1
=====
A:ces-A#

*A:ces-A# show service sap-using sap 1/2/1.1
=====
Service Access Points
=====
PortId              SvcId      Ing.   Ing.   Egr.   Adm   Opr
                   QoS      Fltr   Fltr
-----
1/2/1.1             1          1     none    none    Up    Up
-----
Number of SAPs : 1
=====

*A:ces-A# show service sap-using encap-type cem
=====
Service Access Points Using Encap Type 'cem'
=====
PortId              SvcId      Adm     Opr     Alarm
-----
1/2/1.1             1          Up      Up      No
1/2/2.1             2          Up      Up      No
1/2/3.1             3          Up      Down    Yes
1/2/4.1             4          Up      Down    Yes
-----

```

```
Number of SAPS : 4
-----
=====
```

Table 35: Output fields: SAP-using

Label	Description
Port ID	The ID of the access port where the SAP is defined.
Svc ID	The service identifier.
MTU	The port MTU value.
Ing. QoS	The SAP ingress QoS policy number specified on the ingress SAP.
Ing Fltr	The MAC or IP filter policy ID applied to the ingress SAP.
Egr. QoS	The SAP egress QoS policy number specified on the egress SAP.
Egr. Fltr	The MAC or IP filter policy ID applied to the egress SAP.
Adm	The administrative state of the SAP.
Opr	The operational state of the SAP.

sdp

Syntax

sdp [*sdp-id* | *far-end ip-address*] [**detail** | **keep-alive-history**]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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.

far-end ip-address

Displays only SDPs matching with the specified far-end IP address.

Default SDPs with any far-end IP address.

detail

Displays detailed SDP information.

Default SDP summary output.

keep-alive-history

Displays the last fifty SDP keepalive events for the SDP.

Default SDP summary output.

Output

The following output is an example of SDP information, and [Table 36: Output fields: SDP](#) describes the output fields.

Sample output

```
*A:ALA-12# show service sdp
=====
Services: Service Destination Points
=====
SdpId    Adm MTU   Opr MTU   IP address    Adm  Opr      Deliver Signal
-----
10        4462     4462      10.20.1.3     Up   Dn NotReady MPLS   TLDP
40        4462     1534      10.20.1.20    Up   Up        MPLS   TLDP
60        4462     1514      10.20.1.21    Up   Up        MPLS   TLDP
100       4462     4462      10.0.0.2      Down Down      MPLS   TLDP
500       4462     4462      10.20.1.50    Up   Dn NotReady MPLS   TLDP
-----
Number of SDPs : 5
=====
*A:ALA-12#

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

```

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

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
=====
*A:ALA-12#
*A:ALA-12# show service sdp 8
=====
Service Destination Point (Sdp Id : 8)
=====
SdpId   Adm MTU   Opr MTU   IP address   Adm  Opr       Deliver Signal
-----
8       4462    4462     10.10.10.104 Up   Dn NotReady MPLS   TLDP
=====
*A:ALA-12#

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

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

Associated LSP LIST :
Lsp Name              : to-104
Admin State            : Up                    Oper State            : Down
Time Since Last Tran* : 01d07h36m
=====
* indicates that the corresponding row element may have been truncated.
*A:ALA-12#

```

Table 36: Output fields: SDP

Label	Description
SDP Id	The SDP identifier.

Label	Description
Adm 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.
Opr 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.
IP address	The IP address of the remote end of the MPLS tunnel defined by this SDP.
Adm Admin State	The desired state of the SDP.
Opr Oper State	The operating state of the SDP.
Deliver Delivery	The type of delivery used by the SDP: MPLS.
Flags	All the conditions that affect the operating status of this SDP.
Signal Signaling	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
Last Status Change	The time of the most recent operating status change to this SDP.
Last Mgmt Change	The time of the most recent management-initiated change to this SDP.
Number of SDPs	The total number of SDPs displayed according to the criteria specified.
Hello Time	How often the SDP echo request messages are transmitted on this SDP.
Number of SDPs	The total number of SDPs displayed according to the criteria specified.
Hello Time	How often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	The length of the SDP echo request messages transmitted on this SDP.
Hello Timeout	The number of seconds to wait for an SDP echo response message before declaring a timeout.
Unmatched Replies	The number of SDP unmatched message replies.

Label	Description
Max Drop Count	The maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	The amount of time to wait before the keepalive operating status is eligible to enter the alive state.
TX Hello Msgs	The number of SDP echo request messages transmitted since the keepalive was administratively enabled or the counter was cleared.
Rx Hello Msgs	The number of SDP echo request messages received since the keepalive was administratively enabled or the counter was cleared.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field.

sdp-using

Syntax

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

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Display 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

Specifies the virtual circuit identifier.

Values 1 to 4294967295

far-end ip-address

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

Default Services with any far-end IP address.

Output

The following output is an example of service SDP information, and [Table 37: Output fields: SDP-using](#) describes the output fields.

Sample output

```
*A:ALA-1# show service sdp-using 300
=====
Service Destination Point (Sdp Id : 300)
=====
SvcId      SdpId      Type Far End      Opr State I.Label  E.Label
-----
1          300:1      Mesh 10.0.0.13    Up      131071  131071
2          300:2      Spok 10.0.0.13    Up      131070  131070
100        300:100    Mesh 10.0.0.13    Up      131069  131069
101        300:101    Mesh 10.0.0.13    Up      131068  131068
102        300:102    Mesh 10.0.0.13    Up      131067  131067
-----
Number of SDPs : 5
-----
*A:ALA-1#
*A:ces-A# show service sdp-using
=====
SDP Using
=====
SvcId      SdpId      Type  Far End      Opr S* I.Label  E.Label
-----
1          12:1      Spok  2.2.2.2      Up    131063  131062
2          12:2      Spok  2.2.2.2      Up    131062  131069
3          122:3     Spok  2.2.2.2      Up    131069  131068
4          12:4      Spok  2.2.2.2      Up    131061  131061
-----
Number of SDPs : 4
-----
=====
*A:ces-A#
```

Table 37: Output fields: SDP-using

Label	Description
Svc ID	The service identifier.
Sdp ID	The SDP identifier.
Type	Type of SDP: spoke or mesh.
Far End	The far end address of the SDP.
Oper State	The operational state of the service.

Label	Description
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.

service-using

Syntax

service-using [**sdp** *sdp-id*] [**b-vpls**] [**i-vpls**] [**m-vpls**] [**sdp** *sdp-id*] [**customer** *customer-id*]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays the services matching certain usage properties.

If no optional parameters are specified, all services defined on the system are displayed.

Parameters

[service]

Displays information for the specified service type.

b-vpls

Specifies the B-component instance of the Provider Backbone Bridging (PBB/IEEE 802.1ah) feature. It represents the multi-point tunneling component that multiplexes multiple customer VPNs (ISIDs) together. It is similar to a regular VPLS instance that operates on the backbone MAC addresses.

i-vpls

Specifies the I-component instance of the Provider Backbone Bridging (PBB/IEEE 802.1ah) feature. It identifies the specific VPN entity associated to a customer multipoint (ELAN) service. It is similar to a regular VPLS instance that operates on the customer MAC addresses.

m-vpls

Specifies the M-component (managed VPLS) instance of the Provider Backbone Bridging (PBB/IEEE 802.1ah) feature.

sdp *sdp-id*

Displays only services bound to the specified SDP ID.

Values 1 to 17407

Default Services bound to any SDP ID.

customer customer-id

Displays services only associated with the specified customer ID.

Values 1 to 2147483647

Default Services associated with any customer.

Output

The following output is an example of service information, and [Table 38: Output fields: Service-using](#) describes the output fields.

Sample output

```
*A:ALA-12# show service service-using customer 10
=====
Services
=====
ServiceId  Type      Adm   Opr      CustomerId  Last Mgmt Change
-----
1          VPLS      Up    Up        10          09/05/2006 13:24:15
300        Epipe     Up    Up        10          09/05/2006 13:24:15
-----
Matching Services : 2
=====

*A:ALA-12#

*A:ALA-12# show service service-using
=====
Services
=====
ServiceId  Type      Adm   Opr      CustomerId  Last Mgmt Change
-----
1          uVPLS     Up    Up        1           10/26/2006 15:44:57
2          Epipe     Up    Down      1           10/26/2006 15:44:57
10         mVPLS     Down  Down      1           10/26/2006 15:44:57
11         mVPLS     Down  Down      1           10/26/2006 15:44:57
100        mVPLS     Up    Up        1           10/26/2006 15:44:57
101        mVPLS     Up    Up        1           10/26/2006 15:44:57
102        mVPLS     Up    Up        1           10/26/2006 15:44:57
999        uVPLS     Down  Down      1           10/26/2006 16:14:33
-----
Matching Services : 8
=====

*A:ALA-12#
```

Table 38: Output fields: Service-using

Label	Description
Service Id	The service identifier.

Label	Description
Type	The service type configured for the service ID.
Adm	The desired state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

3.9.2.3 Clear commands

id

Syntax

id *service-id*

Context

clear>service
clear>service>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears commands for a specific service.

Parameters

service-id

Specifies the ID that uniquely identifies a service.

Values service-id: 1 to 214748364
 svc-name: A string up to 64 characters.

spoke-sdp

Syntax

spoke-sdp *sdp-id:vc-id ingress-vc-label*

Context

clear>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command clears and resets the spoke-SDP bindings for the service.

Parameters

sdp-id

Specifies the spoke-SDP ID to be reset.

Values 1 to 17407

vc-id

Specifies the virtual circuit ID on the SDP ID to be reset.

Values 1 to 4294967295

ingress-vc-label

Specifies the ingress VC label to be cleared.

sap

Syntax

sap *sap-id* {all | cem | counters | stp}

Context

clear>service>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears SAP statistics for a SAP.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

all

Clears all SAP queue statistics and STP statistics.

counters

Clears all queue statistics associated with the SAP.

stp

Clears all STP statistics associated with the SAP.

l2pt

Clears all L2PT statistics associated with the SDP.

sdp**Syntax**

sdp *sdp-id* **keep-alive**

Context

clear>service>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command clears keepalive statistics associated with the SDP ID.

Parameters

sdp-id

Specifies the SDP ID for which to clear keepalive statistics.

Values 1 to 17407

counters**Syntax**

counters

Context

clear>service>statistics>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears all traffic counters associated with the service ID.

spoke-sdp

Syntax

spoke-sdp *sdp-id[:vc-id]* {**all** | **counters** | **stp**}

Context

clear>service>statistics>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command clears statistics for the spoke-SDP bound to the service.

Parameters

sdp-id

Specifies the spoke-SDP ID for which to clear statistics.

Values 1 to 17407

vc-id

Specifies 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.

stp

Clears all STP statistics associated with the SDP.

stp

Syntax

stp

Context

clear>service>statistics>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears all spanning tree statistics for the service ID.

statistics

Syntax

statistics

Context

clear>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context clear statistics for a specific service entity.

3.9.2.4 Debug commands

id

Syntax

id *service-id*

Context

debug>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command debugs commands for a specific service.

Parameters

service-id

The ID that uniquely identifies a service.

sap

Syntax

[no] sap *sap-id*

Context

debug>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables debugging for a particular SAP.

Parameters

sap-id

Specifies the SAP ID.

event-type

Syntax

[no] event-type {arp | config-change | oper-status-change}

Context

debug>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables a particular debugging event type.

The **no** form of this command disables the event type debugging.

Parameters

arp

Displays ARP events. This parameter is not supported for use with VLL services.

config-change

Debugs configuration change events.

svc-oper-status-change

Debugs service operational status changes.

Output

The following output is an example of event type debugging information.

Sample output

```
A:bksim180# debug service id 1000 sap 1/7/1 event-type arp
DEBUG OUTPUT show on CLI is as follows:
3 2008/11/17 18:13:24.35 UTC MINOR: DEBUG #2001 Base Service 1000 SAP 1/7/
1 "Service 1000 SAP 1/7/1:
RX: ARP_REQUEST (0x0001)
hwType      : 0x0001
prType      : 0x0800
hwLength    : 0x06
prLength    : 0x04
srcMac      : 8c:c7:01:07:00:03
destMac     : 00:00:00:00:00:00
srcIp       : 10.1.1.2
destIp      : 10.1.1.1
"

4 2008/11/17 18:13:24.35 UTC MINOR: DEBUG #2001 Base Service 1000 SAP 1/7/
1 "Service 1000 SAP 1/7/1:
TX: ARP_RESPONSE (0x0002)
hwType      : 0x0001
prType      : 0x0800
hwLength    : 0x06
prLength    : 0x04
srcMac      : 00:03:0a:0a:0a:0a
destMac     : 8c:c7:01:07:00:03
srcIp       : 10.1.1.1
destIp      : 10.1.1.2
"
```

sdp**Syntax**

[no] sdp *sdp-id:vc-id*

Context

debug>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables debugging for a particular SDP.

Parameters***sdp-id***

Specifies the SDP ID.

4 Ethernet Virtual Private Networks

This chapter provides information about Ethernet Virtual Private Networks (EVPN) on the 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone mode).

4.1 EVPN applications

EVPN, as described in RFC 7432, *BGP MPLS-Based Ethernet VPN*, is an IETF technology that uses a new BGP address family and allows Virtual Private LAN Services (VPLS) to operate in a similar manner to IP-VPNs, in which the MAC addresses and information to set up flooding trees are distributed by BGP.

EVPN is designed to fill the gaps of traditional L2VPN technologies, such as VPLS. The main objective of EVPN is to build E-LAN services similar to IP-VPNs defined in RFC 4364, while supporting MAC learning in the control plane (distributed using multi-protocol BGP (MP-BGP)), efficient multi-destination traffic delivery, and single-active/all-active multi-homing.

EVPN can be used as the control plane for different data plane encapsulations. The Nokia implementation supports EVPN for MPLS tunnels (EVPN-MPLS), where PEs are connected by any type of MPLS tunnel. EVPN-MPLS is generally used as an evolution for VPLS services. The EVPN-MPLS functionality is standardized in RFC 7432.

4.1.1 EVPN for MPLS tunnels in E-LAN services


The following figure shows the use of EVPN for MPLS tunnels on the 7210 SAS. In this example, EVPN is used as the control plane for E-LAN services.



Note:

The following figure shows generic EVPN capabilities. It does not imply support on all 7210 SAS platforms referenced in this guide. For more information about the supported EVPN capabilities for the platforms referenced in this guide, see the following sections in this chapter.

Figure 39: EVPN for MPLS in VPLS services



Service providers that offer E-LAN services request EVPN for its multi-homing capabilities and to leverage the optimization EVPN provides.

EVPN supports single-active multi-homing (per-service load-balancing multi-homing). Although VPLS already supports single-active multi-homing, EVPN single-active multi-homing is seen as the superior technology because of its mass-withdrawal capabilities to speed up convergence in scaled environments.

EVPN technology provides significant benefits, including:

- IP-VPN-like operation and control for E-LAN services
- reduction and (in some cases) suppression of the broadcast, unknown unicast, and multicast (BUM) 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
- potential to use single unified control-plane for both L2 VPN services and L3 VPN services
- superior multi-homing capabilities

The 7210 SAS EVPN-MPLS implementation is compliant with RFC 7432.

4.1.2 EVPN for MPLS tunnels in E-Line services

The MPLS network used by EVPN for E-LAN services can also be shared by E-Line services using EVPN in the control plane. EVPN for E-Line services (EVPN-VPWS) is a simplification of the RFC 7432 procedures, and is supported on the 7210 SAS in compliance with *draft-ietf-bess-evpn-vpws*.

4.2 EVPN for MPLS tunnels

This section provides information about EVPN for MPLS tunnels (EVPN-MPLS).

4.2.1 BGP-EVPN control plane for MPLS tunnels

The following table lists the supported EVPN routes and their usage in the 7210 SAS EVPN-MPLS.

Table 39: EVPN routes and usage


EVPN route	Usage	EVPN-MPLS 7210 SAS support
Type 1 - Ethernet auto-discovery route (A-D)	Mass-withdraw, ESI labels, aliasing	Yes ¹⁴
Type 2 - MAC/IP advertisement route	MAC/IP advertisement, IP advertisement for ARP resolution	Yes
Type 3 - Inclusive multicast Ethernet Tag route	Flooding tree setup (BUM flooding)	Yes
Type 4 - Ethernet segment (ES) route	ES discovery and DF election	Yes

RFC 7432 describes the BGP-EVPN control plane for MPLS tunnels. If EVPN multi-homing is not required, two route types are needed to set up a basic EVPN Instance (EVI): MAC/IP Advertisement and the Inclusive Multicast Ethernet Tag routes. If multi-homing is required, the ES and the Auto-Discovery routes are also needed.

When EVPN multi-homing is enabled in the system, two additional routes are required. The following figure shows the fields in routes type 1 and 4 and their associated extended communities.

¹⁴ ESI labels and aliasing are not supported on 7210 SAS platforms.

Figure 40: EVPN routes type 1 and 4



4.2.1.1 EVPN route type 3 — inclusive multicast Ethernet tag route

The 7210 SAS router generates this route type for setting up the flooding tree (BUM flooding) for a specified VPLS service. The received inclusive multicast routes add entries to the VPLS flood list. When BGP-EVPN MPLS is enabled, the standard supports ingress replication, p2mp mLDP, and composite tunnels as tunnel types in route type 3. On the 7210 SAS, only ingress replication is supported.

4.2.1.2 EVPN route type 2 — MAC/IP advertisement route

The 7210 SAS router generates this route type for advertising MAC addresses (and IP addresses if proxy-ARP/proxy-ND is enabled). The router generates MAC advertisement routes for the following entities:

- learned MACs on SAPs, if the **mac-advertisement** command is enabled
- conditional static MACs, if the **mac-advertisement** command is enabled
- learned MACs on spoke-SDPs, if the **mac-advertisement** command is enabled



Note:

- The **unknown-mac-route** command is not supported for EVPN-MPLS services.
- Proxy-ARP and proxy-ND commands are not supported for IP/MAC associations learned over spoke-SDP.

The route type 2 generated by a router uses the following fields and values:

- route distinguisher**

This is taken from the route distinguisher (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)**

This is zero (0) for MACs learned from single-homed CEs and different from zero for MACs learned from multi-homed CEs.

- **ethernet tag ID**
This is zero (0).
- **MAC address length**
This is always 48.
- **MAC address**
This is learned or statically configured.
- **IP address and IP address length**
 - This is the IP address associated with the MAC being advertised with a length of 32 (or 128 for IPv6).
 - In general, any MAC route without IP has IPL (IP length) = 0 and the IP is omitted.
 - When received, any IPL value not equal to zero, 32, or 128 discards the route.
- **MPLS label 1**
This 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 type 3 routes for the same service, unless **bgp-evpn mpls ingress-replication-bum-label** is configured in the service.
- **MPLS label 2**
This is 0.
- **MAC mobility extended community**
This is used for signaling the sequence number in case of mac moves and the "sticky" bit in case of advertising conditional static MACs. If a MAC route is received with a MAC mobility of *ext-community*, the sequence number and the "sticky" bit are considered for the route selection.

4.2.1.3 EVPN route type 1 — Ethernet Auto-Discovery route

The 7210 SAS router generates this route type to advertise for multi-homing functions. The system can generate two types of Ethernet Auto-Discovery (AD) routes:

- Ethernet AD route per-ESI
- Ethernet AD route per-EVI

The Ethernet AD per-ESI route uses the following fields and values:

- **route distinguisher**
This is taken from the service-level RD.
- **ethernet segment identifier (ESI)**
This contains a 10-byte identifier as configured in the system for a specified **ethernet-segment**.
- **ethernet tag ID**
This value, MAX-ET (0xFFFFFFFF), is reserved and used only for AD routes per ESI.
- **MPLS label**
This is zero (0).
- **ESI label extended community**

This includes the single-active bit (0 for all-active and 1 for single-active) and ESI label for all-active multi-homing split-horizon. The 7210 SAS always sets the single-active bit to "1", as long as single-active is supported. In addition, 7210 SAS does not allocate and send the ESI label for single-active multi-homing.

- **route-target extended community**

This is taken from the service level route target (RT).

The system can send only a separate Ethernet AD per-ESI route per service.

The Ethernet AD per-EVI route uses the following fields and values:

- **route distinguisher**

This is taken from the service level RD.

- **ethernet segment identifier (ESI)**

This contains a 10-byte identifier, as configured in the system for a specified **ethernet-segment**.

- **ethernet tag ID**

This is zero (0).

- **MPLS label**

This encodes the unicast label allocated for the service (high-order 20 bits).

- **route-target extended community**

This is taken from the service level RT.



Note:

The AD per-EVI route is not sent with the ESI label Extended Community.

4.2.1.4 EVPN route type 4 — ES route

The 7210 SAS router generates this route type for multi-homing ES discovery and DF (Designated Forwarder) election:

- **Route Distinguisher**

This is taken from the system level RD.

- **Ethernet Segment Identifier (ESI)**

This contains a 10-byte identifier, as configured in the system for a specified **ethernet-segment**.

- **ES-import route-target community**

This value is automatically derived from the MAC address portion of the ESI.

This extended community is treated as an RT and is supported by RT-constraint (route-target BGP family).

4.2.1.5 BGP tunnel encapsulation extended community

The following routes are sent with the RFC 5512 BGP Encapsulation Extended Community:

- MAC/IP
- Inclusive Multicast Ethernet Tag

- AD per-EVI routes

ES routes and AD per-ESI routes are not sent with this extended community.

The router processes MPLS encapsulation: 10, the BGP tunnel encapsulation tunnel value registered by IANA for RFC 5512. Any other tunnel value makes the route "treat-as-withdraw".

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 (as defined in RFC 5512) is not present in a received route, BGP treats the route as an MPLS. On the 7210 SAS, only MPLS encapsulation is supported.

4.2.2 EVPN for MPLS tunnels in VPLS services

EVPN can be used in MPLS networks where provider edge (PE) routers are interconnected through any type of MPLS tunnel, including RSVP-TE, LDP, BGP, Segment Routing IS-IS, or Segment Routing OSPF. As with VPRN services, the tunnel selection for a VPLS service (with BGP-EVPN MPLS enabled) is based on the **auto-bind-tunnel** command.

The EVPN-MPLS VPLS service uses a regular VPLS service where EVPN-MPLS "bindings" can coexist with SAPs.

Example: Sample configuration output that shows a VPLS service with EVPN-MPLS

```
*A:PE-1>config>service>vpls# info
-----
description "evpn-mpls-service"
bgp
  bgp-evpn
    evi 10
    mpls
    auto-bind-tunnel resolution any
    no shutdown
sap 1/1/1:1 create
exit
-----
```

First configure a **bgp-evpn** context as **mpls**. In addition, the minimum set of commands that must be configured to set up the EVPN-MPLS instance are the **evi** and the **auto-bind-tunnel resolution** commands.



Note:

Ensure that the EVI and the system IP are configured before executing the **configure>service/vpls>bgp-evpn>mpls>no shutdown** command.

The **evi** value, which is the EVPN instance (EVI) identifier, is unique in the system. It is used by the service-carving algorithm for multi-homing (if configured) and for auto-deriving RT and RDs in EVPN-MPLS services.

If the **evi** value is not specified, the value is zero and no RD or RTs are auto-derived from it. If it is specified, and no other RD or RT are configured in the service, the following applies:

- the RD is derived from: *system_ip:evi*
- the RT is derived from: *autonomous-system:evi*

**Note:**

When vsi-import and vsi-export polices are configured, the RT must be configured in the policies, and those values take precedence over the auto-derived RTs. The operational RT for a service is displayed by the **show service id svc-id bgp** command output. Nokia recommends that the user should not configure a VPLS ID using the **bgp-ad>vpls-id** command in the service.

When the **evi** command is configured, a **config>service>vpls>bgp** node (even empty) is required to output correct information using the **show service id 1 bgp** and **show service system bgp-route-distinguisher** commands.

The configuration of an EVI is enforced for EVPN services with SAPs in an **ethernet-segment**. See [EVPN multi-homing in VPLS services](#) for more information about ESs.

The following options are specific to EVPN-MPLS and are configured in the **config>service>vpls>bgp-evpn>mpls** context:

- **control-word**

Enable or disable the **control-word** command to guarantee interoperability to other vendors. In accordance with RFC 7432, this command is required to avoid frame disordering.

- **auto-bind-tunnel**

This command is used to select the type of MPLS transport tunnel used for a specific instance. The command is used in the same way as in VPRN services. See [auto-bind-tunnel](#) for more information.

For BGP-EVPN MPLS, **bgp** must be 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 pBits of the service-delimiting qtag in a new tag added in the customer frame before sending the frame to the EVPN core.

**Note:**

Nokia recommends that the user should not configure the **force-vlan-vc-forwarding** command on the 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE. Instead, Nokia recommends using the no force-vlan-vc-forwarding configuration, which is the default setting.

- **split-horizon-group**

This command associates a user-created split-horizon group to all the EVPN-MPLS destinations. See [EVPN and VPLS integration](#) for more information.

**Note:**

The **split-horizon-group** command is supported only on the 7210 SAS-Mxp.

- **ingress-replication-bum-label**

When this command is enabled, it allows the PE to advertise a label for BUM traffic (inclusive multicast routes) that is different from the label advertised for unicast traffic (with the MAC/IP routes). This is useful to avoid potential transient packet duplication in all-active multi-homing.

**Note:**

On the 7210 SAS, because all-active multi-homing is not supported by default, the **ingress-replication-bum-label** command is disabled. The user has the option to enable this command.

In addition to the preceding options, the following **bgp-evpn** commands are also available for EVPN-MPLS services:

- **[no] mac-advertisement**
- **mac-duplication** and settings

When EVPN-MPLS is established among some PEs in the network, EVPN unicast and multicast "bindings" to the remote EVPN destinations are created on each PE. A specified ingress PE creates the following:

- 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 with a BUM label (only possible if the egress PE is configured with **ingress-replication-bum-label**)

These bindings, as well as the MACs learned on them, can be checked using the **show** commands in the following output example, where 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**. As a result, the device 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). The following is a sample configuration output.

Example: EVPN-MPLS configuration output

```
*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:	EvpnS	06/11/15 19:59:57


```

1          00:ca:fe:ca:fe:72 192.0.2.70:262140      EvpnS      06/11/15 19:59:57
          eMpls:
          192.0.2.72:262141
-----
No. of MAC Entries: 3
-----
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
=====

```

4.2.2.1 EVPN and VPLS integration

In accordance with *draft-ietf-bess-evpn-vpls-seamless-integ*, the 7210 SAS EVPN implementation supports the integration of EVPN-MPLS and VPLS to the same network within the same service. Because EVPN is not deployed in greenfield deployments, this feature is useful for facilitating the integration between both technologies and for migrating VPLS services to EVPN-MPLS.


The following behavior enables the integration of EVPN and SDP-bindings in the same VPLS network:

- 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 applies to spoke-SDPs (manual).
 - If an EVPN endpoint to a specified far-end exists and a spoke-SDP establishment is attempted, the spoke-SDP is set up but is kept down with an operational flag, indicating that there is an EVPN route to the same far-end.
 - If an spoke-SDP exists and a valid or used EVPN route arrives, the EVPN endpoint is set up and the spoke-SDP is brought down with an operational flag indicating that there is an EVPN route to the same far-end.
 - In the case of an SDP-binding and EVPN endpoint to different far-end IPs on the same remote PE, both links are up. This may occur if the SDP-binding is terminated in an IPv4 address that is different from the system address where the EVPN endpoint is terminated.
- Users can add spoke-SDPs and all the EVPN-MPLS endpoints in the same split-horizon group (SHG):
 - On the 7210 SAS-Mxp, the following CLI command is added under the **bgp-evpn>mpls** context so that the EVPN-MPLS endpoints can be added to an SHG: **bgp-evpn mpls [no] split-horizon-group group-name**.
 - On the 7210 SAS-Mxp, the **bgp-evpn mpls split-horizon-group** command must reference a user-configured SHG. User-configured SHGs can be configured within the service context. The same *group-name* can be associated with SAPs, spoke-SDPs, pw-template-bindings, and EVPN-MPLS endpoints.
 - On the 7210 SAS-Mxp, if the **bgp-evpn mpls split-horizon-group** command is not used, the default SHG (that contains all the EVPN endpoints) is still used but cannot be referred to using SAPs/spoke-SDPs.
 - On the 7210 SAS-Sx/S 1/10GE, the **use-evpn-default-shg** CLI parameter is available during spoke-SDP creation. It allows the spoke-SDP to be added to the default EVPN binding SHG group.
 - On the 7210 SAS-Sx/S 1/10GE, by default, all EVPN bindings are part of the default SHG that contains all the EVPN binding endpoints instantiated in the service.
 - On the 7210 SAS-Sx/S 1/10GE, a SAP cannot be added to the default EVPN SHG group.

- On the 7210 SAS-Mxp, the system disables the advertisement of MACs learned on spoke-SDPs or SAPs that are part of an EVPN split-horizon group. On the 7210 SAS-Sx/S 1/10GE, the system disables the advertisement of MACs learned on spoke-SDPs that are part of an EVPN split-horizon group:
 - When the SAPs or spoke-SDPs (manual) on the 7210 SAS-Mxp or the spoke-SDPs (manual) on the 7210 SAS-Sx/S 1/10GE are configured within the same SHG as the EVPN endpoints, MAC addresses are still learned on them, but they are not advertised in EVPN.
 - On the 7210 SAS-Mxp, the preceding statement is also true if proxy-ARP/proxy-ND is enabled and an IP-to-MAC pair is learned on a SAP that belongs to the EVPN split-horizon group.
 - On the 7210 SAS-Mxp, the SAPs added to an EVPN split-horizon group should not be part of any EVPN multi-homed ES. If that occurs, the PE still advertises the AD per-EVI route for the SAP, attracting EVPN traffic that could not possibly be forwarded to that SAP.

The following figure shows an example of EVPN-VPLS integration.

Figure 41: EVPN-VPLS integration



Example

The following is a sample configuration of the 7210 SAS-Mxp for PE1, PE5, and PE2 from the EVPN-VPLS integration example in the preceding figure.

```
*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
spoke-sdp 12:1 create
```

```

exit
spoke-sdp 13:1 split-horizon-group "SHG-1" create
exit
spoke-sdp 14:1 split-horizon-group "SHG-1" create
exit
spoke-sdp 15:1 split-horizon-group "SHG-1" create
exit
sap 1/1/1:1 create
exit

*A:PE5>config>service# info
-----
split-horizon-group "SHG-1" create
vpls 1 customer 1 create
  bgp
    route-target target:65000:1
  spoke-sdp 52:1 create
  exit
  spoke-sdp 51:1 split-horizon-group "SHG-1" create
  exit
  spoke-sdp 53:1 split-horizon-group "SHG-1" create
  exit
  spoke-sdp 54:1 split-horizon-group "SHG-1" create
  exit

*A:PE2>config>service# info
-----
vpls 1 customer 1 create
  end-point CORE create
    no suppress-standby-signaling
  spoke-sdp 21:1 end-point CORE
    precedence primary
  spoke-sdp 25:1 end-point CORE

```

Example

The following is a sample configuration of the 7210 SAS-S 1/10GE for PE1, PE5, and PE2 from the EVPN-VPLS integration example in [Figure 41: EVPN-VPLS integration](#). This example uses the **use-evpn-default-shg** parameter to include spoke SDPs in the default EVPN SHG.

```

*A:PE1>config>service# info
-----
vpls 1 customer 1 create
  bgp
    route-target target:65000:1
  bgp-evpn
    evi 1
    mpls
      no shutdown
  spoke-sdp 12:1 create
  exit
  spoke-sdp 13:1 use-evpn-default-shg create
  exit
  spoke-sdp 14:1 use-evpn-default-shg create
  exit
  spoke-sdp 15:1 use-evpn-default-shg create
  exit
  sap 1/1/1:1 create
  exit

*A:PE5>config>service# info
-----
split-horizon-group "SHG-1" create

```

```

vpls 1 customer 1 create
  bgp
    route-target target:65000:1
  spoke-sdp 52:1 create
  exit
  spoke-sdp 51:1 split-horizon-group "SHG-1" create
  exit
  spoke-sdp 53:1 split-horizon-group "SHG-1" create
  exit
  spoke-sdp 54:1 split-horizon-group "SHG-1" create
  exit
*A:PE2>config>service# info
-----
vpls 1 customer 1 create
  end-point CORE create
    no suppress-standby-signaling
  spoke-sdp 21:1 end-point CORE
    precedence primary
  spoke-sdp 25:1 end-point CORE

```

The following applies to the configuration described in the preceding examples:

- PE1, PE3, and PE4 have BGP-EVPN enabled in VPLS-1. PE2 has active/standby spoke SDPs to PE1 and PE5. In this configuration:
 - PE1, PE3, and PE4 have manual spoke SDPs, but they are kept operationally down as long as there are EVPN endpoints active among them
 - manual spoke SDPs on PE1, PE3, and PE4 and EVPN endpoints are instantiated within the same SHG, for example, the default SHG
 - manual spoke SDPs from PE1 and PE5 to PE2 are not part of the default SHG
- EVPN MAC advertisements
 - for spoke SDPs and EVPN in the same SHG, MACs learned locally on a spoke SDP are not advertised in EVPN
- BUM traffic operation on PE1
 - when CE1 sends BUM traffic, PE1 floods to all the active bindings
 - when CE2 sends BUM traffic, PE2 sends it to PE1 (active spoke SDP) and PE1 floods to all the bindings and SAPs
 - when CE5 sends BUM traffic, PE5 floods to the three EVPN PEs. PE1 floods to the active spoke SDP and SAPs, never to the EVPN PEs because they are part of the same SHG.

4.2.2.2 Auto-Derived RD in services with multiple BGP families

A single RD is used per service and not per BGP family or protocol. On the 7210 SAS, BGP-AD is not supported with BGP-EVPN.



Note:

On the 7210 SAS, to prevent auto-derived RD in services from using BGP-AD information, Nokia recommends that the user should not configure the **bgp-ad>vpls-id** command.

The following rules apply.

- The VPLS RD is selected based on the following precedence:
 - manual RD always takes precedence when configured

- if no manual-rd configuration exists, the RD is derived from the **bgp-evpn>evi**
- if no manual-rd or **bgp-evpn>evi** configuration exists, there is no RD and the service fails
- The selected RD (see the preceding selection criteria) is displayed by the Oper Route Dist field of the **show service id bgp** command.
- The service supports dynamic RD changes; for example, the manual RD can be updated dynamically, even if it is currently in use as the service RD.

**Note:**

When the RD changes, the active routes for that VPLS are withdrawn and re-advertised with the new RD.

- If one of the mechanisms to derive the RD for a specified service is removed from the configuration, the system selects a new RD based on the preceding rules. For example, if the manual RD is removed from the configuration, the routes are withdrawn, the new RD is selected from the EVI, and the routes re-advertised with the new RD. See [Auto-Derived RD in services with multiple BGP families](#) for more information about rules governing the RD selection.

**Note:**

The reconfiguration fails if the new RD already exists in a different VPLS or Epipe service.

4.2.3 EVPN multi-homing in VPLS services

EVPN multi-homing implementation is based on the concept of the Ethernet Segment (ES). An ES is a logical structure that can be defined in one or more PEs and identifies the CE (or access network) multi-homed to the EVPN PEs. An ES is associated with port or LAG objects and shared by all the services defined on those objects. On the 7210 SAS, only the following service objects are allowed to be configured as an ES: port and LAG.

Each ES has a unique Ethernet Segment Identifier (ESI) that is 10 bytes long and is manually configured in the router.

**Note:**

Because the **esi** command is advertised in the control plane to all the PEs in the EVPN network, it is 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 ES with esi = 0 (single-homed ESs are not explicitly configured).

4.2.4 EVPN all-active multi-homing

**Note:**

The 7210 SAS supports only single-active multi-homing. All-active multi-homing (with aliasing) is not supported on any 7210 SAS platform described in this document. References to all-active multi-homing (and aliasing) are only included in this section for completeness of feature description and are not intended to imply support on the 7210 SAS. See the *7210 SAS Software Release Notes 26.x.Rx*, part number 3HE 27852 000x TQZZA, for information about features supported in each load of the Release 26.x.Rx software.

In accordance with RFC 7432, all-active multi-homing is only supported on access LAG SAPs, and it is mandatory to configure the CE with a LAG to avoid duplicated packets to the network. LACP is optional.

When PE3 installs MAC1 in the Forwarding Database (FDB), it associates MAC1 not only with the advertising PE (PE1), but also with all the PEs advertising the same esi (ESI2) for the service. In this example, PE1 and PE2 advertise an AD per-EVI route for ESI2; therefore, PE3 installs the two next-hops associated with MAC1.

To enable aliasing, configure ECMP greater than 1 in the **bgp-evpn>mpls** context.

4.2.4.1 All-active multi-homing procedures


This section describes procedures implemented in 7210 SAS to provide all-active multi-homing for a specified ES.

4.2.4.1.1 Designated Forwarder election

Using Designated Forwarder (DF) election in EVPN all-active multi-homing prevents duplicate packets on the multi-homed CE. The DF election procedure elects one DF PE per ESI per service; the rest of the PEs are non-DF for the ESI and service. Only the DF forwards BUM traffic from the EVPN network toward the ES SAPs (the multi-homed CE). The non-DF PEs do not forward BUM traffic to the local ES SAPs.

The following figure shows the need for DF election in all-active multi-homing.

Figure 42: DF election



Note:


BUM traffic from the CE to the network and known unicast traffic in any direction is allowed on both the DF and non-DF PEs.

4.2.4.1.2 Split-horizon

The EVPN split-horizon procedure ensures that BUM traffic originated by the multi-homed PE and sent from the non-DF to the DF is not replicated back to the CE in the form of echoed packets. To avoid echoed packets, the non-DF (PE1) sends all the BUM packets to the DF (PE2) with an indication of the source ES. That indication is the ESI Label (ESI2 in the following figure), previously signaled by PE2 in the AD per-ESI route for the ES. When it receives an EVPN packet (after the EVPN label lookup), PE2 finds the ESI label that identifies its local ES ESI2. The BUM packet is replicated to other local CEs but not to the ESI2 SAP.

The following figure shows the EVPN split-horizon concept for all-active multi-homing.


Figure 43: Split-horizon



4.2.4.1.3 Aliasing

The following figure shows the EVPN aliasing procedure for all-active multi-homing. Because CE2 is multi-homed to PE1 and PE2 using an all-active ES, "aliasing" is the procedure by which PE3 can load-balance the known unicast traffic between PE1 and PE2, even if the destination MAC address is only advertised by PE1.

Figure 44: Aliasing



4.2.4.2 All-active multi-homing service model

The following examples show output of the PE1 and PE2 configurations that provides all-active multi-homing to the CE2 shown in [Figure 44: Aliasing](#).

Example: PE1 configuration output

```
*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 10.1.1.1:0
ethernet-segment "ESI2" create
esi 01:12:12:12:12:12:12:12:12
multi-homing all-active
service-carving
mode auto
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
ingress-replication-bum-label
sap lag-1:1 create
exit

```

Example: PE2 configuration output

```

*A:PE1>config>lag(1)# info
-----
mode access
encap-type dot1q
port 1/1/1
lacp active administrative-key 1 system-id 00:00:00:00:00:22
no shutdown

*A:PE1>config>service>system>bgp-evpn# info
-----
route-distinguisher 10.1.1.1:0
ethernet-segment "ESI2" create
esi 01:12:12:12:12:12:12:12:12
multi-homing all-active
service-carving
mode auto
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
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 following considerations apply when the all-active multi-homing procedure is enabled:

- The **ethernet-segment** command must be configured with a name and a 10-byte **esi** using the **config service system bgp-evpn ethernet-segment es_name create** and **config service system bgp-evpn ethernet-segment esi value** commands.
- When configuring the **esi**, the system enforces that the 6 high-order octets after the type are not zero, which ensures that the auto-derived route-target for the ES route is not zero). In addition, the entire ESI value must be unique in the system.
- Only a LAG can be associated with the all-active ES. LAG is used exclusively for EVPN multi-homing. Other LAG ports in the system can continue to 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 can respond as though it is connected to the same system.
- Only one SAP per service can be part of the same **ethernet-segment**.

4.2.4.3 ES discovery and DF election procedures

The ES discovery and DF election are implemented in three logical steps, as shown in the following figure.

Figure 45: ES discovery and DF election



4.2.4.3.1 Step 1 — ES advertisement and discovery

ES ESI-1 is configured with all the required parameters, as described in [All-active multi-homing service model](#). When **ethernet-segment no shutdown** is executed, PE1 and PE2 advertise an ES route for ESI-1. They both include the route-target auto-derived from the MAC portion of the configured ESI. If the route-target address family is configured in the network, this 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 announces the ES 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 what services (EVIs) are associated with the ESI. These routes are used by PE3 for its aliasing and backup procedures.

4.2.4.3.2 Step 2 — DF election

When ES routes exchange between PE1 and PE2 is complete, both run the DF election for all the services in the **ethernet-segment**.

PE1 and PE2 elect a Designated Forwarder (DF) per ESI service. The default DF election mechanism in the SR OS is **service-carving** (as per RFC 7432). The following applies when the mechanism is enabled on a specified PE:

- An ordered list of PE IPs where ESI-1 resides is built. The IPs are derived from the origin IP fields of all the ES routes received for ESI-1, as well as the local system address. The lowest IP is considered ordinal "0" in the list.
- The local IP can only be considered a "candidate" after successful **ethernet-segment no shutdown** for a specified service.



Note:

The remote PE IPs must be present in the local PE RTM so that they can participate in the DF election.

- A PE only considers a specified remote IP address as candidate for the DF election algorithm for a specified service if, as well as the ES route, the corresponding AD routes per-ESI and per-EVI for that PE have been received and properly activated.
- All remote PEs that receive 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 route per-ESI with the single-active flag set or the local ESI configuration is single-active, the ESI behaves as single-active.
- An **es-activation-timer** can be configured at the **redundancy>bgp-evpn-multi-homing>es-activation-timer** level or at the **service>system>bgp-evpn>eth-seg>es-activation-timer** level. This timer, which is 3 seconds by default, delays the transition from non-DF to DF for a specified service after the DF election has run:
 - This use of the **es-activation-timer** is different from zero and minimizes the risks of loops and packet duplication due to "transient" multiple DFs.
 - The same **es-activation-timer** should be configured in all PEs that are part of the same ESI. It is up to the user to configure either a long timer to minimize the risks of loops/duplication or even **es-activation-timer=0** to speed up the convergence for non-DF to DF transitions. When the user configures a specific value, the value configured at the ES 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 route per-EVI (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 ES and running the DF algorithm. The boot-timer is configured at the system level, using the **config redundancy bgp-evpn-multi-homing boot-timer** command, and should use a value that is long enough to allow the node (with any cards, if available) to boot up and BGP sessions to come up before exchanging ES routes and running the DF election for each EVI/ISID:
 - 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 it needs to.
 - The following **show** command displays the configured boot-timer and the remaining timer, if the system is still in boot-stage.

```
A:PE1# show redundancy bgp-evpn-multi-homing
```

```
=====
```

```
Redundancy BGP EVPN Multi-homing Information
```

```
=====
```

```
Boot-Timer           : 10 secs
Boot-Timer Remaining : 0 secs
ES Activation Timer   : 3 secs
=====
```

- When **service-carving mode auto** is configured (default mode), the DF election algorithm runs the function $[V(\text{evi}) \bmod N(\text{peers}) = i(\text{ordinal})]$ to identify the DF for a specified service and ESI, as described in the following example:
 - As shown in [Figure 45: ES discovery and DF election](#), PE1 and PE2 are configured with ESI-1. Given that $V(10) \bmod N(2) = 0$, PE1 are elected DF for VPLS-10 (because its IP address is lower than PE2's and it is the first PE in the candidate list).



Note:

The algorithm uses the configured **evi** in the service and not the *service-id*. The **evi** for a service must match in all PEs that are part of the ESI. This guarantees that the election algorithm is consistent across all PEs of the ESI. The **evi** must be always configured in a service with SAPs that are created in an ES.

- A **service-carving** command is supported to manually configure the EVI identifiers for which the PE is primary: **service-carving mode manual/manual evi start-evi to end-evi**. The following considerations apply:
 - The system is the PE forwarding/multicasting traffic for the **evi** identifiers included in the configuration. The PE is secondary (non-DF) for the non-specified **evi** identifiers.
 - If a range is configured but **service-carving** is not **mode manual**, the range has no effect.

- Only two PEs are supported when **service-carving mode manual** is configured. If manual mode is configured for a third PE for an ESI, the two non-primary PEs remain non-DF regardless of the primary status.
- For example, as shown in [Figure 45: 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**, PE1 is the primary PE for service VPLS 10 and PE2 the secondary PE.
- If **service-carving** is disabled, the lowest originator IP wins the election for a specified service and ESI. Use the **config service system bgp-evpn eth-seg service-carving mode off** command to disable service-carving.

The following sample configuration output shows the **ethernet-segment** configuration and DF status for all EVIs configured in the **ethernet-segment**.

```
*A:Dut-B# /show service system bgp-evpn ethernet-segment name "eslag1" all
=====
Service Ethernet Segment
=====
Name                : eslag1
Admin State         : Enabled           Oper State          : Up
ESI                 : 00:bc:01:00:00:00:00:00:01
Multi-homing        : allActive         Oper Multi-homing    : allActive
Lag Id              : 1
ES Activation Timer  : 3 secs (default)
Exp/Imp Route-Target : target:bc:01:00:00:00:00
Svc Carving          : auto
ES SHG Label        : 131070
=====
EVI Information
=====
EVI          SvcId          Actv Timer Rem    DF
-----
1             1              0                no
-----
Number of entries: 1
=====
DF Candidate list
-----
EVI          DF Address
-----
1             10.20.1.2
1             10.20.1.3
-----
Number of entries: 2
-----
```

4.2.4.3.3 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 datapath to remove the LAG SAP (associated with the ESI) from the default flooding list for 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 sample configuration output, 00:ca:ca:ba:ce:03 is learned on both PE1 and PE2 access LAG (on ESI-1). However, PE1 learns the MAC as "Learned" whereas PE2 learns it as "Evpn". This is because CE2 hashes the traffic for that source MAC to PE1. And PE2 learns the MAC through EVPN but associates the MAC to the ESI SAP, because the MAC 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 BUM packets for **evi 1**, the packets are sent including the ESI label at the bottom of the stack (in both directions). The ESI label advertised by each PE for ESI-1 can be displayed using the following command.

```
*A:PE1# show service system bgp-evpn ethernet-segment name "ESI-1"
=====
Service Ethernet Segment
=====
```

Name	: ESI-1	Oper State	: Up
Admin State	: Up		
ESI	: 01:00:00:00:00:71:00:00:00:01		
Multi-homing	: allActive	Oper Multi-homing	: allActive
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
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
=====

```

4.2.4.4 Aliasing

As shown in the example in [Figure 45: ES discovery and DF election](#), if the service configuration on PE3 has ECMP > 1, PE3 adds PE1 and PE2 to the list of next-hops for ESI-1. As soon as PE3 receives a MAC for ESI-1, it starts load-balancing between PE1 and PE2 the flows to the remote ESI CE.

Example: Configuration output for the FDB in PE3

The following is a sample configuration output that shows the FDB in PE3.



Note:

MAC 00:ca:ca:ba:ce:03 is associated with the **ethernet-segment** eES:01:00:00:00:00:71:00:00:00:01 (esi configured on PE1 and PE2 for ESI-1).

```

*A:PE3# show service id 1 fdb detail
=====
Forwarding Database, Service 1
=====
ServId  MAC                Source-Identifier      Type      Last Change
-----
1       00:ca:ca:ba:ce:03  eES:                  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
=====

```

Example: Configuration output for all EVPN-MPLS destination bindings on PE3

The following is a sample configuration output that shows all the EVPN-MPLS destination bindings on PE3, including the ES destination bindings.

**Note:**

The **ethernet-segment** eES:01:00:00:00:00:71:00:00:00:01 is resolved to PE1 and PE2 addresses.

```
*A:Dut-B# /show service id 1 evpn-mpls
=====
BGP EVPN-MPLS Dest
=====
TEP Address      Egr Label      Num. MACs      Mcast          Last Change
      Transport
-----
10.20.1.3        131069         0              Yes            02/02/2014 15:29:40
                  rsvp
10.20.1.4        131069         0              Yes            02/02/2014 15:29:33
                  rsvp
10.20.1.5        131059         0              Yes            02/02/2014 15:29:42
                  rsvp
-----
Number of entries : 3
-----
=====
BGP EVPN-MPLS Ethernet Segment Dest
=====
Eth SegId              Num. Macs      Last Change
-----
00:de:01:00:00:00:00:00:01  1            02/02/2014 15:47:04
-----
Number of entries: 1
-----
```

Example: PE3 configuration


PE3 performs aliasing for all the MACs associated with that ESI. This is possible because PE1 is configured with ECMP parameter >1. The following is a sample configuration output.

```
*A:PE3>config>service>vpls# info
-----
      bgp
      exit
      bgp-evpn
      evi 1
      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
```

4.2.4.5 Network failures and convergence for all-active multi-homing

The following figure shows the behavior on the remote PEs (PE3) when there is an **ethernet-segment** failure.

Figure 46: All-active multi-homing ES failure



The following steps describe the unicast traffic behavior on PE3:

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. In case of a failure between CE2 and PE2, PE2 withdraws its set of Ethernet AD routes and ES route, and PE3 forwards traffic destined for CE2 to PE1 only. PE3 does not need to wait for the withdrawal of the individual MAC.
3. The same handling is used if the failure was at PE1.
4. If after step2, PE2 withdraws its MAC advertisement route, PE3 treats traffic to MAC DA = CE2 as unknown unicast, unless PE1 has previously advertised the MAC.


For BUM traffic, the following events trigger a DF election on a PE, and only the DF forwards BUM traffic after the **esi-activation-timer** expires (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/withdraw
- new AD-EVI route update/withdraw
- local ES port/SAP/service shutdown
- service carving range change (affecting the EVI)
- multi-homing mode change (single/all active to all/single-active)

4.2.4.6 Logical failures on ESs and black holes

Specific “failure scenarios” in the network can trigger effects. The following figure shows some of these scenarios.

Figure 47: Black hole caused by SAP/SVC shutdown



If an individual VPLS service is **shutdown** in PE1 (the example is also valid for PE2), the corresponding LAG SAP goes operationally down. This event triggers the withdrawal of the AD per-EVI route for that 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 black-holing the traffic that CE2 “hashes” to the link to PE1. Because traffic sent from CE2 to PE2 or traffic from the rest of the CEs to CE2 is unaffected, the 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 the **bgp-evpn mpls shutdown** command is executed, the SAP associated with the ES goes operationally down (**StandbyforMHPprotocol**). If no other SAPs or SDP-bindings are configured in the service, the service also goes operationally down. However, if other SAPs or SDP-bindings are present, the service remains operationally up.

4.2.4.7 Transient issues because of MAC route delays

The following figure shows scenarios that may cause potential transient issues in the network.

Figure 48: Transient issues caused by “slow” MAC learning



In the preceding figure, the scenario on the left shows an example of transient packet duplication caused by delay in PE3 to learn MAC1.

In an all-active multi-homing scenario, if a specified MAC address (for example, MAC1), is not yet learned 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.

Configuring **ingress-replication-bum-label** in PE1 and PE2 resolves the issue. PE1 and PE2 know that the received packet is an unknown unicast packet; consequently, the NDF (PE1) does not send the packets to the CE, which prevents transient and duplication.

In the preceding figure, the scenario on the right shows an example of transient black hole caused by delay in PE1 to learn MAC1.

In an all-active multi-homing scenario, MAC1 is known in PE3 and aliasing is applied to MAC1. However, MAC1 is not yet known in PE1, which is the NDF for the ES. If PE3 hashing picks up PE1 as the destination of the aliased MAC1, the packets are blackholed. To resolve this issue, unknown unicast traffic that arrives with a unicast label should not be blocked on the NDF. If PE1 and PE2 are configured using **ingress-replication-bum-label**, PE3 sends unknown unicast packets with a BUM label and known unicast with a unicast label. In the latter case, PE1 considers it safe to forward the frame to the CE, even if it is unknown unicast.



Note:

This is a transient issue that is resolved as soon as MAC1 is learned in PE1 and the frames are forwarded as known unicast.

4.2.5 EVPN single-active multi-homing

The 7210 SAS SR OS supports only single-active multi-homing on access LAG SAPs and regular SAPs for a specified VPLS service. For LAG SAPs, the CE is configured with a different LAG to each PE in the ES (in contrast to a single LAG in an all-active multi-homing).

The following SR OS procedures support EVPN single-active multi-homing for a specified ES:

- **DF election**


The DF election in single-active multi-homing determines the forwarding for BUM traffic from the EVPN network to the ES CE. DF election also determines the forwarding of any traffic (unicast or BUM) in any direction (to or from the CE).

- **backup PE**

In single-active multi-homing, the remote PEs do not perform aliasing to the PEs in the ES. The remote PEs identify the DF based on the MAC routes and send the unicast flows for the ES to the PE in the DF. The remote PEs also program a backup PE as an alternative next-hop for the remote ESI in case of failure. This is in accordance with the Backup PE procedure, defined in RFC 7432.

The following figure shows an example backup PE for PE3.

Figure 49: Backup PE



al_0716

4.2.5.1 Single-active multi-homing service model

The following example shows a PE1 configuration that provides single-active multi-homing to CE2, as shown in [Figure 49: Backup PE](#).

Example: PE1 configuration

```
*A:PE1>config>service>system>bgp-evpn# info
-----
route-distinguisher 10.1.1.1:0
ethernet-segment "ESI2" create
esi 01:12:12:12:12:12:12:12:12
multi-homing single-active no-esi-label
service-carving
mode auto
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 single-active multihoming"
bgp
bgp-evpn
evi 10
mpls
no shutdown
auto-bind-tunnel resolution any
lag 1:1 create
exit
```

The PE2 example configuration for this scenario is as follows.

Example: PE2 configuration

```
*A:PE1>config>service>system>bgp-evpn# info
-----
route-distinguisher 10.1.1.1:0
ethernet-segment "ESI2" create
esi 01:12:12:12:12:12:12:12:12:12:12:12:12:12:12:12
multi-homing single-active no-esi-label
service-carving
lag 2
no shutdown

*A:PE1>config>redundancy>evpn-multi-homing# info
-----
boot-timer 120
es-activation-timer 10

*A:PE1>config>service>vpls# info
-----
description "evpn-mpls-service with single-active multihoming"
bgp
bgp-evpn
evi 10
mpls
no shutdown
auto-bind-tunnel resolution any
lag 2:1 create
exit
```

In single-active multi-homing, the non-DF PEs for a specified ESI block unicast and BUM traffic in both directions (upstream and downstream) on the object associated with the ESI. Otherwise, single-active multi-homing is similar to all-active multi-homing with the following differences:

- The **ethernet-segment** is configured for single-active: **service>system>bgp-evpn>eth-seg>multi-homing single-active**.
- The advertisement of the ESI-label in an AD per-ESI is optional in standards for **single-active** ESs. Use the **service system bgp-evpn eth-seg multi-homing single-active no-esi-label** command to control the ESI label advertisement. By default on the 7210 SAS, the ESI label is not used for single-active ESs, and there is no option available to enable the use of the ESI label.



Note:

The 7210 SAS ignores the ESI label received from an EVPN peer, which means that BUM traffic sent by the 7210 SAS to a peer DF node is always sent without the ESI label advertised by the DF. On the 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE, only **multi-homing single-active no-esi-label** can be configured.

- For single-active multi-homing, the ES can be associated with a **port** or **lag-id**, as shown in [Figure 49: Backup PE](#), where:
 - **port** is used for single-active SAP redundancy without the need for LAG
 - **lag** is used for single-active LAG redundancy



Note:

For a LAG configured with single-active homing, the LAG parameters **key**, **system-id**, and **system-priority** must be different on the PEs that are part of the ES.

- For single-active multi-homing, when the PE is non-DF for the service, the SAPs on the **ethernet-segment** are down and show **StandByForMHPProtocol** as the reason.
- From a service perspective, single-active multi-homing can provide redundancy to CEs (MHD, Multi-Homed Devices) with the following setup:
 - **LAG with or without LACP**
In this case, the multi-homed ports on the CE are part of the different LAGs (a LAG per multi-homed PE is used in the CE).
 - **regular Ethernet 802.1q/ad ports**
In this case, the multi-homed ports on the CE/network are not part of any LAG.

4.2.5.2 ES and DF election procedures

In all-active multi-homing, the non-DF keeps the SAP up, although it removes it from the default flooding list. In the single-active multi-homing implementation, the non-DF brings the SAP operationally down. For more information, see [ES discovery and DF election procedures](#).

The following **show** command output is an example status of the single-active ESI-7413 in the non-DF.

```
*A:Dut-B# show service system bgp-evpn ethernet-segment name "eslag1"
=====
Service Ethernet Segment
=====
Name                : eslag1
Admin State         : Enabled           Oper State         : Up
ESI                 : 00:bc:01:00:00:00:00:00:01
Multi-homing        : singleActiveNoEsi* Oper Multi-homing : singleActive*
Lag Id              : 1
ES Activation Timer  : 3 secs (default)
Exp/Imp Route-Target : target:bc:01:00:00:00:00
Svc Carving         : auto
ES SHG Label        : N/A
=====
* indicates that the corresponding row element may have been truncated.
=====
*A:Dut-D# /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   03/13/2000 11:43:16
=====
DF Candidates                                Time Added
-----
10.20.1.4                                03/13/2000 12:00:30
10.20.1.5                                03/13/2000 11:43:16
-----
Number of entries: 2
=====
*A:Dut-D#
```

4.2.5.3 Backup PE function

In the example in [Figure 49: Backup PE](#), the remote PE3 imports 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 route per-ESI with the single-active flag set. MACs for a specified service and ESI are learned from a single PE, that is, the DF for that <ESI, EVI>.

The remote PE installs a single EVPN-MPLS destination (TEP, label) for a received MAC address and a backup next-hop to the PE for which the AD routes per-ESI and per-EVI are received. For example, in the following **show** command sample 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. This ES is resolved to PE(192.0.2.73), which is the DF on the ES.

```
*A:PE3# show service id 1 fdb detail
=====
Forwarding Database, Service 1
=====
```

ServId	MAC	Source-Identifier	Type 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:Dut-D# /show service id 1 evpn-mpls
=====
BGP EVPN-MPLS Dest
=====
```

TEP Address	Egr Label Transport	Num. MACs	Mcast	Last Change
10.20.1.2	131061 rsvp	0	Yes	03/13/2000 11:26:29
10.20.1.2	131062 rsvp	1	No	03/13/2000 12:10:04
10.20.1.5	131061 rsvp	0	Yes	03/13/2000 11:18:23

```
-----
Number of entries : 3
-----
=====
BGP EVPN-MPLS Ethernet Segment Dest
=====
```

Eth SegId	Num. Macs	Last Change
00:de:01:00:00:00:00:00:00:01	453	03/13/2000 12:10:14

```
-----
Number of entries: 1
-----
=====
```

```
*A:Dut-D# /show service id 1 evpn-mpls esi 00:de:01:00:00:00:00:00:00:01
```

BGP EVPN-MPLS Ethernet Segment Dest		
Eth SegId	Num. Macs	Last Change
00:de:01:00:00:00:00:00:01	453	03/13/2000 12:10:14
BGP EVPN-MPLS Dest TEP Info		
TEP Address	Egr Label Transport	Last Change

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/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 datapath. Upon receiving an AD per-ES/per-EVI route withdrawal for the ESI, it flushes the MACs associated with the ESI.




Note:

On the 7210 SAS, an ES can be multi-homed to up to two PEs.

4.2.5.4 Network failures and convergence for single-active multi-homing

The following figure shows an example of remote PE (PE3) behavior when there is an **ethernet-segment** failure.

Figure 50: Single-active multi-homing ES failure



The following steps list the behavior of the remote PE3 for unicast traffic:

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. PE3 does not need to wait for the withdrawal of the individual MAC, and immediately forwards the traffic destined for CE2 to PE1 (the backup PE) only.
3. After the (2) 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.

A DF election on PE1 is also triggered. A DF election is triggered by the same events as all-active multi-homing. In this case, the DF forwards traffic to CE2 when the **esi-activation-timer** expires; the timer is triggered when a transition from non-DF to DF occurs.


4.2.6 EVPN-VPWS for MPLS tunnels

This section describes the EVPN-VPWS for MPLS tunnels.

4.2.6.1 BGP-EVPN control plane for EVPN-VPWS

EVPN-VPWS uses route type 1 and route type 4; it does not use route types 2, 3, or 5. The following figure shows the encoding of the extensions required for the Ethernet A-D per-EVI routes. The encoding follows guidelines described in *draft-ietf-bess-evpn-vpws*.

Figure 51: EVPN-VPWS BGP extensions



sw0438

Assuming that the advertising PE has an access SAP 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:


- The Ethernet tag ID field is encoded with the value configured by the user in the **config>service>bgp-evpn>local-ac-name>eth-tag value** command.
- RD and MPLS label values are encoded as specified in RFC 7432.
- The ESI is 0.
- The route is sent along with an EVPN L2 attributes extended community, as specified in *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
- L2 MTU is encoded with **service-mtu** configured in the Epipe service

4.2.6.2 EVPN for MPLS tunnels in Epipe services (EVPN-VPWS)

BGP-EVPN can be enabled in Epipe services with either SAPs or spoke SDPs at the access, as shown in the following figure.

Figure 52: EVPN-MPLS VPWS



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, MAC/IP routes, or IP prefix routes.
- AD Ethernet 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 destination is created for the Epipe.

Example: Epipe 2 as a EVPN-VPWS service between PE2 and PE4

In the following configuration example, Epipe 2 is an EVPN-VPWS service between PE2 and PE4 (as shown in [Figure 52: EVPN-MPLS VPWS](#)):

```
PE2>config>service>epipe(2)#
-----
bgp-evpn
```

```

    evi 2
    mpls
        auto-bind-tunnel resolution any
        ecmp 2
        local-ac-name "AC-1"
        eth-tag 100
        remote-ac-name "AC-2"
        eth-tag 200
        no shutdown
    sap 1/1/1:1 create
PE4>config>service>epipe(2)#
-----
    bgp-evpn
    evi 2
    mpls
        auto-bind-tunnel resolution any
        local-ac-name "AC-2"
        eth-tag 200
        remote-ac-name "AC-1"
        eth-tag 100
        no shutdown
    sap 1/1/1:1

```

The following considerations apply to the example configuration:

- The **evi** 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 to which 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 force-vlan-vc-forwarding**
 - **mpls shutdown**
- The following **bgp-evpn** commands identify the local and remote attachment circuits, with the configured **eth-tags** encoded in the advertised and received AD Ethernet per-EVI routes:
 - **local-ac-name** *name*
 - **local-ac-name** *name* **eth-tag** *tag-value*; where tag-value [1..16777215]
 - **remote-ac-name** *name*
 - **remote-ac-name** *name* **eth-tag** *tag-value*; where tag-value [1..16777215]
 - Changes on remote Ethernet tags are allowed without shutting down **bgp-evpn mpls** or the Epipe service. The **local-ac eth-tag** value cannot be changed without **bgp-evpn mpls shutdown**.
 - Both local and remote Ethernet tags are mandatory to bring up the Epipe service.

EVPN-VPWS Epipes can also be configured with the following characteristics:

- Access attachment circuits can be SAPs or spoke SDPs. Only manually configured spoke SDPs are supported, **endpoints** are not supported. The **vc-switching** configuration is not supported on **bgp-evpn** enabled pipes.
- EVPN-VPWS Epipes support **control-word**.

When **bgp-evpn>mpls>control-word** 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 does not setup the EVPN destination. As a result, the service does 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 L2 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 setup 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 L2 MTU value is 0, the MTU is ignored.




Note: The 7210 SAS supports the **no service-mtu-check** command to disable the MTU checks in the dataplane to allow for interop with other PE routers. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide* for more information about the packet length used for service MTU enforcement on the 7210 SAS.

4.2.6.3 Using A/S PW and MC-LAG with EVPN-VPWS Epipes

The use of A/S PW (for access spoke SDPs) and MC-LAG (for access SAPs) provides an alternative redundant solution for EVPN-VPWS that do not use the EVPN multihoming procedures described in *draft-ietf-bess-evpn-vpws*. The following figure shows the use the mechanism in a single Epipe.

Figure 53: A/S PW and MC-LAG support on EVPN-VPWS



In the preceding figure, an 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). As 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 (in PE1 and PE2).

- 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, it 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, it advertises the AD EVI route for **eth-tag** 1.

- Users can configure MC-LAG for access SAPs using the example configuration of PE3 and PE4 shown in Figure 165. 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 the PE3 advertises the AD per-EVI route for **eth-tag** 2.

4.2.6.4 LAG-based LLF for EVPN-VPWS services

The 7210 SAS supports CE-to-CE fault propagation in EVPN-VPWS services by using LAG standby signaling, which can be LACP-based or power-off when the remote CE is using LAG. That is, when detecting a CE failure, an EVPN-VPWS PE withdraws the corresponding autodiscovery per-EVI route, which then triggers the remote PE to signal the fault to the connected CE using LAG standby signaling (either LACP-based or power-off).

The following figure shows an example of link loss forwarding for EVPN-VPWS.

Figure 54: Link loss forwarding for EVPN-VPWS



Example: PE1 configuration

```
*A:Dut>config> system> oper-group "llf-1" create
*A:Dut>config> system> oper-group> info detail
-----
    hold-time
      group-down nn
      group-up  mm
    exit
-----
```

```

A:PE1>config>lag(1)# info
-----
mode access
encap-type null
port 1/1/1
port 1/1/2
standby-signaling power-off
monitor-oper-group "llf-1"
no shutdown
-----

*A:PE1>config>service>epipe# info
-----
bgp
exit
bgp-evpn
    evi 1
        local-ac-name ac-1
        eth-tag 1
    exit
    remote-ac-name ac-2
    eth-tag 2
    exit
mpls
    oper-group "llf-1"
    auto-bind-tunnel
    resolution any
    exit
    no shutdown
exit
sap lag-1 create
    no shutdown
exit
no shutdown
-----

```

The following applies to the PE1 configuration:

- The EVPN Epipe service is configured on PE1 with a null LAG SAP and the operational group "llf-1" under **bgp-evpn>mpls**. This is the only member of operational group "llf-1".
- The operational group monitors the status of the BGP-EVPN instance in the Epipe service. The status of the BGP-EVPN instance is determined by the existence of an EVPN destination at the Epipe.
- In **access** mode and encap-type **null**, the LAG is configured with the **monitor-oper-group "llf-1"** command.

As shown in [Figure 54: Link loss forwarding for EVPN-VPWS](#), upon failure on CE2, the following events occur:

1. PE2 withdraws the EVPN route.
2. The EVPN destination is removed in PE1 and operational group "llf-1" also goes down.
3. Because lag-1 is monitoring "llf-1", the operational group that is becoming inactive triggers standby signaling on the LAG; that is, power-off or LACP out-of-sync signaling to the CE1.


When the SAP or port is down because of the LAG monitoring of the operational group, PE1 does not trigger an AD per-EVI route withdrawal, even if the SAP is brought operationally down.

4. After CE2 recovers and PE2 re-advertises the AD per-EVI route, PE1 creates the EVPN destination and operational group "llf-1" comes up. As a result, the monitoring LAG stops signaling standby and the LAG is brought up.

4.2.6.5 LAG or port standby signaling to the CE on non-DF EVPN PEs (single-active)

As described in [EVPN for MPLS tunnels](#), the EVPN single-active multihoming PEs elected as non-DF must notify their attached CEs, so the CE does not send traffic to the non-DF PE. This scenario is shown in the following figure.

Figure 55: LACP standby signaling from the non-DF



As shown in the preceding figure, the multihomed PEs are configured with multiple EVPN services that use ES-1. ES-1 and its associated LAG is configured as follows:

Example: ES-1 and associated LAG configuration

```
*A:Dut>config> system> oper-group "SA-1" create
*A:Dut>config> system> oper-group> info detail
-----
hold-time
group-down nn
group-up mm
exit
-----

*A:Dut>config>lag# info
-----
mode access
encap-type dot1q
port 1/1/3
lACP active administrative-key 32772 system-id 00:00:00:00:00:01 system-priority 1
monitor-oper-group SA-1
no shutdown
-----

*A:Dut>config>service>system>bgp-evpn# info
-----
ethernet-segment "toI" create
esi 14:13:12:08:00:00:00:00:08
```

```

service-carving
  mode off
exit
multi-homing single-active
oper-group SA-1
lag 2
no shutdown
exit
-----

```

When the operational group is configured on the ES and monitored on the associated LAG:

- The operational group status is driven by the ES DF status (defined by the number of DF SAPs or oper-up SAPs owned by the ES).
- The operational group goes down if all the SAPs in the ES go down (this happens in PE2 in [Figure 55: LACP standby signaling from the non-DF](#)). The ES operational group goes up when at least one SAP in the ES goes up.

As a result, if PE2 becomes non-DF on all SAPs in the ES, they all go operationally down, including the ES-1 operational group.

- Because LAG-1 is monitoring the operational group, when its status goes down, LAG-1 signals LAG standby state to the CE. The standby signaling can be configured as LACP or power-off.
- The ES and AD routes for the ES are not withdrawn because the router recognizes that the LAG becomes standby for the ES operational group.



Note:

- The **config>lag>monitor-oper-group name** command is supported for ports configured in either **access** or **hybrid** mode. Any **encap-type** can be used.
- With an Epipe configured on the CE to use a single LAG (with active/standby member links) per service to provide uplink redundancy, the 7210 SAS PEs must be configured with a LAG in **access** port mode, and the system ID and system priority of the LAG should be configured to the same value across all those PEs.
- With an Epipe or a VPLS configured on the CE to use multiple LAGs per service (with active/standby LAGs) to provide uplink redundancy, the 7210 SAS PEs must be configured with a LAG in either **access** port mode or **hybrid** port mode, and the system ID and system priority of the LAG should be configured to a different value across all those PEs.

Operational groups cannot be assigned to ESs that are configured as **all-active** or **service-carving mode auto**.

4.2.7 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 has to make a selection (for instance, the same MAC 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/MAC but different route keys are received, BGP hands the routes over to EVPN. EVPN selects the route based on the following tie-break order:

1. Conditional static MACs (local protected MACs)
 2. EVPN static MACs (remote protected MACs)
 3. Data plane learned MACs (regular learning on SAPs/SDP bindings)
 4. EVPN MACs with higher SEQ number
 5. Lowest IP (next-hop IP of the EVPN NLRI)
 6. Lowest Ethernet tag (that is zero for MPLS and may be different from zero for VXLAN)
 7. Lowest RD
- **BGP route selection for MAC routes with the same route key**
The following priority order is applied:
 1. EVPN static MACs (remote protected MACs)
 2. EVPN MACs with higher sequence number
 3. Regular BGP selection (local preference, AIGP metric, shortest AS-path, lowest IP)
 - **BGP route selection for remaining EVPN routes**
The regular BGP selection is followed.



Note: If BGP has to run an actual selection and a specific (otherwise valid) EVPN route "loses" to another EVPN route, the non-selected route is displayed by the **show router bgp routes evpn x detail** command with a "tie-breaker" reason.

4.3 General EVPN topics

This section provides information about general topics related to EVPN.



Note:
Hash labels (that is, the Flow Aware Transport label (RFC 6391)) are not supported with 7210 SAS EVPN VPLS services.

4.3.1 ARP and ND snooping and proxy support


VPLS services support proxy-Address Resolution Protocol (proxy-ARP) and proxy-Neighbor Discovery (proxy-ND) functions cannot be enabled or disabled per service. When enabled, the **config>service>system>evpn-proxy-arp-nd** command populates the corresponding proxy-ARP or proxy-ND table with IP-to-MAC entries learned from the following sources:

- EVPN-received IP-to-MAC entries
- user-configured static IP-to-MAC entries
- snooped dynamic IP-to-MAC entries (learned from ARP, GARP, or NA messages received on local SAPs; snooped dynamic IP-to-MAC entries on spoke-SDP bindings are not supported)

In addition, any ingress ARP or ND frame on a SAP are intercepted and processed. The system answers ARP requests and Neighbor Solicitation messages if the requested IP address is present in the proxy table.

The following figure shows an example proxy-ARP usage in an EVPN network. Proxy-ND functions in a similar way. The MAC address notation in the diagram is shortened for readability.

Figure 56: Proxy-ARP example usage in an EVPN network



al_0626

In the preceding figure, PE1 is configured as follows:

```
*A:Dut-B>config>service>system# info
-----
evpn-proxy-arp-nd
-----
*A:Dut-B>config>service>vpls# info
-----
description "Vpls 1 "
service-mtu 1400
split-horizon-group "vpls1" create
description "Default description for SHG vpls1"
exit
bgp
route-distinguisher auto-rd
route-target export target:100:1 import target:100:1
pw-template-binding 100
exit
exit
bgp-evpn
evi 1
mpls
split-horizon-group "vpls1"
ingress-replication-bum-label
auto-bind-tunnel
resolution-filter
ldp
```

```

        exit
        resolution filter
    exit
    no shutdown
exit
stp
    shutdown
exit
sap lag-1:1 create
    description "Default sap description for service id 1"
    no shutdown
exit
proxy-arp
    age-time 600
    send-refresh 200
    dup-detect window 3 num-moves 3 hold-down max anti-spoof-
mac 00:aa:aa:aa:aa:aa
    dynamic-arp-populate
    no shutdown
exit
no shutdown
-----
*A:Dut-B>config>service>vpls#

```

The preceding figure shows the following steps, assuming proxy-ARP is **no shutdown** on PE1 and PE2, and the tables are empty:

1. ISP-A sends ARP-request for 10.10.10.3.
2. PE1 learns the MAC 00:01 in the FDB as usual and advertises it in EVPN without any IP. Optionally if the MAC is configured as a Cstatic MAC, it is advertised as a protected MAC to other PEs with the sticky bit set.
3. The ARP-request is sent to the CPM, where it is handled as follows:
 - An ARP entry (IP 10.1/MAC 00:01) is populated into the proxy-ARP table.
 - EVPN advertises MAC 00:01 and IP 10.1 in EVPN with the same SEQ number and protected bit as the previous route-type 2 for MAC 00:01.
 - A GARP is also issued to other SAPs/SDP-bindings (assuming they are not in the same split-horizon group as the source). If the **garp-flood-evpn** command is enabled, the GARP message is also sent to the EVPN network.
 - The original ARP-request can still be flooded to the EVPN or not based on the **unknown-arp-request-flood-evpn** command.
4. Assuming PE1 was configured with **unknown-arp-request-flood-evpn**, the ARP-request is flooded to PE2 and delivered to ISP-B. ISP-B replies with its MAC in the ARP-reply. The ARP-reply is finally delivered to ISP-A.
5. PE2 learns MAC 00:01 in the FDB and the entry 10.1'00:01 in the proxy-ARP table, based on the EVPN advertisements.
6. When ISP-B replies with its MAC in the ARP-reply, the MAC is handled as follows:
 - MAC 00:03 is learned in FDB at PE2 and advertised in EVPN.
 - MAC 00:03 and IP 10.3 are learned in the proxy-ARP table and advertised in EVPN with the same SEQ number as the previous MAC route.
 - ARP-reply is unicasted to MAC 00:01.

7. EVPN advertisements are used to populate PE1's FDB (MAC 00:03) and proxy-ARP (IP 10.3 to MAC 00:03) tables as mentioned in 5.

From this point onward, the PEs reply to any ARP-request for 00:01 or 00:03 without the need for flooding the message in the EVPN network. By replying to known ARP-requests and Neighbor Solicitations, the PEs help to significantly reduce the flooding in the network.

Use the following commands to customize proxy-ARP/proxy-ND behavior:

- **dynamic-arp-populate** and **dynamic-nd-populate**

These commands enable the addition of dynamic entries to the proxy-ARP or proxy-ND table (disabled by default). When executed, the system populates proxy-ARP/proxy-ND entries from snooped GARP/ARP/NA messages on SAPs/SDP-bindings, in addition to the entries coming from EVPN (if EVPN is enabled). These entries are shown as dynamic.

- **static ipv4-address mac-address**, **static ipv4-address mac-address**, and **static ipv6-address mac-address {host | router}**

These commands configure static entries to be added to the table.



Note:

A static IP-to-MAC entry requires the addition of the MAC address to the FDB as either learned or CStatic (conditional static mac) to become active (*Status active*).

- **age-time seconds**

This command specifies the aging timer per proxy-ARP/proxy-ND entry. When the aging expires, the entry is flushed. The age is reset when a new ARP/GARP/NA for the same IP-to-MAC is received.

- **send-refresh seconds**

If this command is enabled, the system sends ARP-request or Neighbor Solicitation (NS) messages at the configured time, which enables the owner of the IP to reply and, therefore, refresh its IP-to-MAC (proxy-ARP entry) and MAC (FDB entry).

- **table-size table-size**

This command enables the user to limit the number of entries learned on a specified service. By default, the table-size limit is 250.

Flooding unknown ARP-requests, NS messages, or unsolicited GARPs and NA messages in an EVPN network can be configured using the following commands:

- **proxy-arp [no] unknown-arp-request-flood-evpn**
- **proxy-arp [no] garp-flood-evpn**
- **proxy-nd [no] unknown-ns-flood-evpn**
- **proxy-nd [no] host-unsolicited-na-flood-evpn**
- **proxy-nd [no] router-unsolicited-na-flood-evpn**

- **dup-detect [anti-spoof-mac mac-address] window minutes num-moves count hold-down minutes | max**

This command enables a mechanism that detects duplicate IPs and ARP/ND spoofing attacks. The following is a summary of the **dup-detect** command mechanism:

- Attempts (relevant to dynamic and EVPN entry types) to add the same IP (different MAC) are monitored for **window minutes** value and when the **count** value is reached within the configured

window, the proxy-ARP/proxy-ND entry for the IP is suspected and marked as duplicate. An alarm is also triggered.

- The condition is cleared when **hold-down** time expires (*max* does not expire) or a **clear** command is issued.
- If the **anti-spoof-mac** command is configured, the proxy-ARP or proxy-ND offending entry's MAC is replaced by the configured *mac-address* and advertised in an unsolicited GARP/NA for local SAP or SDP-bindings and in EVPN to remote PEs.
- This mechanism assumes that the same anti-spoof-mac is configured in all PEs for the service, and that traffic with destination **anti-spoof-mac** received on SAPs/SDP-bindings is dropped. An ingress MAC filter must be configured to drop traffic to the **anti-spoof-mac**.

The following table shows the combinations that produce a **Status = Active** proxy-ARP entry in the table. The system only replies to proxy-ARP requests for active entries. Any other combination result in a **Status = inActiv** entry. If the service is not active, the proxy-ARP entries are not active, regardless of the FDB entries



Note:

A static entry is active in the FDB even when the service is down.

Table 40: Proxy-ARP entry combinations

Proxy-ARP entry type	FDB entry type (for the same MAC)
Dynamic	learned
Static	CStatic
EVPN	EVPN, EVPNS with matching ESI
Duplicate	—

When proxy-ARP or proxy-ND is enabled on services with multi-homed ESs, a proxy-ARP entry type "EVPN" may be associated with a "learned" FDB entry because the CE can send traffic for the same MAC to all the multi-homed PEs in the ES. In such cases, the entry is inactive, in accordance with the preceding table.

4.3.1.1 Proxy-ARP/ND periodic refresh, unsolicited refresh, and confirm-messages

When proxy-ARP or proxy-ND is enabled, the system starts populating the proxy table and responding to ARP-requests or NS messages. To keep the active IP-to-MAC entries alive and ensure that all the host/routers in the service update their ARP/ND caches, the system may generate the following three types of ARP/ND messages for a specified IP-to-MAC entry:

- **periodic refresh messages (ARP-requests or NS for a specified IP)**

These messages are activated by the **send-refresh** command and their objective is to keep the existing FDB and proxy-ARP/ND entries alive, in order to minimize EVPN withdrawals and re-advertisements.

- **unsolicited refresh messages (unsolicited GARP or NA messages)**

These messages are sent by the system when a new entry is learned or updated. Their objective is to update the attached host/router caches.

- **confirm messages (unicast ARP-requests or unicast NS messages)**

These messages are sent by the system when a new MAC is learned for an existing IP. The objective of the confirm messages is to verify that a specified IP has moved to a different part of the network and is associated with the new MAC. If the IP has not moved, it forces the owners of the duplicate IP to reply and triggers **dup-detect**.

4.3.1.2 Proxy-ND and the Router flag in Neighbor Advertisement messages

RFC 4861 describes the use of the (R) or "Router" flag in NA messages as follows:

- a node capable of routing IPv6 packets must reply to NS messages with NA messages where the R flag is set (R=1)
- hosts must reply with NA messages where R=0

The use of the R flag in NA messages impacts how the hosts select their default gateways when sending packets "off-link". Therefore, it is important that the proxy-ND function on the 7210 SAS meet one of the following criteria:

1. provide the appropriate R flag information in proxy-ND NA replies
2. flood the received NA messages if it cannot provide the appropriate R flag when replying

Because of the use of the R flag, the procedure for learning proxy-ND entries and replying to NS messages differs from the procedures for proxy-ARP in IPv4: the router or host flag is added to each entry, and that determines the flag to use when responding to a NS.

4.3.1.3 Procedure to add the R Flag to a specified entry

The procedure to add the R flag to a specified entry is as follows:

- Dynamic entries are learned based on received NA messages. The R flag is also learned and added to the proxy-ND entry so that the appropriate R flag is used in response to NS requests for a specified IP.
- Static entries are configured as host or router as per the **[no] static ip-address ieee-address {host | router}** command.
- EVPN entries are learned from BGP and the **evpn-nd-advertise {host | router}** the R flag added to them.
- In addition, the **evpn-nd-advertise {host | router}** command indicates what static and dynamic IP-to-MAC entries the system advertises in EVPN. If **evpn-nd-advertise router** is configured, the system should flood the received unsolicited NA messages for hosts. This is controlled by the **[no] host-unsolicited-na-flood-evpn** command. The opposite is also recommended, so that the **evpn-nd-advertise host** is configured using the **router-unsolicited-na-flood-evpn** command.

4.3.1.4 Configuration guidelines for proxy-ARP and proxy-ND

On the 7210 SAS, users can enable or disable proxy-ARP and proxy-ND commands for all EVPN services configured on the node; however, the option to enable or disable proxy-ARP or proxy-ND per service is not available.

Use the following syntax to enable or disable proxy-ARP or proxy-ND capability per node.

```
configure>service>system>[no]evpn-proxy-arp-nd
```

If the per-node **evpn-proxy-arp-nd** command is disabled, it is not possible to enable the **proxy-arp** or **proxy-nd** command per service, and **proxy-arp** or **proxy-nd** is disabled for all EVPN services on the node.

When the **evpn-proxy-arp-nd** command is enabled, the user must run the following sequence of commands to disable the per-node **proxy-arp** and **proxy-nd** commands, if required:

1. Run the **config>service>system>no evpn-proxy-arp-nd** command.
2. Run the **config>service>vpls>no proxy-arp** command.
3. Run the **config>service>vpls>no proxy-nd** command.

The following example shows the command usage to enable **proxy-arp** and **proxy-nd** for all services, with all parameters set to default values:

```
-----
configure
  service
    system
      evpn-proxy-arp-nd
    exit
  exit
  service
    vpls 100
      proxy-arp
        no age-time
        dup-detect window 3 num-moves 5 hold-down 9
        no dynamic-arp-populate
        ... (so on for other parameters supported under proxy-arp)
        no shutdown
      exit
      proxy-nd
        no age-time
        dup-detect window 3 num-moves 5 hold-down 9
        no dynamic-nd-populate
        ... (so on for other parameters supported under proxy-nd)
        no shutdown
      exit
    exit
  exit
-----
```

When the user runs the **config>service>system>evpn-proxy-arp-nd** command, the software automatically runs a **no shutdown** command for **proxy-arp** and **proxy-nd** for all services. Similarly, when the user runs the **config>service>system>no evpn-proxy-arp-nd** command, the software automatically runs a **shutdown** command for **proxy-arp** and **proxy-nd** for all services.

When the **evpn-proxy-arp-nd** command is enabled, the **proxy-arp** or **proxy-nd** command is enabled for all services and cannot be disabled for individual services. Configuring service-specific **proxy-arp** or **proxy-nd** command parameters is supported only when the **evpn-proxy-arp-nd** command is enabled.

4.3.1.4.1 Proxy-ARP and proxy-ND support for spoke-SDP bindings

The 7210 SAS does not support proxy-ARP and proxy-ND on spoke-SDP bindings. ARP or ND packets received on spoke-SDP bindings are not identified or sent to the CPU for further processing.

When using spoke-SDP bindings on the 7210 SAS, Nokia recommends that users disable proxy-ARP and proxy-ND functionality on all PEs belonging to the EVPN service, or enable the forwarding of ARP or ND packets over EVPN bindings on all PEs that belong to the EVPN service, so that ARP or ND packets are flooded throughout the EVPN instance.

4.3.2 BGP-EVPN MAC mobility

EVPN defines a mechanism to allow the smooth mobility of MAC addresses from one CE/NVE to another. The 7210 SAS supports this procedure and the MAC mobility extended community in MAC advertisement routes:

- The router honors and generates the Sequence (SEQ) number in the MAC mobility extended community for MAC moves.
- When a MAC is EVPN-learned and it is attempted to be learned locally, a BGP update is sent with SEQ number changed to "previous SEQ"+1 (exception: **mac-duplication detect num-moves** value is reached).
- A SEQ number = zero or no MAC mobility *ext-community* are interpreted as sequence zero.
- In case of mobility, the following MAC selection procedure is followed:
 - If a PE has two or more active remote EVPN routes for the same MAC, the highest SEQ number is selected. The tie-breaker is the lowest IP (BGP NH IP).
 - If a PE has two or more active EVPN routes and it is the originator of one of them, the highest SEQ number is selected. The tie-breaker is the lowest IP (BGP NH IP of the remote route is compared to the local system address).



Note:

When EVPN multi-homing is used in EVPN-MPLS, the ESI is compared to determine whether a MAC received from two different PEs should be processed within the context of MAC mobility or multi-homing. Two MAC routes that are associated with the same remote or local ESI but different PEs are considered reachable through all those PEs. Mobility procedures are not triggered if the MAC route still belongs to the same ESI.

4.3.3 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 7210 SAS supports an enhanced version of this procedure, which is described in this section.

In a scenario where two or more hosts are misconfigured using the same (duplicate) MAC address, the duplicate MAC address is learned by the PEs in the VPLS. As a result, the traffic originating from the hosts triggers continuous MAC moves among the PEs attached to the hosts. It is important to recognize such situations and avoid incrementing the sequence number (in the MAC Mobility attribute) to infinity.

To remedy accidentally duplicated MAC addresses, a router that detects a MAC mobility event through local learning starts a **window in-minutes** timer (the default value is 3). If the configured **num-moves num**

value is detected before the timer expires (the default value is 5), the router concludes that a duplicate MAC situation has occurred and sends a trap message to alert the operator. Use the **show service id svc-id bgp-evpn** command to display the MAC addresses. The following is a sample configuration output.

```
10 2014/01/14 01:00:22.91 UTC MINOR: SVCMMGR #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
MPLS Admin Status     : Enabled          Creation Origin      : manual
MAC Dup Detn Moves    : 5                MAC Dup Detn Window  : 3
MAC Dup Detn Retry    : 9                Number of Dup MACs   : 1
-----
Detected Duplicate MAC Addresses      Time Detected
-----
00:00:00:00:00:12                   01/14/2014 01:00:23
-----
=====
```

After a duplicate MAC address is detected, the router stops sending and processing BGP MAC advertisement routes for that MAC address until one of the following occurs:

1. The MAC is flushed because of a local event (SAP or SDP-binding associated with the MAC fails) or the reception of a remote update with better SEQ number (because of a MAC flush at the remote router).
2. The **retry in-minutes** timer expires, which flushes the MAC and restarts the process.



Note:

The other routers in the VPLS instance forward the traffic for the duplicate MAC address to the router advertising the best route for the MAC.

The values of **num-moves** and **window** can be configured for different environments. In scenarios where BGP rapid-update EVPN is configured, the operator should configure a shorter window timer than scenarios where BGP updates are sent per the configured **min-route-advertisement** interval, which is the default.

The preceding MAC duplication parameters can be configured per VPLS service under the **bgp-evpn mac-duplication** context. The following is a sample configuration output.

Example

```
A:Dut-B>config>service>vpls>bgp-evpn# info
-----
    evi 1
    mac-duplication
      detect num-moves 5 window 2
      retry 10
    exit
    mpls
      split-horizon-group "vpls1"
      ingress-replication-bum-label
      auto-bind-tunnel
      resolution-filter
      ldp
      exit
      resolution filter
    exit
    no shutdown
```



```
exit
-----
```

4.3.4 Conditional static MAC and protection

In RFC 7432, the MAC Mobility Extended Community section defines the use of the sticky bit to signal static MAC addresses. These addresses must be protected to prevent attempts to dynamically learn them in a different place in the EVPN-MPLS VPLS service.



Note:

On the 7210 SAS, the conditional static MACs are not protected using MAC-protect functionality. A Cstatic MAC is advertised to other PEs with the sticky bit set so that it is prevented from being learned dynamically at a different place in the EVPN-MPLS VPLS service. MAC frames whose source MAC address matches the statically configured MAC address are forwarded based on destination MAC address lookup and are not dropped.

In the 7210 SAS, any conditional static MAC address that is defined in an EVPN-MPLS VPLS service is advertised by BGP-EVPN as a static address (that is, with the sticky bit set). The following is a sample output that shows the configuration of a conditional static MAC.

Example

```
*A:Dut-B>config>service>vpls# info
-----
description "evpn mpls service "
.....
sap lag-1:1 create
description "Default sap description for service id 1"
no shutdown
exit
static-mac
mac 00:ca:ca:ca:ca:00 create sap lag-1:1 monitor fwd-status
exit
```

```
A:Dut-C# show router bgp routes evpn mac hunt mac-address 00:ca:ca:ca:ca:00
```

```
.....
=====
BGP EVPN MAC Routes
=====
-----
RIB In Entries
-----
Network       : n/a
Nexthop       : 10.20.1.2
From          : 10.20.1.2
Res. Nexthop  : 10.10.3.2
Local Pref.   : 100
Aggregator AS : None
Atomic Aggr.  : Not Atomic
AIGP Metric   : None
Connector     : None
Community     : target:100:1 bgp-tunnel-encap:MPLS
               mac-mobility:Seq:0/Static
Cluster       : No Cluster Members
Originator Id : None
Flags         : Used Valid Best IGP
Route Source  : Internal
Interface Name : ip-10.10.3.3
Aggregator    : None
MED           : 0
IGP Cost      : 400
Peer Router Id : 10.20.1.2
```

```

AS-Path      : No As-Path
EVPN type    : MAC
ESI          : 00:bc:01:00:00:00:00:00:01
Tag          : 0
IP Address   : n/a
Route Dist.  : 2.2.2.2:1
Mac Address  : 00:ca:ca:ca:ca:00
MPLS Label1  : LABEL 131056      MPLS Label2   : n/a
Route Tag    : 0
Neighbor-AS  : n/a
Orig Validation: N/A
Add Paths Send : Default
Last Modified : 00h02m02s

```

```

-----
RIB Out Entries
-----

```

```

Routes : 1
=====

```

4.3.5 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 has to make a selection (for example, the same MAC 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 as follows:

- **EVPN route selection for MAC routes**

When two or more routes with the same *mac-length/mac* but different route key are received, BGP transfers the routes to EVPN. EVPN selects the route based on the following tie-breaking order:

1. conditional static MACs (local protected MACs)
2. EVPN static MACs (remote protected MACs)
3. data plane learned MACs (regular learning on SAPs/SDP-bindings)
4. EVPN MACs with higher SEQ number
5. lowest IP (next-hop IP of the EVPN NLRI)
6. lowest Ethernet tag (that is zero for MPLS)
7. lowest RD

- **BGP route selection for MAC routes with the same route-key**

The priority order is as follows:

1. EVPN static MACs (remote protected MACs)
2. EVPN MACs with higher sequence number
3. regular BGP selection (local-pref, aigp metric, shortest as-path, ..., lowest IP)

- **BGP route selection for the rest of the EVPN routes follows regular BGP selection**

**Note:**

If BGP runs through the selection criteria and a specified and valid EVPN route is not selected in favor of another EVPN route, the non-selected route is displayed by the **show router bgp routes evpn evpn-type detail** command with a tie-breaker reason.

4.3.6 EVPN interaction with other features

This section describes the interaction of EVPN with other features.

4.3.6.1 EVPN-MPLS with existing VPLS features

When enabling existing VPLS features in an EVPN-MPLS-enabled service, the following considerations apply:

- EVPN-MPLS is only supported in regular VPLS. Other VPLS types, such as **m-vpls**, are not supported.
- In general, no router-generated control packets are sent to the EVPN destination bindings, except for proxy-ARP/proxy-ND confirm messages for EVPN-MPLS.
- For xSTP and M-VPLS services, the following applies:
 - xSTP can be configured in BGP-EVPN services. BPDUs are not sent over the EVPN bindings.
 - BGP-EVPN is blocked in M-VPLS services; however, a different M-VPLS service can manage a SAP or spoke-SDP in a BGP-EVPN-enabled service.
- For BGP-EVPN-enabled VPLS services, **mac-move** can be used in SAPs/SDP-bindings; however, MACs learned through BGP-EVPN are not considered.

**Note:**

MAC duplication provides protection against MAC moves between EVPN and SAPs/SDP-bindings.

- The **disable-learning** command and other FDB-related tools only work for data-plane-learned MAC addresses.
- MAC OAM tools (**mac-ping**, **mac-trace**, **mac-populate**, **mac-purge**, and **cpe-ping**) are not supported for BGP-EVPN services.
- SAPs that belong to a specified ES but are configured on non-BGP-EVPN-MPLS-enabled VPLS or Epipe services are kept down using the **StandByForMHPProtocol** flag.
- CPE ping is not supported on EVPN services.
- Other features not supported in conjunction with BGP-EVPN are:
 - endpoints and attributes
 - BPDU translation
 - L2PT termination
 - MAC-pinning
 - IGMP snooping in VPLS services when BGP-EVPN MPLS is enabled (in the service)
 - DHCP snooping
 - ETH-CFM (MEPs, vMEPs, MIPs)


- **allow-ip-int-bind** (R-VPLS)

4.3.6.2 EVPN with G.8032 in an access ring

It is possible to use the G.8032 operation in an access ring with EVPN. The only supported configuration is a G.8032 sub-ring with a non-virtual link and without MAC flush propagation from the EVPN network to the G.8032 sub-ring. This section provides a sample configuration and guidelines about the configuration.

The following figure shows the network topology of an access ring with EVPN. It shows a G.8032 sub-ring formed by nodes G, B, C on the left side of the figure and nodes A, D, E on the right side of the figure, connected to the EVPN network formed by nodes B, C, D, E.

Figure 57: Network topology of an access ring



Nodes B, C, D, E connect to both the EVPN network using network ports and to the G.8032 ring using access ports. For example, on node B, network ports 1/1/7, 1/1/1, and 1/1/6 connect PE-B to remote EVPN nodes D, E, C, respectively. Additionally, on node B, access port 1/1/22 is part of the G.8032 access ring that connects PE-B to the G.8032 ring formed with access CE node G.

EVPN bindings are protected by using fast reroute (FRR) paths; however, in the event a failure occurs in the EVPN network, MAC flush is not propagated from the EVPN network to the G.8032 ring.

G.8032 data SAPs and control SAPs on the EVPN PE nodes (B, C, D, E) can be configured only on non-ES ports. Non-ES LAGs cannot be used with G.8032 on 7210 SAS.

Example

The following is a sample configuration of the access CE node, node A in the preceding figure, which is part of the G.8032 access ring.

```
#-----
echo "System Configuration"
#-----
system
  name "Dut-A"
#-----
```

```

echo "Ethernet Rings Configuration"
#-----
eth-ring 124
exit
eth-ring 124
description "Ethernet Ring 124"
guard-time 20
revert-time 60
rpl-node owner
path a 1/1/1 raps-tag 124
description "Ethernet Ring : 124 Path : pathA"
rpl-end
eth-cfm
mep 6 domain 1 association 1241
ccm-enable
control-mep
control-sap-tag 724
no shutdown
exit
exit
no shutdown
exit
path b 1/1/2 raps-tag 124
description "Ethernet Ring : 124 Path : pathB"
eth-cfm
mep 7 domain 1 association 1242
ccm-enable
control-mep
control-sap-tag 724
no shutdown
exit
exit
no shutdown
exit
no shutdown
exit
#-----
-----snipped-----
#-----
echo "Service Configuration"
#-----
service
customer 1 create
description "Default customer"
exit
vpls 1 customer 1 svc-sap-type any create
description "Default tls description for service id 1"
disable-learning

stp
shutdown
exit
sap 1/1/5:1 create
description "Default sap description for service id 1"
egress
exit
exit
sap 1/1/1:1 eth-ring 124 create
stp
shutdown
exit
egress
exit
exit

```

```

        sap 1/1/2:1 eth-ring 124 create
        stp
            shutdown
        exit
        egress
        exit
    exit
    no shutdown
exit
vpls 124 customer 1 vpn 124 svc-sap-type any create
description "Default tls description for service id 124"
stp
    shutdown
exit
sap 1/1/1:124 eth-ring 124 create
description "SAP 1/1/1:124 on Ethernet Ring 124 "
stp
    shutdown
    exit
    egress
    exit
exit
sap 1/1/2:124 eth-ring 124 create
description "SAP 1/1/2:124 on Ethernet Ring 124 "
stp
    shutdown
    exit
    egress
    exit
    exit
    no shutdown
    exit
exit
#-----

```

Example

The following is a sample configuration of an EVPN PE node, node D in the preceding figure.

```

#-----
echo "System Configuration"
#-----
    system
        name "Dut-D"
    -----snipped-----
#-----
echo "Ethernet Rings Configuration"
#-----
    eth-ring 124
    exit
    eth-ring 124
        description "Ethernet Ring 124"
        guard-time 20
        path a 1/1/21 raps-tag 124
            description "Ethernet Ring : 124 Path : pathA"
            eth-cfm
                mep 5 domain 1 association 1243
                ccm-enable
                control-mep
                control-sap-tag 724
                no shutdown
            exit
        exit
    exit

```

```

        no shutdown
    exit
    no shutdown
exit
#-----
-----snipped-----
#-----
echo "Service Configuration"
#-----
service
    sdp 42 mpls create
        far-end 10.20.1.2
        ldp
        path-mtu 1600
        keep-alive
        shutdown
    exit
    no shutdown
exit
sdp 43 mpls create
    far-end 10.20.1.3
    ldp
    path-mtu 1600
    keep-alive
    shutdown
    exit
    no shutdown
exit
sdp 45 mpls create
    far-end 10.20.1.5
    ldp
    path-mtu 1600
    keep-alive
    shutdown
    exit
    no shutdown
exit
customer 1 create
    description "Default customer"
exit
system
    bgp-evpn
        ethernet-segment "esPort1" create
            esi 00:de:03:00:00:00:00:00:03
            service-carving
            mode auto
            exit
            multi-homing single-active no-esi-label
            shutdown
        exit
    exit
exit
vpls 1 customer 1 svc-sap-type any create
    description "Default tls description for service id 1"
    split-horizon-group "vpls1" create
        description "Default description for SHG vpls1"
    exit
    bgp-evpn
        evi 1
        mpls
            control-word
            force-vlan-vc-forwarding
            split-horizon-group "vpls1"
            ingress-replication-bum-label

```

```

        auto-bind-tunnel
        resolution any
        exit
        no shutdown
    exit
exit
stp
    shutdown
exit
sap 1/1/21:1 eth-ring 124 create
    stp
        shutdown
    exit
    egress
    exit
exit
no shutdown
exit
vpls 124 customer 1 vpn 124 svc-sap-type any create
    description "Default tls description for service id 124"
    stp
        shutdown
    exit
    sap 1/1/21:124 eth-ring 124 create
        description "SAP 1/1/21:124 on Ethernet Ring 124 "
        stp
            shutdown
        exit
        egress
        exit
    exit
    no shutdown
exit
exit
#-----
echo "Router (Service Side) Configuration"
#-----
    router Base
#-----
echo "BGP Configuration"
#-----
    bgp
        connect-retry 1
        min-route-advertisement 1
        rapid-withdrawal
        bfd-enable
        group "bgpEvpn"
            peer-as 100
            bfd-enable
            neighbor 10.20.1.2
                family evpn
                peer-as 100
                bfd-enable
            exit
            neighbor 10.20.1.3
                family evpn
                peer-as 100
                bfd-enable
            exit
            neighbor 10.20.1.5
                family evpn
                peer-as 100
                bfd-enable
            exit
        exit

```



```

        exit
        no shutdown
    exit
#-----

```

4.3.7 Routing policies for BGP EVPN routes

Routing policies match on specific fields when importing or exporting EVPN routes. These matching fields are the following:

- communities (*comm-val*), extended communities (*ext-comm*), and large communities (*large-comm*)
- well-known communities (*well-known-comm*); **no-export** | **no-export-subconfed** | **no-advertise**
- family EVPN
- protocol BGP-VPN (this term also matches VPN-IPv4/6 routes)
- BGP attributes that are applicable to EVPN routes (such as AS-path, local-preference, next-hop)

4.4 Configuring an EVPN service with CLI

This section provides information to configure EVPN services using the CLI for 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone mode).

4.4.1 EVPN-MPLS configuration examples

This section provides EVPN-MPLS configuration examples.

4.4.1.1 EVPN single-active multi-homing example

To use single-active multi-homing on PE-1 and PE-2 instead of all-active multi-homing perform the following:

- Change the LAG configuration to **multi-homing single-active**.
The CE-12 is now configured with two different LAGs; therefore, the key, system ID, and system priority values must be different on PE-1 and PE-2.
- Change the Ethernet segment configuration to **multi-homing single-active**

No changes are needed at the service level on any of the three PEs.

Example: Single-active and all-active multi-homing

The following configuration example shows the differences between single-active multi-homing and all-active multi-homing.

```

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
esi 00:de:01:00:00:00:00:00:01
    service-carving
        mode auto
    exit
    multi-homing single-active
    lag 6
    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 00:00:00:00:72:72
        no shutdown
-----
A:PE2>config>lag# /configure service system bgp-evpn
A:PE2>config>service>system>bgp-evpn# info
esi 00:de:01:00:00:00:00:00:01
    service-carving
        mode auto
    exit
    multi-homing single-active
    lag 6
    no shutdown

```

4.5 EVPN command reference

This section describes the EVPN commands for 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone mode).

4.5.1 Command hierarchies

- [EVPN configuration commands](#)
- [EVPN show commands](#)
- [EVPN clear commands](#)
- [EVPN tools commands](#)

4.5.1.1 EVPN configuration commands



Note:

See [VPLS service configuration commands](#) for information about configuring commands in the `config>vpls>bgp` context.

```

config
- service

```

```

- system
  - [no] evpn-proxy-arp-nd
config
- service
  - epipe service-id [customer customer-id] [create] [vpn vpn-id][vc-switching]
  - epipe service-id [customer customer-id] [create] [vpn vpn-id][svc-sap-type {null-
star | dot1q | dot1q-preserve | any | dot1q-range | qinq-inner-tag-preserve}] [customer-
vid vlan-id]
  - no epipe service-id
  - bgp
  - no bgp
    - route-distinguisher rd
    - no route-distinguisher
    - route-target {ext-community | export ext-community | import ext-community}
    - no route-target
  - bgp-evpn
  - no bgp-evpn
    - evi value
    - no evi
    - local-ac-name ac-name
    - no local-ac-name
      - [no] eth-tag tag-value
    - remote-ac-name ac-name
    - no remote-ac-name
      - [no] eth-tag tag-value
    - mpls
      - auto-bind-tunnel
        - resolution {disabled | any | filter}
        - resolution-filter
          - [no] bgp
          - [no] ldp
          - [no] rsvp
          - [no] sr-isis
          - [no] sr-ospf
      - [no] control-word
      - [no] force-vlan-vc-forwarding
      - oper-group name
      - no oper-group
      - [no] shutdown
  - vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls]
  - no vpls service-id
    - bgp-evpn
    - no bgp-evpn
      - evi value
      - no evi
      - [no] mac-advertisement
      - mac-duplication
        - detect num-moves num-moves window minutes
        - retry minutes
        - no retry
    - mpls
      - auto-bind-tunnel
        - resolution {disabled | any | filter}
        - resolution-filter
          - [no] bgp
          - [no] ldp
          - [no] rsvp
          - [no] sr-isis
          - [no] sr-ospf
      - [no] control-word
      - [no] force-vlan-vc-forwarding
      - [no] ingress-replication-bum-label
      - [no] shutdown
      - split-horizon-group name

```

```

- no split-horizon-group
- [no] proxy-arp
  - age-time seconds
  - no age-time
  - dup-detect [anti-spoof-mac mac-address] window minutes num-moves count hold-
down minutes | max
  - [no] dynamic-arp-populate
  - [no] garp-flood-evpn
  - [no] send-refresh seconds
  - static ip-address ieee-address
  - no static ip-address
  - table-size table-size
  - [no] unknown-arp-request-flood-evpn
  - [no] shutdown
- [no] proxy-nd
  - age-time seconds
  - no age-time
  - dup-detect [anti-spoof-mac mac-address] window minutes num-moves count hold-
down minutes | max
  - [no] dynamic-nd-populate
  - evpn-nd-advertise {host | router}
  - [no] host-unsolicited-na-flood-evpn
  - [no] router-unsolicited-na-flood-evpn
  - [no] send-refresh seconds
  - [no] static ip-address ieee-address {host | router}
  - table-size table-size
  - [no] unknown-ns-flood-evpn
  - [no] shutdown

```

```

config
- service
- system
- bgp-evpn
  - ethernet-segment name [create]
  - no ethernet-segment name
  - es-activation-timer seconds
  - no es-activation-timer
  - esi esi
  - no esi
  - lag lag-id
  - no lag
  - multi-homing single-active no-esi-label
  - no multi-homing
  - oper-group name
  - no oper-group
  - port port-id
  - no port
  - service-carving
    - manual
      - evi start [to to]
      - no evi start
      - mode {auto | manual | off}
    - [no] shutdown
  - route-distinguisher rd
  - no route-distinguisher

config
- redundancy
  - bgp-evpn-multi-homing
  - boot-timer seconds
  - es-activation-timer seconds

```

4.5.1.2 EVPN show commands

```
show
- service
  - evpn-mpls
  - id service-id
    - bgp-evpn
    - evpn-mpls [esi esi]
    - proxy-arp [ip-address] [detail]
    - proxy-nd [ip-address] [detail]
  - system
    - bgp-evpn
    - bgp-evpn ethernet-segment
    - bgp-evpn ethernet-segment name name [all]
    - bgp-evpn ethernet-segment name name evi [evi]
```

```
show
- redundancy
  - bgp-evpn-multi-homing
```

4.5.1.3 EVPN clear commands

```
clear
- service
  - id service-id
    - proxy-arp [duplicate] [dynamic]
    - proxy-nd [duplicate] [dynamic]
```

4.5.1.4 EVPN tools commands

```
tools
- dump
  - service
    - proxy-arp
    - proxy-nd
```

4.5.2 Command descriptions

4.5.2.1 EVPN configuration commands

evpn-proxy-arp-nd

Syntax

[no] evpn-proxy-arp-nd

Context

config>service>system

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables proxy-ARP and proxy-ND capability per node.

When this command is enabled, the **proxy-arp** and **proxy-nd** commands are enabled for all services and cannot be disabled for individual services. Using the per-service CLI context and commands under the **proxy-arp** or **proxy-nd** command, users can configure the service-specific **proxy-arp** or **proxy-nd** command parameters. Only when the **evpn-proxy-arp-nd** command is enabled can the per-service CLI commands be used to configure **proxy-arp** or **proxy-nd** parameters.

When this command is disabled, it is not possible to enable the **proxy-arp** or **proxy-nd** command per service, and **proxy-arp** and **proxy-nd** is disabled for all EVPN services on the node.

The **no** form of this command disables proxy-ARP and proxy-ND capability per node.



Note:

The **no** form of this command reverts all configured **proxy-arp** and **proxy-nd** command parameters to the default values and shuts down proxy-ARP and proxy-ND for all services.

Default

no evpn-proxy-arp-nd

epipe

Syntax

epipe *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**vc-switching**]

epipe *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**svc-sap-type** {**null-star** | **dot1q** | **dot1q-preserve** | **any** | **dot1q-range** | **qinq-inner-tag-preserve**}] [**customer-vid** *vlan-id*]

no epipe *service-id*

Context

config>service

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures a point-to-point Epipe service instance. An Epipe connects two endpoints defined as SAPs. Both SAPs may be defined in one or separate devices connected over the service provider network. When the endpoint SAPs are separated by the service provider network, the far-end SAP is generalized into an SDP. This SDP describes a destination and the encapsulation method used to reach it. In addition to the SDPs, endpoint SAPs can also be connected by EVPN destinations. A VPLS connects

multiple customer sites together acting like a zero-hop, Layer 2 switched domain. A VPLS is always a logical full mesh.

No MAC learning or filtering is provided 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 (created using the **customer** command in the service context). The customer association cannot be edited; the service must be deleted and recreated with a new customer association. More than one VPLS may be created for a single customer ID. By default, no VPLS instances exist until they are explicitly created.

After 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 results in an error.

By default, no Epipe services exist until they are explicitly created with this command.

The **no** form of this command deletes the Epipe service instance with the specified *service-id*. The service cannot be deleted until the service has been shutdown.

Parameters

service-id

Specifies the unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id* — 1 to 2147483647
 svc-name — a string up to 64 characters

customer *customer-id*

Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn *vpn-id*

Specifies the VPN ID number which allows you to identify VPNs by a VPN identification number.

Values 1 to 2147483647

Default null (0)

vc-switching

Specifies whether pseudowire switching signaling is used for the spoke-SDPs configured in the service. This is not supported for EVPN-VPWS.

svc-sap-type

Keyword that specifies the type of access SAPs and access-uplink SAPs allowed in the service.

Values **null-star** — Specifies that the allowed SAP in the service can be null SAPs, dot1q default, Q.* SAP, 0.* SAP, or default QinQ SAP (also known as *.* SAP).

dot1q — Specifies that the allowed SAPs in the service are dot1q SAPs and dot1q explicit null SAPs.

dot1q-preserve — Specifies that the allowed SAPs in the service are dot1q. The dot1q ID is not stripped after packets match the SAP.

dot1q-range — Specifies that the access SAP in the service can use VLAN ranges as the SAP tags. The VLAN ranges are configured using the **config>connection-profile** command. On ingress of the access dot1q SAP using VLAN ranges, the outermost tag is not removed before forwarding.

any — Specifies that any SAP type is allowed.

qinq-inner-tag-preserve — Specifies that an Epipe service processes and forwards packets received with three or more tags on a QinQ SAP.

Default any

customer-vid *vlan-id*

Specifies the dot1q VLAN ID used while creating the local dot1q SAP for **svc-sap-type dot1q-preserve**. Applicable only for access-uplink mode.

Values 1 to 4094

bgp

Syntax

bgp

no bgp

Context

config>service>epipe

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures BGP related parameters for BGP EVPN.

The **no** form of this command disables EVPN-VPWS.

Default

no bgp

route-distinguisher

Syntax

```
route-distinguisher rd
no route-distinguisher
```

Context

```
config>service>vpls>bgp
config>service>epipe>bgp
```

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the Route Distinguisher (RD) component signaled in the MP-BGP NLRI for L2VPN and EVPN families. This value is used for BGP-AD, BGP VPLS, and BGP multi-homing NLRI if these features are configured.

If this command is not configured, the RD is automatically built using the BGP-AD VPLS ID. The following rules apply:

- if BGP AD VPLS-ID is configured and no RD is configured under the BGP node - RD=VPLS-ID
- if BGP AD VPLS-id is not configured, an RD value must be configured under the BGP node (this is the case when only BGP VPLS is configured)
- if BGP AD VPLS-id is configured and an RD value is also configured under BGP node, the configured RD value prevails

Alternatively, the **auto-rd** option allows the system to automatically generate an RD based on the **bgp-auto-rd-range** command configured at the service level. For **bgp-evpn** enabled VPLS and Epipe services, the **route-distinguisher** value 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. See the **evi** command description for more information.

Default

```
no route-distinguisher
```

Parameters

rd

Specifies all routes in the specified BGP community.

Values	<i>ip-addr:comm-val 2byte-asnumber:ext-comm-val> 4byte-asnumber:comm-val</i>
ip-addr	a.b.c.d
comm-val	0 to 65535

2byte-asnumber 1 to 65535
 ext-comm-val 0 to 4294967295
 4byte-asnumber 0 to 4294967295

route-target

Syntax

route-target {*ext-community* | **export** *ext-community* | **import** *ext-community*}
no route-target

Context

config>service>vpls>bgp
 config>service>epipe>bgp

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the Route Target (RT) component that is signaled in the related MP-BGP attribute used for BGP auto-discovery, BGP VPLS, BGP multi-homing, and EVPN if these features are configured in this VPLS service.

If this command is not used, the RT is built automatically using the VPLS ID. The extended community can have the same two formats as the VPLS ID, a two-octet AS-specific extended community, IPv4 specific extended community. For BGP EVPN enabled VPLS and Epipe 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. See the [evi](#) command description for more information.

Default

no route-target

Parameters

export *ext-community*

Specifies communities allowed to be sent to remote PE neighbors, up to 72 characters.

import *ext-community*

Specifies communities allowed to be accepted from remote PE neighbors, up to 72 characters.

vpls

Syntax

vpls *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**m-vpls**] [**name** *name*] [**create**]
no vpls *service-id*

Context

config>service

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

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 connects multiple customer sites together acting as a zero-hop, Layer 2 switched domain. A VPLS is always a logical full mesh.

If the **create** command is enabled in the **environment** context, the **create** keyword must be specified when the service is created. Specify the **customer** keyword and *customer-id* to associate the service with a customer. The *customer-id* must already exist (created using the **customer** command in the service context). After 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.

After 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 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 shutdown and deleted, and the service has been shutdown.

Parameters

service-id

Specifies the unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id* — 1 to 2147483648
 svc-name — a string up to 64 characters

customer *customer-id*

Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn *vpn-id*

Specifies the VPN ID number which allows you to identify VPNs by a VPN identification number.

Values 1 to 2147483647

Default null (0)

m-vpls

Specifies a management VPLS.

bgp-evpn**Syntax**

bgp-evpn

no bgp-evpn

Context

config>service>vpls

config>service>epipe

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables BGP-EVPN in the base instance.

The **no** form of this command disables BGP-EVPN.

evi**Syntax**

evi *value*

no evi

Context

config>service>vpls>bgp-evpn

config>service>epipe>bgp-evpn

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures a 2-byte EVPN instance that is unique in the system. It is used by the service-carving algorithm for multihoming and auto-deriving RTs and RDs.

If not specified, the value is zero and no RD or RTs are autoderived. If the **evi value** is specified and no other RD or RT is configured in the service, the following rules apply:

- the RD is derived from **<system_ip>:evi**
- the RT is derived from **<autonomous-system>:evi**



Note:

If VSI import and export policies are configured, the RT must be configured in the policies, and those values take precedence over the auto-derived RTs. The operational RT for a service is shown in the **show service id bgp** command.

The **no** form of this command reverts the **evi value** to zero.

Default

no evi

Parameters

value

Specifies the EVPN instance.

Values 1 to 65535

local-ac-name

Syntax

local-ac-name *ac-name*

no local-ac-name

Context

config>service>epipe>bgp-evpn

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure the local Ethernet tag value.

The **no** form of this command disables the local Ethernet tag value.

Default

no local-ac-name

Parameters

ac-name

Specifies the name of the local attachment circuit, up to 32 characters.

eth-tag**Syntax**

[no] **eth-tag** *tag-value*

Context

config>service>epipe>bgp-evpn>local-ac-name

config>service>epipe>bgp-evpn>remote-ac-name

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

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 Ethernet tag value of the imported AD per-EVI routes for the service. If there is a match, the system creates an EVPN destination for the Epipe.

The **no** form of this command disables the local Ethernet tag value.

Default

no eth-tag

Parameters

tag-value

Specifies the Ethernet tag value of the attachment circuit.

Values 1 to 16777215

remote-ac-name**Syntax**

remote-ac-name *ac-name*

no remote-ac-name

Context

```
config>service>epipe>bgp-evpn
```

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure the remote Ethernet tag value.

The **no** form of this command disables the remote Ethernet tag value.

Default

no remote-ac-name

Parameters

ac-name

Specifies the name of the remote attachment circuit, up to 32 characters.

mpls

Syntax

mpls

Context

```
config>service>vpls>bgp-evpn
```

```
config>service>epipe>bgp-evpn
```

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure the BGP EVPN MPLS parameters.

auto-bind-tunnel

Syntax

auto-bind-tunnel

Context

```
config>service>vpls>bgp-evpn>mpls
```

```
config>service>epipe>bgp-evpn>mpls
```

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure automatic binding of a BGP-EVPN service using tunnels to MP-BGP peers.

The **resolution** mode must be configured to enable auto-bind resolution to tunnels in TTM. The following configurations are available:

- If **resolution** is explicitly set to **disabled**, the auto-binding to the tunnel is removed.
- If **resolution** is set to **any**, any supported tunnel type in the EVPN context is selected, following TTM preference.
- The **resolution-filter** option is used to specify one or more explicit tunnel types; only the specified tunnel types are selected again following the TTM preference.

The following tunnel types are supported in a BGP-EVPN MPLS context, in order of preference: RSVP, LDP, SR-ISIS, SR-OSPF, and BGP.

The **rsvp** value specifies that BGP searches for the best metric RSVP LSP to the address of the BGP next hop. This address can correspond to the system interface or to another loopback used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel ID.

The **ldp** value specifies that BGP searches for an LDP LSP with a FEC prefix corresponding to the address of the BGP next hop.

The **sr-isis** (**sr-ospf**) value specifies that an SR tunnel to the BGP next hop is selected in the TTM from the lowest numbered ISIS (OSPF) instance.

The **bgp** value specifies 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, the inter-area or inter-as prefixes is not resolved.

To activate the list of tunnel types configured under **resolution-filter**, the **resolution** must be set to **filter**.

resolution

Syntax

resolution {**disabled** | **any** | **filter**}

Context

```
config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel
```

```
config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel
```

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the resolution mode in the automatic binding of a BGP-EVPN MPLS service to tunnels to MP-BGP peers.

Default

resolution disabled

Parameters**disabled**

Disables the automatic binding of a BGP-EVPN MPLS service to tunnels to MP-BGP peers.

any

Enables the binding to any supported tunnel type in a BGP-EVPN MPLS context following TTM preference.

filter

Enables the binding to the subset of tunnel types configured under **resolution-filter**.

resolution-filter**Syntax**

resolution-filter

Context

config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel

config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure 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 following tunnel types are supported in a BGP-EVPN MPLS context, in order of preference: RSVP, LDP, Segment Routing (SR), BGP, and UDP.

bgp**Syntax**

[no] **bgp**

Context

config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the BGP tunnel type.

BGP EVPN searches for a BGP LSP to the address of the BGP next hop. If the user does not enable the BGP tunnel type, the inter-area or inter-as prefixes are not resolved.

The **no** form of this command disables BGP as a tunnel type to consider.

Default

no bgp

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

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the LDP tunnel type.

BGP searches for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.

The **no** form of this command disables LDP as a tunnel type to consider.

Default

no ldp

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

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the RSVP-TE tunnel type.

BGP searches for the best metric RSVP LSP to the address of the BGP next hop. This address can correspond to the system interface or to another loopback used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel ID.

The **no** form of this command disables RSVP as a tunnel type to consider.

Default

no rsvp

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

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the Segment Routing (SR) tunnel type programmed by an ISIS instance in TTM.

The **no** form of this command disables SR-ISIS as a tunnel type to consider.

Default

no sr-isis

sr-ospf

Syntax

[no] sr-ospf

Context

config>service>vpls>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

config>service>epipe>bgp-evpn>mpls>auto-bind-tunnel>resolution-filter

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the SR tunnel type programmed by an OSPF instance in TTM.

The SR tunnel to the BGP next hop is selected in the TTM from the lowest numbered IS-IS (OSPF) instance.

The **no** form of this command disables SR-OSPF as a tunnel type to consider.

Default

no sr-ospf

shutdown

Syntax

shutdown

no shutdown

Context

config>service>vpls>bgp-evpn>mpls

config>service>epipe>bgp-evpn>mpls

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

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

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

Default

shutdown

mac-advertisement

Syntax

[no] mac-advertisement

Context

config>service>vpls>bgp-evpn

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables the advertisement in BGP of the learned MACs on SAPs and SDP bindings. When the **mac-advertisement** command is disabled, the local MACs will be withdrawn in BGP.

The **no** form of this command disables **mac-advertisement**.

Default

mac-advertisement

mac-duplication

Syntax

mac-duplication

Context

config>service>vpls>bgp-evpn

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure the BGP EVPN MAC duplication parameters.

detect

Syntax

detect num-moves *num-moves* **window** *minutes*

Context

config>service>vpls>bgp-evpn>mac-duplication

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command modifies the default behavior of the **mac-duplication** feature, which is always enabled by default. The command specifies the number of moves (**num-moves**) to monitor within a period of time (**window**).

Default

detect num-moves 5 window 3

Parameters

num-moves

Specifies the number of MAC moves in a VPLS. The counter is incremented when a specified MAC is locally relearned in the FDB or flushed from the FDB because of the reception of a better remote EVPN route for that MAC.

Values 3 to 10

Default 5

minutes

Specifies the length of the window, in minutes.

Values 1 to 15

Default 3

retry

Syntax

retry *minutes*

no retry

Context

config>service>vpls>bgp-evpn>mac-duplication

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the timer after which the MAC in hold-down state is automatically flushed and the **mac-duplication** process starts again. This value is expected to be equal to two times or more than that of **window**.

If the **no** form of this command is configured and **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

retry 9

Parameters

minutes

Specifies the BGP EVPN MAC duplication retry, in minutes.

Values 2 to 60

control-word

Syntax

[no] control-word

Context

config>service>vpls>bgp-evpn>mpls

config>service>epipe>bgp-evpn>mpls

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables the transmission and reception of the control word, as defined in RFC 7432, which helps avoid frame disordering.

This command is enabled or disabled for all EVPN-MPLS destinations at the same time.

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

Default

no control-word

force-vlan-vc-forwarding

Syntax

[no] force-vlan-vc-forwarding

Context

config>service>vpls>bgp-evpn>mpls

config>service>epipe>bgp-evpn>mpls

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

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.



Note:

When the **force-vlan-vc-forwarding** command is enabled, the VC VLAN ID is always set to 0.

This command is disabled on the 7210 SAS. It is set to the **no** form by default and cannot be enabled. If the ingress SAP/SDP binding is null-encapsulated, the output VLAN ID and pBits are zero.

Default

no force-vlan-vc-forwarding

oper-group

Syntax

oper-group *name*

no oper-group

Context

config>service>epipe>bgp-evpn>mpls

config>service>system>bgp-evpn>ethernet-segment

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command adds the **bgp-evpn mpls** instance as a member of the operational group. The operational group is up when either this command is not yet configured or an EVPN destination is created under the EVPN instance added as member. When configured, no other objects (for example, SAP, SDP-bind, BGP-EVPN instance) can be configured as part of the operational group within the same or different service.

The **no** form of this command disables the operational group.

Default

no oper-group

Parameters

name

Specifies the name of the operational group, up to 32 characters.

ingress-replication-bum-label

Syntax

[no] ingress-replication-bum-label

Context

config>service>vpls>bgp-evpn>mpls

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the system to send a separate label for Broadcast, Unknown unicast, and Multicast (BUM) traffic in a specified service. By default (**no ingress-replication-bum-label**), the same label is used for unicast and flooded BUM packets when forwarding traffic to remote PEs.

Saving labels may cause transient traffic duplication for all-active multihoming. If **ingress-replication-bum-label** is enabled, the system advertises two labels per EVPN VPLS instance, one for unicast and one for BUM traffic. The ingress PE uses the BUM label for flooded traffic to the advertising egress PE, which allows the egress PE to determine whether unicast traffic has been flooded by the ingress PE. Depending on the scale required in the network, the user may choose between saving label space or avoiding transient packet duplication sent to an all-active multi-homed CE for certain MACs.

The **no** form of this command uses the same label for unicast and flooded BUM packets.

Default

no ingress-replication-bum-label

split-horizon-group

Syntax

split-horizon-group *name*

no split-horizon-group

Context

config>service>vpls>bgp-evpn>mpls

Platforms

7210 SAS-Mxp

Description

This command configures 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 **bgp-evpn mpls split-horizon-group** command is not used, the default split-horizon group (that contains all the EVPN destinations) is still used, but it is not possible to refer to it on SAPs/spoke-SDPs.

User-configured split-horizon groups can be configured within the service context. The same group name can be associated with SAPs, spoke-SDPs, pw-templates, pw-template-bindings, and EVPN-MPLS destinations.

The configuration of the **bgp-evpn mpls split-horizon-group** command is only allowed if **bgp-evpn>mpls** is shut down; no changes are allowed when **bgp-evpn>mpls** is **no shutdown**.

If the SAPs or spoke-SDPs (manual) are configured within the same split-horizon group as the EVPN-MPLS endpoints, MAC addresses will still be learned but not advertised in BGP-EVPN. If an EVPN-MPLS

provider tunnel is enabled in the service, the SAPs and SDP-bindings that share the same split-horizon group of the EVPN-MPLS provider-tunnel will be brought operationally down if the point-to-multipoint tunnel is operationally up.

The **no** form of this command configures the EVPN-MPLS destinations to use the default split-horizon group.

Default

no split-horizon-group

Parameters

name

Specifies the split-horizon group name.

proxy-arp

Syntax

[no] proxy-arp

Context

config>service>vpls

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables proxy-ARP in an VPLS service.

On the 7210 SAS, users can enable or disable proxy-ARP commands for all EVPN services configured on the node; however, the option to enable or disable proxy-ARP per service is not available.

To enable or disable proxy-ARP capability, use the **config>service>system>evpn-proxy-arp-nd** command.

The **no** form of this command removes the proxy-ARP context.



Note:

If the **config>service>system>evpn-proxy-arp-nd** command is configured, it must be disabled to run the **no proxy-arp** command. See [Configuration guidelines for proxy-ARP and proxy-ND](#) for more information.

Default

no proxy-arp

proxy-nd

Syntax

[no] proxy-nd

Context

config>service>vpls

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables proxy-ND in a VPLS service.

On the 7210 SAS, users can enable or disable proxy-ND commands for all EVPN services configured on the node; however, the option to enable or disable proxy-ND per service is not available.

To enable or disable proxy-ND capability, use the **config>service>system>evpn-proxy-arp-nd** command.

The **no** form of this command removes the proxy-ND context.



Note:

If the **config>service>system>evpn-proxy-arp-nd** command is configured, the **no proxy-nd** command cannot be run. See [Configuration guidelines for proxy-ARP and proxy-ND](#) for more information.

Default

no proxy-nd

age-time

Syntax

age-time *seconds*

no age-time

Context

config>service>vpls>proxy-arp

config>service>vpls>proxy-nd

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command specifies the aging timer per proxy-ARP and proxy-ND entry for dynamic entries. When the aging expires, the entry is flushed. The age is reset when a new ARP, GARP, or NA for the same MAC-IP is received.

If the corresponding FDB MAC entry is flushed, the proxy-ARP or proxy-ND entry becomes inactive and subsequent ARP or NS lookups are treated as "missed". EVPN withdraws the IP-to-MAC if the entry becomes inactive. The **age-time** should be set at the **send-refresh seconds** value * 3 to ensure that no active entries are unnecessarily removed.

The **no** form of this command disables the aging timer.

Default

no age-time

Parameters

seconds

Specifies the aging time, in seconds.

Values 60 to 86400

dup-detect

Syntax

dup-detect [**anti-spoof-mac** *mac-address*] **window** *minutes* **num-moves** *count* **hold-down** [*minutes* | **max**]

Context

config>service>vpls>proxy-arp

config>service>vpls>proxy-nd

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables the mechanism that detects duplicate IPs and ARP/ND spoofing attacks. Attempts (relevant to dynamic and EVPN entry types) to add the same IP (different MAC) are monitored for **window** *minutes*. When *count* is reached within that **window**, the proxy-ARP or proxy-ND entry for the suspected IP is marked as duplicate. An alarm is also triggered. This condition is cleared when **hold-down** time expires (max does not expire) or a **clear** command is issued.

If the **anti-spoof-mac** keyword is configured, the proxy-ARP or proxy-ND MAC address of the offending entry is replaced with the configured anti-spoof *mac-address* and advertised in an unsolicited GARP/NA for local SAPs/SDP-bindings, and in EVPN to remote PEs. This mechanism assumes that the same **anti-spoof-mac** is configured in all the PEs for the same service, and that traffic with destination **anti-spoof-mac** received on SAPs/SDP-bindings will be dropped. An ingress **mac-filter** may be configured to drop traffic to the **anti-spoof-mac**.

Default

dup-detect window 3 num-moves 5 hold-down 9

Parameters**window *minutes***

Specifies the window size, in minutes.

Values 1 to 15

Default 3

count

Specifies the number of moves required so that an entry is declared duplicate.

Values 3 to 10

Default 5

hold-down *minutes*

Specifies the hold-down time, in minutes, for a duplicate entry.

Values 2 to 60 | max

Default 9

mac-address

Specifies the MAC address to use as the optional anti-spoof-mac.

dynamic-arp-populate**Syntax**

[no] dynamic-arp-populate

Context

config>service>vpls>proxy-arp

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables the addition of dynamic entries to the proxy-ARP table.

When enabled, the system populates proxy-ARP entries from snooped GARP or ARP messages on SAPs/SDP-bindings. These entries are shown as dynamic.

When disabled, dynamic ARP entries are flushed from the proxy-ARP table. Enabling **dynamic-arp-populate** is only recommended in networks where this command is consistently configured in all PEs.

The **no** form of this command disables the addition of dynamic entries to the proxy-ARP table.

Default

no dynamic-arp-populate

dynamic-nd-populate**Syntax**

[no] dynamic-nd-populate

Context

config>service>vpls>proxy-nd

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables the addition of dynamic entries to the proxy-ND table.

When enabled, the system populates proxy-ND entries from snooped Neighbor Advertisement (NA) messages on SAPs or SDP-bindings, in addition to the entries coming from EVPN (if the EVPN is enabled). These entries are shown as dynamic, and not as EVPN or static entries.

When disabled, dynamic ND entries are flushed from the proxy-ND table. Enabling **dynamic-nd-populate** is only recommended in networks where this command is consistently configured in all PEs.

The **no** form of this command disables the addition of dynamic entries to the proxy-ND table.

Default

no dynamic-nd-populate

evpn-nd-advertise**Syntax**

evpn-nd-advertise {host | router}

Context

config>service>vpls>proxy-nd

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables the advertisement of static or dynamic entries that are learned as a host or router. Only one option (host or router) is possible in a specified service. This command also determines the R flag (host or router) when sending NA messages for existing EVPN entries in the proxy-ND table.

This command can only be modified if **proxy-nd** is shut down.

Default

evpn-nd-advertise router

Parameters**host**

Keyword to enable the advertisement of static or dynamic entries that are learned as host.

router

Keyword to enable the advertisement of static or dynamic entries that are learned as routers.

garp-flood-evpn

Syntax

[no] garp-flood-evpn

Context

config>service>vpls>proxy-arp

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command controls whether the system floods GARP-requests and GARP-replies to the EVPN. The GARPs impacted by this command are messages in which the sender IP is equal to the target IP and the MAC DA is broadcast.

The **no** form of this command only floods to local SAPs/SDP-bindings but not to EVPN destinations. The use of the **no** form is only recommended in networks where CEs are routers that are directly connected to the PEs. Networks using aggregation switches between the host/routers and the PEs should flood GARP messages in the EVPN to ensure that the remote caches are updated and the BGP does not miss the advertisement of these entries.

Default

garp-flood-evpn

host-unsolicited-na-flood-evpn

Syntax

[no] host-unsolicited-na-flood-evpn

Context

config>service>vpls>proxy-nd

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command controls whether the system floods host unsolicited Neighbor Advertisement (NA) messages to the EVPN. The NA messages with the following flags are impacted by this command:

- S=0
- R=0

The **no** form of this command only floods to local SAPs/SDP-bindings but not to the EVPN destinations. The use of the **no** form is only recommended in networks where CEs are routers that are directly connected to the PEs. Networks using aggregation switches between the host/routers and the PEs should flood unsolicited NA messages in the EVPN to ensure that the remote caches are updated and the BGP does not miss the advertisement of these entries.

Default

host-unsolicited-na-flood-evpn

router-unsolicited-na-flood-evpn

Syntax

[no] router-unsolicited-na-flood-evpn

Context

config>service>vpls>proxy-nd

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command controls whether the system floods router unsolicited NAs to EVPN. The NA messages impacted by this command are NA messages with the following flags:

- S=0
- R=1

The **no** form of this command only floods to local SAPs/SDP-bindings but not to EVPN destinations. This is only recommended in networks where CEs are routers directly connected to the PEs. Networks using aggregation switches between the host/routers and the PEs should flood unsolicited NA messages in EVPN to ensure that the remote caches are updated and the BGP does not miss the advertisement of these entries.

Default

router-unsolicited-na-flood-evpn

send-refresh

Syntax

send-refresh *seconds*

no send-refresh

Context

config>service>vpls>proxy-arp

config>service>vpls>proxy-nd

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables the system to send a refresh message at the configured time. A refresh message is an ARP-request message that uses 0s as the sender IP for the case of a proxy-ARP entry. For proxy-ND entries, a refresh is a regular NS message that uses the chassis MAC address as the MAC source address.

The **no** form of this command suppresses the refresh messages.

Default

no send-refresh

Parameters

seconds

Specifies the time to send a refresh message, in seconds.

Values 120 to 86400

static

Syntax

static *ip-address ieee-address*

no static *ip-address*

Context

config>service>vpls>proxy-arp

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures static entries to be added to the table. A static MAC-IP entry requires the addition of the MAC address to the FDB as either learned or CStatic (conditional static MAC) to become active.

The **no** form of this command removes the specified static entry.

Parameters

ip-address

Specifies the IPv4 address for the static entry.

ieee-address

Specifies a 48-bit MAC address in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx represents a hexadecimal number.

static

Syntax

static *ipv6-address* *ieee-address* {**host** | **router**}

no static *ipv6-address*

Context

config>service>vpls>proxy-nd

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures static entries to be added to the table. A static MAC-IP entry requires the addition of the MAC address to the FDB as either dynamic or CStatic (Conditional Static MAC) to become active. Along with the IPv6 and MAC address, the entry must also be configured as either host or router. This determines whether the received NS for the entry is replied with the R flag set to 1 (router) or 0 (host).

The **no** form of this command removes the specified static entry.

Parameters

ipv6-address

Specifies the IPv6 address for the static entry.

ieee-address

Specifies a 48-bit MAC address in the form xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx, where xx represents a hexadecimal number.

host

Specifies that the entry is type "host".

router

Specifies that the entry is type "router".

table-size

Syntax

table-size *table-size*

Context

config>service>vpls>proxy-arp

config>service>vpls>proxy-nd

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command adds a table-size limit per service. By default, the limit is 250; it can be set up to 16k entries per service. A non-configurable implicit high watermark of 95% and low watermark of 90% exists, per service and per system.

When those watermarks are reached, a syslog or trap is triggered. When the system or service limit is reached, entries for a specified IP can be replaced (a different MAC can be learned and added) but no new IP entries are added, regardless of the type (Static, evpn, dynamic). If the user attempts to change the *table-size* value to a value that cannot accommodate the number of existing entries, the attempt fails.

Default

table-size 250

Parameters

table-size

Specifies the table-size as the number of entries for the service.

Values 1 to 16384

unknown-arp-request-flood-evpn

Syntax

[no] **unknown-arp-request-flood-evpn**

Context

config>service>vpls>proxy-arp

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command controls whether unknown ARP requests are flooded into the EVPN network. By default, the system floods ARP requests, including EVPN (with source squelching), if there is no active proxy-ARP entry for the requested IP.

The **no** form of this command only floods to local SAPs/SDP-bindings and not to EVPN destinations.

Default

unknown-arp-request-flood-evpn

unknown-ns-flood-evpn

Syntax

[no] unknown-ns-flood-evpn

Context

config>service>vpls>proxy-nd

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables unknown Neighbor Solicitation (NS) messages to be flooded into the EVPN network. By default, the system floods NS (with source squelching) to SAPs/SDP-bindings including EVPN, if there is no active proxy-ND entry for the requested IPv6.

The **no** form of this command only floods to local SAPs/SDP-bindings but not to EVPN destinations.

Default

unknown-ns-flood-evpn

shutdown

Syntax

[no] shutdown

Context

config>service>vpls>proxy-arp

config>service>vpls>proxy-nd

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables and disables the proxy-ARP and proxy-ND functionalities. ARP, GARP, and ND messages are snooped and redirected to the CPM for lookup in the proxy-ARP or proxy-ND table. The proxy-ARP or proxy-ND table is populated with IP-to-MAC pairs received from different sources (EVPN, static, dynamic). When the **shutdown** command is issued, the system stops snooping ARP or ND frames and the dynamic/EVPN dup proxy-ARP or proxy-ND table entries are flushed. All the static entries are kept in the table as "inactive", regardless of their previous "Status".



Note:

The **proxy-arp shutdown** and **no shutdown**, and **proxy-nd shutdown** and **no shutdown** commands cannot be executed if the **config>service>system>evpn-proxy-arp-nd** command is configured.

The **no** form of this command enables the proxy-ARP and proxy-ND functionalities.

Default

shutdown

ethernet-segment

Syntax

ethernet-segment *name* [**create**]

no ethernet-segment *name*

Context

config>service>system>bgp-evpn

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures an ES instance and its corresponding name.

The **no** form of this command deletes the specified ES.

Parameters

name

Specifies the ES name, up to 28 characters.

create

Keyword to create an ES.

es-activation-timer

Syntax

es-activation-timer *seconds*

no es-activation-timer

Context

config>service>system>bgp-evpn>ethernet-segment

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the ES activation timer for the specified **ethernet-segment**. The **es-activation-timer** delays the activation of a specified **ethernet-segment** on a specified PE that has been elected as DF (Designated Forwarder). Only when the **es-activation-timer** has expired, the SAP associated to an **ethernet-segment** can be activated (in case of single-active multihoming).

The **no** form of this command specifies that the system uses the value in the **config>redundancy>bgp-evpn-multi-homing>es-activation-timer** context, if configured. Otherwise the system uses the default value of 3 seconds.

Default

no es-activation-timer

Parameters

seconds

Specifies the number of seconds for the **es-activation-timer**.

Values 0 to 100

Default 3

esi

Syntax

esi *esi*

no esi

Context

config>service>system>bgp-evpn>ethernet-segment

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the 10-byte Ethernet segment identifier (ESI) associated with the **ethernet-segment** that will be signaled in the BGP-EVPN routes. The ESI value cannot be changed unless the **ethernet-segment** is **shutdown**. Reserved ESI values, 0 and MAX-ESI, are not allowed.

The **no** form of this command deletes the ESI from the Ethernet segment.

Default

no esi

Parameters

esi

Specifies the 10-byte ESI in the form 00-11-22-33-44-55-66-77-88-99, using "-", ":", or " " as separators.

lag

Syntax

lag *lag-id*

no lag

Context

config>service>system>bgp-evpn>ethernet-segment

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures a lag ID associated to the ES. When the **ethernet-segment** is configured as **all-active**, only a LAG can be associated to the ES. When the **ethernet-segment** is configured as **single-active**, a LAG or port can be associated to the ES. In either case, only one of the two objects can be configured in the ES. A specified LAG can be part of only one ES

The **no** form of this command removes the association of the Ethernet segment to LAG ports.

Default

no lag

Parameters

lag-id

Specifies the lag ID associated with the ES.

Values 1 to 25

multi-homing

Syntax

multi-homing single-active no-esi-label

no multi-homing

Context

config>service>system>bgp-evpn>ethernet-segment

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the multi-homing mode for the specified **ethernet-segment** as **single-active** multi-homing, as defined in RFC7432.



Note:

The **esi-label** option cannot be enabled for **single-active**.

When **single-active no-esi-label** is specified, the system does not allocate an ESI label and advertise ESI label 0 to peers. The 7210 SAS does not use the ESI label received from a peer to send traffic to that peer.

The **multi-homing** command must be configured for the Ethernet segment to be enabled.

The **no** form of this command disables multi-homing on the Ethernet segment.

Default

no multi-homing

Parameters

single-active

Specifies single-active mode for the ES.

no-esi-label

Specifies that the system does not send an ESI label for **single-active** mode.

port

Syntax

port *port-id*

no port

Context

config>service>system>bgp-evpn>ethernet-segment

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures a port ID associated with the ES. If the **ethernet-segment** is configured as **single-active**, a LAG or port can be associated to the ES. In any case, only one of the two objects can be configured in the ES. A specified port can be part of only one **ethernet-segment**. Only Ethernet ports can be added to an **ethernet-segment**.

The **no** form of this command removes the ES association to all ports.

Default

no port

Parameters

port-id

Specifies the port ID associated to the ES.

Values	<i>port-id</i>	<i>slot/mda/port [.channel]</i>

service-carving

Syntax

service-carving

Context

config>service>system>bgp-evpn>ethernet-segment

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure service-carving in the Ethernet segment. The service-carving algorithm determines the PE that is the Designated Forwarder (DF) in a specified ES and for a specific service.

manual

Syntax

manual

Context

config>service>system>bgp-evpn>eth-seg>service-carving

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context manually configure the service-carving algorithm; that is, configure the EVIs for which the PE is DF.

evi

Syntax

evi start [to to] primary

no evi start

Context

config>service>system>bgp-evpn>eth-seg>service-carving>manual

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the EVI ranges for which the PE is DF.



Note:

Multiple individual EVI values and ranges are allowed. The PE will be non-DF for the **evi** values not defined as **primary**.

The **no** form of this command removes the specified EVI range.

Parameters

start

Specifies the initial EVI value of the range for which the PE is DF.

Values 1 to 65535

to

Specifies the end EVI value of the range for which the PD is DF. If not configured, only the individual start value is considered.

Values 1 to 65535

primary

Specifies that the PE is DF for the configured EVI range.

mode**Syntax**

mode {**manual** | **auto** | **off**}

Context

config>service>system>bgp-evpn>eth-seg>service-carving

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the **service-carving** mode. This determines how the DF is elected for a specified ES and service.

Default

mode auto

Parameters**auto**

Specifies the service-carving algorithm defined in RFC 7432. The DF for the service is calculated based on the modulo function of the service (identified by either the EVI or the ISID) and the number of PEs.

manual

Specifies that the DF is elected based on the manual configuration added in the **service-carving>manual** context.

off

Specifies that all the services elect the same DF PE (assuming the same PEs are active for all the configured services). The PE with the lowest IP is elected as DF for the ES.

shutdown**Syntax**

[no] **shutdown**

Context

```
config>service>system>bgp-evpn>ethernet-segment
```

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command changes the administrative status of the **ethernet-segment**.

The user can only configure **no shutdown** when **esi**, **multi-homing**, and **lag/port** are configured. If the ES or the corresponding **lag/port** are **shutdown**, the ES 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
```

```
config>service>epipe>bgp
```

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the route distinguisher (RD) that will be signaled in EVPN Type 4 routes (Ethernet segment routes).

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

Default

no route-distinguisher

Parameters

rd

Specifies the RD in the following format.

- *ip-addr.comm-val*

Values *ip-addr* — a.b.c.d

comm-val — 0 to 65535

Default system-ip: 0

redundancy

Syntax
redundancy

Context
config

Platforms
7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description
Commands in this context configure the global redundancy parameters.

bgp-evpn-multi-homing

Syntax
bgp-evpn-multi-homing

Context
config>redundancy

Platforms
7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description
Commands in this context configure the BGP-EVPN global timers.

boot-timer

Syntax
boot-timer *seconds*

Context
config>redundancy>bgp-evpn-multi-homing

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

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 segments and running the DF algorithm.

The following considerations apply to the functionality:

- The boot-timer is configured at the system level. The configured value must provide enough time to allow the node and the cards (if available) to come up and BGP sessions to come up before exchanging ES routes and running the DF election for each EVI.
- The boot-timer is synchronized across CPMs and is relative to the System UP-time; therefore the boot-timer is not subject to change or reset upon CPM switchover.
- The boot-timer is never interrupted (however, the **es-activation-timer** can be interrupted if there is a new event triggering the DF election).
- The boot-timer runs per EVI on the ES's in the system. While **system-up-time>boot-timer** is true, the system does not run the DF election for any EVI. 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** kicks in.
- 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 number of seconds for the boot-timer.

Values 0 to 600

es-activation-timer

Syntax

es-activation-timer *seconds*

Context

config>redundancy>bgp-evpn-multi-homing

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the global Ethernet segment activation timer. The **es-activation-timer** delays the activation of a specified Ethernet segment on a specified PE that has been elected as the DF

(Designated Forwarder). Only when the **es-activation-timer** has expired, can the SAP/SDP-binding associated to an Ethernet segment be activated (in case of single-active multi-homing) or added to the default-multicast-list (in case of all-active multi-homing).

The **es-activation-timer** configured at the Ethernet-segment level supersedes this global **es-activation-timer**.

Default

es-activation-timer 3

Parameters

seconds

Specifies the number of seconds for the **es-activation-timer**.

Values 0 to 100

4.5.2.2 EVPN show commands

evpn-mpls

Syntax

evpn-mpls

Context

show>service

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays the remote EVPN-MPLS tunnel endpoints in the system.

Output

The following output is an example of EVPN MPLS tunnel endpoint information, and [Table 41: Output fields: EVPN MPLS tunnel endpoints](#) describes the output fields.

Sample output

```
*A:Dut-B# show service evpn-mpls
=====
EVPN MPLS Tunnel Endpoints
=====
EvpnMplsSTEP Address EVPN-MPLS Dest      ES Dest
-----
10.20.1.3      1              0
10.20.1.4      1              0
10.20.1.5      1              0
-----
```

```

Number of EvpnMpls Tunnel Endpoints: 3
-----
=====

```

Table 41: Output fields: EVPN MPLS tunnel endpoints

Label	Description
EvpnMplsTEP	Displays the tunnel endpoint addresses
EVPN-MPLS Dest	Displays the number of EVPN-MPLS destinations
ES Dest	Displays the Ethernet segment destination

bgp-evpn

Syntax

bgp-evpn

Context

show>service>id

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays the **bgp-evpn** configured parameters for a specified service, including the administrative status of MPLS, the configuration for **mac-advertisement** and **unknown-mac-route**, as well as the **mac-duplication** parameters. The command shows the duplicate MAC addresses that **mac-duplication** has detected.

If the service is BGP-EVPN MPLS, this command also shows the parameters corresponding to EVPN-MPLS.

Output

The following output is an example of BGP EVPN information for a specified service, and [Table 42: Output fields: service ID BGP-EVPN](#) describes the output fields.

Sample output

```

*A:Dut-B# /show service id 1 bgp-evpn
=====
BGP EVPN Table
=====
MAC Advertisement   : Enabled
CFM MAC Advertise  : Disabled
MAC Dup Detn Moves  : 5           MAC Dup Detn Window: 3
MAC Dup Detn Retry  : 9           Number of Dup MACs : 0
EVI                 : 1
-----
Detected Duplicate MAC Addresses      Time Detected

```



```

=====
BGP EVPN MPLS Information
=====

```

```

Admin Status      : Enabled
Force Vlan Fwding : Disabled      Control Word      : Disabled
Split Horizon Group: (Not Specified)
Ingress Rep BUM Lbl: Disabled      Max Ecmp Routes   : 0
Ingress Ucast Lbl : 131069        Ingress Mcast Lbl : 131069
=====

```

```

=====
BGP EVPN MPLS Auto Bind Tunnel Information
=====

```

```

Resolution      : any
Filter Tunnel Types: (Not Specified)
=====

```

Table 42: Output fields: service ID BGP-EVPN

Label	Description
BGP EVPN Table	
MAC Advertisement	Displays whether MAC advertisement is enabled or disabled
CFM MAC Advertise	Displays whether CFM MAC advertise is enabled or disabled
MAC Dup Detn Moves	Displays the number of moves that trigger MAC duplication detection
MAC Dup Detn Window	Displays the configured window size used for duplicate MAC detection
MAC Dup Detn Retry	Displays the retry timer value used for MAC duplication detection.
Number of Dup MACs	Displays the number of duplicate MAC addresses
EVI	Displays the EVPN instance ID
BGP EVPN MPLS Information	
Admin Status	Displays the administrative status of the EVPN MPLS
Force Vlan Fwding	Displays the status of force-vlan-forwarding
Control Word	Displays the status of control
Split Horizon Group	Displays the split-horizon group membership information
Ingress Rep BUM Lbl	Displays the label used for Ingress BUM replication
Max Ecmp Routes	Displays the maximum number of ECMP routes
Ingress Ucast Lbl	Displays the ingress unicast label

Label	Description
Ingress Mcast Lbl	Displays the ingress multicast label
BGP EVPN MPLS Auto Bind Tunnel Information	
Resolution	Displays the transport tunnel resolution filter used
Filter Tunnel Types	Displays auto-bind-tunnel resolution filter values, if applicable

evpn-mpls

Syntax

evpn-mpls

evpn-mpls esi esi

Context

show>service>id

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays the existing EVPN-MPLS destinations for a specified service and all related information. Filtering based on **esi** (for EVPN multihoming) is supported to display the EVPN-MPLS destinations associated with an Ethernet Segment Identifier (ESI).

Parameters

esi

Specifies a 10-byte ESI by which to filter the displayed information. For example, ESI-0 | ESI-MAX or 00-11-22-33-44-55-66-77-88-99 with any of these separators ('-',':',';',' ').

Output

The following output is an example of EVPN MPLS information, and [Table 43: Output fields: EVPN MPLS](#) describes the output fields.

Sample output

```
*A:Dut-B# /show service id 1 evpn-mpls
=====
BGP EVPN-MPLS Dest
=====
TEP Address      Egr Label      Num. MACs      Mcast          Last Change
      Transport
-----
10.20.1.3        131069         0              Yes            02/02/2014 15:29:40
                  rsvp
10.20.1.4        131069         0              Yes            02/02/2014 15:29:33
                  rsvp
```

```

10.20.1.5      131059      0      Yes      02/02/2014 15:29:42
                rsvp
-----
Number of entries : 3
=====
BGP EVPN-MPLS Ethernet Segment Dest
=====
Eth SegId      Num. Macs      Last Change
-----
00:de:01:00:00:00:00:00:01  1      02/02/2014 15:47:04
-----
Number of entries: 1
-----
*A:PE-1# show service id 2 evpn-mpls esi 00:10:00:00:00:00:00:00:00

```

Table 43: Output fields: EVPN MPLS

Label	Description
TEP Address	Displays the TEP address
Egr Label	Displays the egress label
Transport	Displays the transport type
Number of entries	Indicates the number of entries
Eth SegId	Displays the Ethernet segment ID
Transport:Tnl-Id	Displays the tunnel type and tunnel ID of the EVPN-MPLS entry
Transport:Tnl	Displays the transport tunnel
Num. MAC	Displays the number of MACs
Mcast	Displays the multicast information
Sup BCast Domain	Displays the Sup BCast Domain

proxy-arp

Syntax

proxy-arp [*ip-address*] [**detail**]

Context

show>service>id

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays, in a table, the existing proxy-ARP entries for a specified service. The table is populated by EVPN MAC routes that contain a MAC and an IP address, as well as static entries or dynamic entries from snooped ARP messages on access SAPs.

A 7210 SAS that receives an ARP request from a SAP performs a lookup in the proxy-ARP table for the service. If a match is found, the router replies to the ARP and does not allow ARP flooding in the VPLS service. If a match is not found, the ARP is flooded within the service if the configuration allows it.

This command allows for specific IP addresses to be displayed. Dynamic IP entries associated with a MAC list are displayed with the corresponding MAC list and resolve timers information.

Parameters

ip-address

Specifies an IP address.

Values a.b.c.d

detail

Displays more information.

Output

The following output is an example of proxy-ARP information for a specified service, and [Table 44: Output fields: proxy-ARP](#) describes the output fields.

Sample output

```
show service id 1 proxy-arp detail
-----
Proxy Arp
-----
Admin State       : enabled
Dyn Populate      : enabled
Age Time          : disabled          Send Refresh      : disabled
Table Size        : 16383             Total              : 2
Static Count      : 0                  EVPN Count         : 1
Dynamic Count     : 1                  Duplicate Count    : 0
Dup Detect
-----
Detect Window     : 3 mins             Num Moves          : 5
Hold down         : 9 mins
Anti Spoof MAC    : None
EVPN
-----
Garp Flood        : disabled           Req Flood          : enabled
-----
=====
VPLS Proxy Arp Entries
=====
IP Address        Mac Address      Type   Status   Last Update
-----
10.1.1.1          00:00:00:00:00:01 dyn    active   03/13/2020 10:25:39
10.1.1.10         00:00:00:00:00:11 evpn    active   03/13/2020 10:25:40
-----
Number of entries : 2
=====
```

Table 44: Output fields: proxy-ARP

Label	Description
Admin State	Displays the admin state: enabled or disabled
Dyn Populate	Displays the status of the ARP dynamic population
Age Time	Displays the configured ARP age timer
Send Refresh	Displays the configured ARP refresh timer
Table Size	Displays the configured ARP table size
Total	Displays the total table used count
Static Count	Displays the static ARP entries count
EVPN Count	Displays the count of ARP entries learned through the EVPN tunnel
Dynamic Count	Displays the count of ARP entries dynamically learned
Duplicate Count	Displays the count of ARP duplicate entries
Detect Window	Displays the configured window value for ARP duplicate detection
Num Moves	Displays the configured count for number of moves used for ARP duplicate detection
Hold Down	Displays the hold-down timer used by ARP duplicate detection
Anti Spoof MAC	Displays the MAC address configured for anti-spoof detection
Garp Flood	Displays the status for GARP flooding
Req Flood	Displays the status of ARP request flooding
IP Address	Displays the IP address of the proxy-ARP entry
Mac Address	Displays the MAC address of the proxy-ARP entry
Type	Displays the type of ARP entry
Status	Displays the status
Last Update	Displays the date and time of the last update

proxy-nd

Syntax

proxy-nd [*ipv6-address*] [**detail**]

Context

show>service>id

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays, in a table, the existing proxy-ND entries for a specified service. The table is populated by the EVPN MAC routes containing a MAC and an IPv6 address, as well as static entries or dynamic entries from snooped NA messages on access SAPs.

A 7210 SAS that receives a Neighbor Solicitation (NS) from a SAP performs a lookup in the proxy-ND table for the service. If a match is found, the router replies to the NS and does not allow NS flooding in the VPLS service. If a match is not found, the NS is flooded in the service, if the configuration allows it.

This command allows specific IPv6 addresses to be displayed. Dynamic IPv6 entries associated with a MAC list are shown with the corresponding MAC list and resolve timer information.

Parameters

ipv6-address

Specifies an IPv6 address.

Values ipv6-address:
 x:x:x:x:x:x:x:x (eight 16-bit pieces)
 x:x:x:x:x:x:d.d.d.d
 where:
 x - [0 to FFFF]H
 d - [0 to 255]D

detail

Displays more information.

Output

The following output is an example of proxy-ND information for a specified service, and [Table 45: Output fields: proxy-ND](#) displays the output fields.

Sample output

```
A:Dut-C# show service id 1 proxy-nd detail
-----
Proxy ND
-----
```

```

Admin State      : enabled
Dyn Populate     : enabled
Age Time        : disabled
Table Size      : 250
Static Count     : 0
Dynamic Count    : 1
Dup Detect      :
-----
Detect Window   : 3 mins
Hold down       : 9 mins
Anti Spoof MAC  : None
EVPN
-----
Unknown NS Flood : enabled
Rtr Unsol NA Flood : disabled
ND Advertise     : Router
Host Unsol NA Fld : disabled
=====
VPLS Proxy ND Entries
=====
IP Address      Mac Address      Type Status Rtr/ Last Update
                                   Host
-----
2000::4         00:00:00:00:00:04 dyn active Rtr 01/14/2020 09:47:43
-----
Number of entries : 1*A:PE-2# show service id 5 proxy-nd

```

Table 45: Output fields: proxy-ND

Label	Description
Admin State	Displays the admin state for proxy-ND: enabled or disabled
Dyn Populate	Displays the status for dynamic populate
Age Time	Displays the aging timer for ND entries
Send Refresh	Displays the refresh timer for ND entries
Table Size	Displays the proxy-ND table size
Total	Displays the count of learned ND entries
Static Count	Displays the count of static ND entries
EVPN Count	Displays the count of ND entries learned from the EVPN binding
Dynamic Count	Displays the count of dynamically learned ND entries
Duplicate Count	Displays the count of duplicate ND entries
Detect Window	Displays the configured value for window size used for duplicate detection
Num Moves	Displays the configured value for number of moves used in duplicate ND detection
Hold Down	Displays the value of the hold-down timer

Label	Description
Anti Spoof MAC	Displays the configured anti-spoof MAC address
Unknown NS Flood	Displays the state of unknown Neighbor Solicitation messages that are flooded to the EVPN network
ND Advertise	Displays the advertisement of static or dynamic entries that are learned as hosts or routers
Rtr Unsol NA Flood	Displays the state of system floods router unsolicited Neighbor Advertisements to EVPN
Host Unsol NA Fld	Displays the state of system floods host unsolicited Neighbor Advertisements to EVPN
IP Address	Displays the IP address of the proxy-ND entry
Mac Address	Displays the MAC address of the proxy-ND entry
Type	Displays the type of ND entry
Status	Displays the status of the ND entry
Last Update	Displays the date and time of the last update

system

Syntax

system

Context

show>service

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure the **system** BGP EVPN **show** commands.

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

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays system BGP EVPN information.

Parameters

ethernet-segment

Displays Ethernet Segment (ES) information.

name

Specifies the ES name, up to 28 characters.

all

Displays all available information for the specified ES.

evi

Displays information for the specified EVI.

Values 1 to 65535

Output

The following output is an example of system BGP EVPN information, and [Table 46: Output fields: system BGP-EVPN](#) describes the output fields.

Sample output

```
*A:Dut-B# /show service system bgp-evpn
=====
System BGP EVPN Information
=====
Evpn Route Dist.           : <none>
Oper Route Dist.           : 10.20.1.2:0
Oper Route Dist Type       : default
=====
```

Table 46: Output fields: system BGP-EVPN

Label	Description
Evpn Route Dist.	Displays the EVPN route distinguisher
Oper Route Dist.	Displays address of the operational route distinguisher
Oper Route Dist Type	Displays the operational route distinguisher type

redundancy

Syntax

redundancy

Context

show

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context display the global redundancy parameters.

bgp-evpn-multi-homing

Syntax

bgp-evpn-multi-homing

Context

show>redundancy

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays information related to the EVPN global timers.

Output

The following output is an example of BGP EVPN multi-homing information, and [Table 47: Output Fields: BGP-EVPN multi-homing](#) displays the output fields.

Sample output

```
*A:Dut-B# 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 47: Output Fields: BGP-EVPN multi-homing

Label	Description
Boot-Timer	Displays the value configured for the boot timer
Boot-Timer Remaining	Displays the amount of time remaining on the boot timer
ES Activation Timer	Displays the value configured for the ES activation timer

4.5.2.3 EVPN clear commands

proxy-arp

Syntax

proxy-arp [duplicate] [dynamic]

Context

clear>service>id

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command clears all entries in the proxy-ARP table if none of the optional parameters is specified. If the duplicate parameter is specified it clears all the duplicate entries in the proxy-ARP table. If the dynamic parameter is specified it clears all the dynamic entries in the proxy-ARP table.

Parameters

duplicate

Clears the proxy ARP duplicate entries.

dynamic

Clears the proxy ARP dynamic entries.

proxy-nd

Syntax

proxy-nd [duplicate] [dynamic]

Context

clear>service>id

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command clears all entries in the proxy-ND table if none of the optional parameters is specified. If the duplicate parameter is specified it clears all the duplicate entries in the hold-down state from the proxy-ND table. If the dynamic parameter is specified it clears all the dynamic entries in the hold-down state from the proxy-ND table.

Parameters

duplicate

Clears the proxy ND duplicate entries.

dynamic

Clears the proxy ND dynamic entries.

4.5.2.4 Tools commands

service

Syntax

service

Context

tools>dump

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures tools to display service dump information.

proxy-arp

Syntax

proxy-arp usage

Context

tools>dump>service

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command provides information about the usage and limit of the system-wide proxy-ARP table for all services. The command also shows if the limit has been exceeded and a trap raised.

Output

The following output is an example of **tools dump service proxy-arp** usage information.

Sample output

```
*A:Dut# tools dump service proxy-arp usage
Proxy arp Usage
      Current Usage      :      10
      System Limit       :    16384
High Usage Trap Raised:   No
      High Usage Threshold:    95 percent
      High Usage Clear Threshold: 90 percent
```

proxy-nd

Syntax

proxy-nd usage

Context

tools>dump>service

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command provides information about the usage and limit of the system-wide proxy-ND table for all the services. The command also shows if the limit has been exceeded and a trap raised.

Output

The following output is an example of **tools dump service proxy-nd** usage information.

Sample output

```
*A:Dut# tools dump service proxy-nd usage
Proxy nd Usage
      Current Usage      :      211
      System Limit       :    16384
High Usage Trap Raised:   No
      High Usage Threshold:    95 percent
      High Usage Clear Threshold: 90 percent
```

5 Virtual Private LAN Service

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

5.1 VPLS service overview

Virtual Private LAN Service (VPLS) is a class of virtual private network service that allows the connection of multiple sites in a single bridged domain over a provider-managed IP/MPLS network. The customer sites in a VPLS instance appear to be on the same LAN, regardless of their location. VPLS uses an Ethernet interface on the customer-facing (access) side that simplifies the LAN/WAN boundary and allows for rapid and flexible service provisioning. The 7210 SAS supports provisioning of access or uplink spokes to connect to the provider edge (PE) IP/MPLS routers.

VPLS provides a balance between point-to-point Frame Relay service and outsourced routed services (VPRN). VPLS enables customers to maintain control of their own routing strategies. All customer routers in the VPLS service are part of the same subnet (LAN), which simplifies the IP addressing plan, especially when compared to a mesh constructed from many separate point-to-point connections. 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 (which is considered a local service) or more (which is considered a distributed service) service routers. The connection appears to be a bridged domain to the customer sites, so protocols, including routing protocols, can traverse the VPLS service.


Other VPLS advantages include the following:

- VPLS is a transparent, protocol-independent service.
- There is no Layer 2 protocol conversion between LAN and WAN technologies.
- There is no need to design, manage, configure, and maintain separate WAN access equipment, which eliminates the need to train personnel on WAN technologies, such as Frame Relay.

5.1.1 VPLS packet walk-through in network mode

This section provides an example of VPLS processing of a customer packet sent across the network from site-A, which is connected to PE-Router-A through a 7210 SAS to site-C, which is connected through 7210 SAS to PE-Router-C (shown in the following figure) in an H-VPLS configuration. This section does not describe the processing on the PE routers, but only on 7210 SAS routers.


Figure 58: VPLS service architecture



1. 7210-A (shown in the following figure)

- a. Service packets arriving at 7210-A are associated with a VPLS service instance based on the combination of the physical port and the IEEE 802.1Q tag (VLAN-ID) in the packet.

Figure 59: Access port ingress packet format and lookup



- b. 7210-A learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the service access point (SAP) on which it was received.
- c. The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. There are two possibilities: either the destination MAC address has already been learned (known MAC address), or the destination MAC address is not yet learned (unknown MAC address).


For a Known MAC Address (shown in the following figure):

- d. If the destination MAC address has already been learned by 7210, an existing entry in the FIB table identifies the far-end PE-Router and the service VC-label (inner label) to be used before sending the packet to PE-Router-A.
- e. 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 (shown in the following figure):

- f. If the destination MAC address has not been learned, 7210 floods the packet to spoke-SDPs that are participating in the service.

Figure 60: Network port egress packet format and flooding



2. Core Router Switching

- a. The PE router encapsulates this packet in an MPLS header and transports it across the core network to the remote 7210-C.

3. 7210-C ([Figure 59: Access port ingress packet format and lookup](#))


- a. 7210-C associates the packet with the VPLS instance based on the VC label in the received packet after the stripping of the tunnel label.
- b. 7210-C learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the spoke-SDP on which the packet was received.
- c. The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. Again, there are two possibilities: either the destination MAC address has already been learned (known MAC address), or the destination MAC address has not been learned on the access side of 7210-C (unknown MAC address).
- d. If the destination MAC address has been learned by 7210-C, an existing entry in the FIB table identifies the local access port and the IEEE 802.1Q tag (if any) to be added before sending the packet to customer Location-C. The egress Q tag may be different from the ingress Q tag.
- e. If the destination MAC address has not been learned, 7210 floods the packet to all the access SAPs that are participating in the service.

5.1.2 VPLS packet walk-through in access-uplink mode

This section provides an example of VPLS processing of a customer packet sent across the network from site-A, which is connected to PE-Router-A through a 7210 SAS to site-C, which is connected through 7210 SAS to PE-Router-C ([Figure 58: VPLS service architecture](#)) in an H-VPLS configuration. This section does not describe the processing on the PE routers but only on 7210 SAS routers:


1. 7210-A (shown in the following figure)
- a. Service packets arriving at 7210-A are associated with a VPLS service instance based on the combination of the physical port and the IEEE 802.1Q tag (VLAN-ID) in the packet.

Figure 61: Access port ingress packet format and lookup



- b. 7210-A learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the service access point (SAP) on which it was received.
- c. The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. There are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address is not yet learned (unknown MAC address).
- For a Known MAC Address (Figure 60: Network port egress packet format and flooding):
- d. If the destination MAC address has already been learned by 7210, an existing entry in the FIB table identifies destination uplink QinQ SAP to be used for sending the packet toward the PE-Router-A.
- e. The customer packet is sent on this uplink SAP when the IEEE 802.1Q tag is stripped and the uplink SAP tag is added to the packet.
- For an Unknown MAC Address (shown in the following figure):
- f. If the destination MAC address has not been learned, 7210 will flood the packet to all the uplink SAP spoke-SDPs that are participating in the service.

Figure 62: Network port egress packet format and flooding



2. Core Router Switching
- a. The PE router will encapsulate this packet in the appropriate MPLS header and transport it across the core network to the remote 7210-C.
3. 7210-C (Figure 59: Access port ingress packet format and lookup)

- a. 7210-C associates the packet with the VPLS instance based on the VLAN tags in the received packet.
- b. 7210-C learns the source MAC address in the packet and creates an entry in the FIB table that associates the MAC address to the access-uplink port on which the packet was received.
- c. The destination MAC address in the packet is looked up in the FIB table for the VPLS instance. Again, there are two possibilities: either the destination MAC address has already been learned (known MAC address) or the destination MAC address has not been learned on the access side of 7210-C (unknown MAC address).
- d. If the destination MAC address has been learned by 7210-C, an existing entry in the FIB table identifies the local access port and the IEEE 802.1Q tag (if any) to be added before sending the packet to customer Location-C. The egress Q tag may be different from the ingress Q tag.
- e. If the destination MAC address has not been learned, 7210 floods the packet to all the access SAPs that are participating in the service.

5.2 VPLS features

5.2.1 VPLS enhancements

The Nokia VPLS implementation includes several enhancements beyond basic VPN connectivity. The following VPLS features can be configured individually for each VPLS service instance:

- extensive MAC and IP filter support (up to Layer 4). Filters can be applied on a per-SAP basis.
- Forwarding Information Base (FIB) management features including:
 - configurable FIB size limit
 - FIB size alarms
 - MAC learning disable
 - discard unknown
 - separate aging timers for locally and remotely learned MAC addresses
- ingress rate limiting for broadcast, multicast, and destination unknown flooding on a per-SAP basis
- implementation of Spanning Tree Protocol (STP) parameters on a per-VPLS, per-SAP, and per-spoke-SDP basis
- optional SAP and spoke-SDP redundancy to protect against node failure
- IGMP snooping on a per-SAP and SDP basis

5.2.2 VPLS over MPLS in network operating mode

The VPLS architecture proposed in *draft-ietf-ppvpn-vpls-ldp-0x.txt* specifies the use of provider equipment (PE) that is capable of learning, bridging, and replication on a per-VPLS basis. The PE routers that participate in the service are connected using MPLS Label Switched Path (LSP) tunnels in a full-mesh composed of mesh SDPs or based on an LSP hierarchy (Hierarchical VPLS (H-VPLS)) composed of mesh SDPs and spoke-SDPs. The 7210 SAS supports only H-VPLS.

Multiple VPLS services can be offered over the same set of LSP tunnels. Signaling specified in *RFC 4905* is used to negotiate a set of ingress and egress VC labels on a per-service basis. The VC labels are used by the PE routers for de-multiplexing traffic arriving from different VPLS services over the same set of LSP tunnels.

VPLS/H-VPLS is provided over MPLS by:

- connecting 7210 SAS to bridging-capable PE routers through a mesh/spoke-SDP. The PE routers are connected using a full mesh of LSPs.
- negotiating per-service VC labels using draft-Martini encapsulation
- replicating unknown and broadcast traffic in a service domain
- enabling MAC learning over tunnel and access ports (see [VPLS MAC learning and packet forwarding](#))
- using a separate FIB per VPLS service

5.2.3 VPLS over QinQ spokes for 7210 SAS devices configured in access-uplink operating mode

7210 SAS devices configured in access-uplink operating mode support QinQ spokes or dot1q spokes, which allows them to connect to upstream PE nodes which provides IP/MPLS transport.

VPLS is provided over QinQ/Dot1q spokes by:

- connecting bridging-capable 7210 SAS devices
- replicating unknown and broadcast traffic in a service domain
- enabling MAC learning over QinQ/Dot1q spokes and access ports (see [VPLS MAC learning and packet forwarding](#))
- using a separate FIB per VPLS service

5.2.4 VPLS MAC learning and packet forwarding

The 7210 SAS edge devices perform the packet replication required for broadcast and multicast traffic across the bridged domain. MAC address learning is performed by the 7210 SAS device to reduce the amount of unknown destination MAC address flooding.

Each 7210 SAS maintains a Forwarding Information Base (FIB) for each VPLS service instance, and learned MAC addresses are populated in the FIB table of the service. All traffic is switched based on MAC addresses and forwarded between all participating nodes using the LSP tunnels unknown destination packets (for example, the destination MAC address has not been learned) are forwarded on all LSPs to all participating nodes for that service until the target station responds and the MAC address is learned by the 7210 SAS associated with that service.

5.2.5 IGMP snooping in a VPLS service

**Note:**

- IGMP snooping is supported on all 7210 SAS platforms as described in this document, including those configured in the access-uplink operating mode.
- This section provides information about IGMP snooping support in a VPLS service. It does not apply to R-VPLS services. IGMP snooping can also be enabled for R-VPLS services. See [R-VPLS and IGMPv3 snooping](#) for more information.

In Layer 2 switches, multicast traffic is treated as an unknown MAC address or broadcast frame, which causes the incoming frame to be flooded out (broadcast) on every port within a VLAN. Although this is acceptable behavior for unknown and broadcast frames, this flooded multicast traffic may result in wasted bandwidth on network segments and end stations because IP multicast hosts can join and be interested in only specific multicast groups.

IGMP snooping uses information in Layer 3 protocol headers of multicast control messages to determine the processing at Layer 2. By doing so, an IGMP snooping switch provides the benefit of conserving bandwidth on those segments of the network in which no node has expressed interest in receiving packets addressed to the group address.

**Note:**

References to SDP in the following section about IGMP snooping are applicable only to 7210 SAS platforms operating in network mode.

IGMP snooping can be enabled in the context of VPLS services. The IGMP snooping optimizes the multicast data flow to only those SAPs or SDPs that are members of the group. The system builds a database of group members for each service by listening to IGMP queries and reports from each SAP or SDP, as follows:

- When the switch receives an IGMP report from a host for a particular multicast group, the switch adds the host port number to the forwarding table entry.
- When the switch receives an IGMP leave message from a host, it removes the host port from the table entry, if no other group members are present. It also deletes entries if it does not receive periodic IGMP membership reports from the multicast clients.

The following is a list of supported IGMP snooping features:

- IGMPv1, IGMPv2, and IGMPv3 are supported in accordance with RFC 1112, *Host Extensions for IP Multicasting*, and RFC 2236, *Internet Group Management Protocol, Version 2*:
 - The 7210 SAS-T configured in the access-uplink operating mode supports IGMPv1, IGMPv2, and IGMPv3 snooping in a VPLS service.
 - All 7210 SAS platforms as described in this document, except those configured in the access-uplink operating mode, support only IGMPv1 and IGMPv2 snooping in a VPLS service.
- IGMP snooping can be enabled and disabled on individual VPLS service instances.
- IGMP snooping can be configured on individual SAPs that are part of a VPLS service. When IGMP snooping is enabled on a VPLS service, all its contained SAPs and SDPs automatically have snooping enabled.
- Fast leave terminates the multicast session immediately, instead of using the standard group-specific query to check if other group members are present on the network.

- SAPs and SDPs can be statically configured as multicast router ports. This allows the operator to control the set of ports to which IGMP membership reports are forwarded.
- Static multicast group membership on a per-SAP and a per-SDP basis can be configured.
- The maximum number of multicast groups (static and dynamic) that a SAP or SDP can join can be configured. An event is generated when the limit is reached.
- The maximum number of multicast groups (static and dynamic) that a VPLS instance simultaneously supports can be configured.
- Proxy summarization of IGMP messages reduces the number of IGMP messages processed by upstream devices in the network.
- IGMP filtering allows a subscriber to a service or the provider to block, receive, or transmit permission (or both) to individual hosts or a range of hosts. The following types of filters can be defined:
 - Filter group membership that reports from a particular host or range of hosts. This filtering is performed by importing a defined routing policy into the SAP or SDP.
 - Filters that prevent a host from transmitting multicast streams into the network. The operator can define a data-plane filter (ACL) that drops all multicast traffic and apply this filter to a SAP or SDP.

5.2.5.1 Configuration guidelines for IGMP snooping in VPLS service

The following IGMP snooping considerations apply:

- Layer-2 multicast is supported in VPLS services.
- IGMP snooping is not supported for VCs (**vc-ether** or **vc-vlan**) with **control-word** enabled.
- IGMP snooping fast leave processing can be enabled only on SAPs and SDPs. IGMP snooping proxy summarization is enabled by default on SAPs and SDPs and cannot be disabled. Proxy summarization and fast leave processing are supported only on SDPs for which VCs are configured to use **vc-type ether** and do not have **control-word** enabled.
- IGMP filtering using policies is available on SAPs and SDPs. It is supported only on SDPs for which VCs are configured to use **vc-type ether** and do not have **control-word** enabled.
- Dynamic learning is only supported on SDPs for which VCs are configured to use **vc-type ether** and do not have **control-word** enabled.
- SDPs that are configured to use VCs of type **vc-vlan** that need to be mrouter ports must be configured statically. Multicast group memberships for such SDPs must be configured statically. Dynamic learning is not available for these SDPs.
- IGMP snooping is not supported for **control-word** enabled SDP.
- All 7210 SAS platforms as described in this document, except those configured in the access-uplink operating mode, support only IGMPv1 and IGMPv2 snooping in a VPLS service. These platforms do not support IGMPv3 snooping in a VPLS service.
- All 7210 SAS platforms as described in this document, except those configured in the access-uplink operating mode, support IGMPv3 in an R-VPLS service only. See [R-VPLS and IGMPv3 snooping](#) for more information.

5.2.6 DHCPv4 snooping

To support a DHCP-based address assignment in a Layer 2 aggregation network, the 7210 SAS supports DHCPv4 snooping. The 7210 SAS can copy packets designated to the standard UDP port for DHCP (port 67) to its control plane for inspection, this process is called DHCPv4 snooping.

DHCPv4 snooping can be performed in two directions:

- From the client to the DHCP server (Discover or Request messages) to insert Option 82 information. For these applications, DHCPv4 snooping must be enabled on the SAP toward the subscriber.
- From the DHCP server (ACK messages), to remove the Option 82 field toward the client. For these applications, DHCPv4 snooping must be enabled on both the SAP toward the network and the SAP toward the subscriber.

5.2.7 DHCPv6 snooping

The 7210 SAS supports DHCPv6 snooping, as described in RFC 6221, *Lightweight DHCPv6 Relay Agent*. DHCPv6 snooping allows the user to enable the Lightweight DHCPv6 Relay Agent (LDRA) in a VPLS service.

An LDRA allows relay agent information to be inserted by an 7210 SAS access node that performs a link-layer bridging function. An LDRA resides on the same IPv6 link as the client and a DHCPv6 relay agent or DHCPv6 server, and is similar in function to the DHCPv4 snooping function. DHCPv6 snooping allows relay agent information, including the interface ID option, to be inserted by the access node so that it can be used by the DHCPv6 server for client identification. It also allows for insertion of the remote ID option, allowing the DHCPv6 server to know where the client is attached to the network and if the location is trusted.

DHCPv6 snooping can be performed in two directions:


- From the client to the DHCPv6 server/relay (for example, Discover or Request messages) to insert relay agent information. For these applications, DHCP snooping must be enabled on the SAP (using the **config>service>vpls>sap>dhcp6>snoop** command) toward the subscriber.
- From the DHCPv6 server or relay (for example, ACK messages) to remove the relay agent information toward the client. For these applications, DHCPv6 snooping must be enabled on both the SAP toward the network (using the **config>service>vpls>sap>dhcp6>snoop** command), the SDP binding toward the network (using the **config>service>vpls>spoke-sdp>snoop** or **config>service>vpls>mesh-sdp>snoop** command), and the SAP toward the subscriber (using the **config>service>vpls>sap>dhcp6>snoop** command).



Note: On the 7210 SAS, DHCPv6 snooping on SDP bindings is only supported for spoke SDPs and mesh SDPs in TLDP VPLS services.

In the following figure, the 7210 SAS is used as an access aggregation device to aggregate business CPE for enterprise services, Digital Subscriber Line Access Multiplexers (DSLAMs), and Optical Line Terminations (OLTs) for residential broadband services that use Layer 2 VPLS services.

Figure 63: DHCPv6 snooping for Layer 2 services



Operators need a mechanism to identify the DHCPv6 messages received on their customer-facing ports so that the appropriate IPv6 address and parameters can be provided by the DHCPv6 server. In a Layer 2 network, this is achieved by using DHCPv6 snooping to insert option-18 and option-37 (relay agent information options) in the DHCPv6 messages received from customer-facing ports before being forwarded toward the DHCPv6 server or relay.

5.2.7.1 Client and network-facing service objects

DHCPv6 snooping enables users to identify client and network-facing service objects, in accordance with RFC 6221. Use the following commands to specify if the service object is client-facing, network-facing, or both:

- `config>service>vpls>sap>dhcp6>snoop [client-facing | network-facing | both]`
- `config>service>vpls>mesh-sdp>dhcp6>snoop [network-facing]`
- `config>service>vpls>spoke-sdp>dhcp6>snoop [network-facing]`

To support ring deployments, which employs a ring protection mechanism and the forwarding state of the ring port can change dynamically, a configured service object (for example, a SAP or SDP binding) on a ring port can be client-facing or network-facing depending on the forwarding state of the ring port. To support this scenario, the **both** option is provided. When the **both** option is configured, the processing rules that apply to both client-facing and network-facing service objects are applied.

On the 7210 SAS, SAPs in a VPLS service are configured as **client-facing** (customer-facing ports) by default. In other words, all SAPs with DHCPv6 snooping enabled are client facing. The option to configure the SAP as **network-facing** or **both** is supported; for example, access ports facing the core as a Layer 2 switch with a high SAP scale mode may need to be configured as **network-facing** or **both**.

SDP bindings that are explicitly configured for DHCPv6 snooping are **network-facing** ports by default.

In accordance with RFC 6221, the 7210 SAS does not forward a DHCPv6 relay-forward message out of client-facing service objects. The message is only forwarded out of network-facing service objects or objects configured as **both**.

DHCP relay-reply messages are only processed when received on a network-facing service object. DHCPv6 snooping intercepts and processes all IP traffic received on the network-facing service object that matches:

- a link-local scoped source address
- a link-local scoped destination address
- protocol type UDP
- destination port 547

DHCPv6 snooping inspects the DHCP message type and only forwards the relay-reply message. The other DHCP message types are silently discarded.

The 7210 SAS processes the DHCPv6 client messages received on client-facing interfaces when DHCPv6 snooping is enabled on the SAP or SDP binding and the **client-facing** option is configured.

All IP traffic received on the client-facing service object (for example, a SAP or SDP binding) that meets the following conditions is intercepted and processed:

- destination IP address is set to "All_DHCP_Relay_Agents_and_Servers (ff02::1:2)"
- protocol type UDP
- destination port 547

When snooping is enabled, all DHCPv6 messages that match the preceding criteria are trapped to the CPU. Only the DHCPv6 client messages are processed further; other messages are silently dropped.

5.2.7.1.1 Trusted and untrusted service objects

The 7210 SAS provides the option to configure a service object, which is configured as **client-facing** or **both**, as trusted or untrusted with the **config>service>vpls>sap>dhcp6> [no] trusted** command.

Use the **trusted** command to process and forward a relay-forward message received on a client-facing service object (when DHCPv6 snooping is enabled).

Use the **no trusted** command to disable processing and forwarding of relay-forward message received on a client-facing service object (when DHCPv6 snooping is enabled).

5.2.7.2 DHCPv6 relay agent options

The 7210 SAS provides an option to add option-18 (interface ID option) and option-37 (relay agent remote ID option).

To configure option-18, use the **config>service>vpls>sap>dhcp6>option> [no] interface-id** command . To configure option-37, use the **config>service>vpls>sap>dhcp6>option> [no] remote-id** command.

The **interface-id** option is added to relay-forward messages when DHCPv6 snooping is enabled on client-facing service objects.

The **remote-id** command is optional and only added to relay-forward messages if configured by the user. Configure the **no remote-id** option to disable the addition of this option. IANA Enterprise-numbers .txt specifies the Nokia Enterprise ID value to use with the DHCPv6 **remote-id** option.

5.2.7.3 DHCPv6 snooping QoS considerations

DHCPv6 messages processed by the node are rate-controlled toward the CPU. A Self-Generated Traffic (SGT) QoS configuration is not supported for DHCPv6 messages. However, in most scenarios the incoming packet markings is used to determine the ingress FC and color (if applicable) and the FC is used

for forwarding it out of the appropriate queue on the egress with the appropriate packet header markings (depending on the egress object that packet is forwarded out of - dot1p, IP DSCP, and MPLS EXP field in the packet is updated).

**Note:**

On the 7210 SAS-Mxp, when SAP egress queues are enabled, CPU-generated and forwarded packets are sent out using default port queues.

5.2.8 Multicast VLAN Registration (MVR) support in VPLS service


Multicast VPLS Registration (MVR) is a bandwidth optimization method for multicast in a broadband services network. MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on one or more network-wide multicast VPLS instances.

MVR assumes that subscribers join and leave multicast streams by sending IGMP join and leave messages. The IGMP leave and join message are sent inside the VPLS to which the subscriber port is assigned. The multicast VPLS is shared in the network while the subscribers remain in separate VPLS services. Using MVR, users on different VPLS cannot exchange information between them but still multicast services are provided.

On the MVR VPLS, IGMP snooping must be enabled. On the user VPLS, IGMP snooping and MVR work independently. If IGMP snooping and MVR are both enabled, MVR reacts only to join and leave messages from multicast groups configured under MVR. Join and leave messages from all other multicast groups are managed by IGMP snooping in the local VPLS. This way, potentially several MVR VPLS instances could be configured, each with its own set of multicast channels.

MVR by proxy — In some situations, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behavior) but to another SAP. This is called MVR by proxy, shown in the following figure.

Figure 64: MVR and MVR by proxy



5.2.8.1 Configuration guidelines for MVR in VPLS services

In an MVR configuration, the **svc-sap-type** of the VPLS service that is the source (also known as MVR VPLS service) and the **svc-sap-type** of the VPLS service that is the sink (also known as user VPLS service) should match.

5.2.9 Layer 2 forwarding table management

The following sections describe VPLS features related to management of the FIB.

5.2.9.1 FIB size

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

- **MAC FIB size limits**

Allows users to specify the maximum number of MAC FIB entries that are learned locally for a SAP or remotely for a spoke-SDP. If the configured limit is reached, no new addresses will be learned from the SAP or spoke-SDP until at least one FIB entry is aged out or cleared, as follows:

- When the limit is reached on a SAP or spoke-SDP, packets with unknown source MAC addresses are still forwarded (this default behavior can be changed by configuration). By default, if the

destination MAC address is known, it is forwarded based on the FIB, and if the destination MAC address is unknown, it will be flooded. Alternatively, if discard unknown is enabled at the VPLS service level, unknown destination MAC addresses are discarded.

- The log event SAP MAC limit reached is generated when the limit is reached. When the condition is cleared, the log event SAP MAC Limit Reached Condition Cleared is generated.
- Disable learning at the VPLS service level allows users to disable the dynamic learning function on the service. Disable Learning is supported at the SAP and spoke-SDP level as well.
- Disable aging allows users to turn off aging for learned MAC addresses. It is supported at the VPLS service level, SAP level and spoke-SDP level.

5.2.9.2 FIB size alarms

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

5.2.9.3 Local and remote aging timers

Like a Layer 2 switch, learned MACs within a VPLS instance can be aged out if no packets are sourced from the MAC address for a specified period of time (the aging time). In each VPLS service instance, there are independent aging timers for locally learned MAC and remotely learned MAC entries in the FIB. A local MAC address is a MAC address associated with a SAP because it ingresses on a SAP. A remote MAC address is a MAC address received by an SDP from another router for the VPLS instance. The local-age timer for the VPLS instance specifies the aging time for locally learned MAC addresses, and the remote-age timer specifies the aging time for remotely learned MAC addresses.

In general, the remote-age timer is set to a longer period than the local-age timer to reduce the amount of flooding required for destination unknown MAC addresses. The aging mechanism is considered a low priority process. In most situations, the aging out of MAC addresses can happen within tens of seconds beyond the age time. To minimize overhead, local MAC addresses on a LAG port and remote MAC addresses, in some circumstances, can take up to two times their respective age timer to be aged out.

5.2.9.4 Disable MAC aging

The MAC aging timers can be disabled, which prevents learned MAC entries from being aged out of the FIB. When aging is disabled, it is still possible to manually delete or flush learned MAC entries. Aging can be disabled for learned MAC addresses on a SAP or a spoke-SDP of a VPLS service instance.

5.2.9.5 Disable MAC learning

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

5.2.9.6 Unknown MAC discard

Unknown MAC discard is a feature that discards all packets ingressing the service where the destination MAC address is not in the FIB. The normal behavior is to flood these packets to all end points in the service.

Unknown MAC discard can be used with the disable MAC learning and disable MAC aging options to create a fixed set of MAC addresses allowed to ingress and traverse the service.

5.2.9.7 VPLS and rate limiting

Traffic that is 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 meters can be defined per forwarding class to provide rate-limiting/policing of broadcast traffic, MAC multicast traffic, and unknown destination MAC traffic.

5.2.9.8 MAC move

The MAC move feature is useful to protect against undetected loops in a VPLS topology as well as the presence of duplicate MACs in a VPLS service.

If two clients in the VPLS have the same MAC address, the VPLS will experience a high relearn rate for the MAC. When MAC move is enabled, the 7210 SAS will shut down the SAP or spoke-SDP and create an alarm event when the threshold is exceeded.

MAC move allows sequential order port blocking. By configuration, some VPLS ports can be configured as "non-blockable" which allows simple level of control which ports are being blocked during loop occurrence.

5.2.9.8.1 Split horizon SAP groups and split horizon spoke-SDP groups



Note:

Per-service split horizon groups are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Within the context of VPLS services, a loop-free topology inside a fully meshed VPLS core is achieved by applying a split-horizon forwarding concept. The packets received from a mesh SDP are never forwarded to other mesh SDPs within the same service. The advantage of this approach is that no protocol is required to detect loops within the VPLS core network.

In applications such as DSL aggregation, it is useful to extend this split-horizon concept also to groups of SAPs and spoke-SDPs. This extension is referred to as a split horizon SAP group. Traffic arriving on a SAP or a spoke-SDP within a split horizon group will not be forwarded to other SAPs and spoke-SDPs configured in the same split horizon group, but will be forwarded to other SAPs/spoke-SDPs, which are not part of the split horizon group.

5.2.9.8.2 Configuration guidelines for use of split horizon group in a VPLS service

The following configuration guidelines for the use of split horizon group in a VPLS service apply:

- On 7210 SAS-T, and 7210 SAS-Sx/S 1/10GE operating in network mode, mesh SDPs cannot be configured in a service which uses split horizon group (SHG). Conversely, if a service has a mesh SDP configured, split horizon group cannot be used in the same service.
- On 7210 SAS-Mxp, service based SHG can be configured along with mesh SDPs and spoke-SDPs.

5.2.10 VPLS and spanning tree protocol

The Nokia VPLS service provides a bridged or switched Ethernet Layer 2 network. Equipment connected to SAPs forward Ethernet packets into the VPLS service. The 7210 SAS participating in the service learns where the customer MAC addresses reside, on ingress SAPs.

Unknown destinations, broadcasts, and multicasts are flooded to all other SAPs in the service. If SAPs are connected together, either through misconfiguration or for redundancy purposes, loops can form and flooded packets can keep flowing through the network. The Nokia implementation of the Spanning Tree Protocol (STP) is designed to remove these loops from the VPLS topology. This is done by putting one or several SAPs in the discarding state.

The Nokia implementation of the Spanning Tree Protocol (STP) incorporates some modifications to make the operational characteristics of VPLS more effective.

The STP instance parameters allow the balancing between resiliency and speed of convergence extremes. Modifying particular parameters can affect the behavior. For information about command usage, descriptions, and CLI syntax, refer to [Configuring a VPLS service with CLI](#).

5.2.10.1 Spanning tree operating modes

For each VPLS instance, a preferred STP variant can be configured. The STP variants supported are:

- **rstp**
Rapid Spanning Tree Protocol (RSTP) compliant with IEEE 802.1D-2004 - default mode
- **dot1w**
compliant with IEEE 802.1w
- **comp-dot1w**
operation as in RSTP but backwards compatible with IEEE 802.1w
(this mode allows interoperability with some MTU types)
- **mstp**
compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q-REV/D5.0-09/2005. This mode of operation is only supported in an mVPLS.

While the 7210 SAS initially uses the mode configured for the VPLS, it will dynamically fall back (on a per-SAP basis) to STP (IEEE 802.1D-1998) based on the detection of a BPDU of a different format. A trap or log entry is generated for every change in spanning tree variant.

Some older 802.1W compliant RSTP implementations may have problems with some of the features added in the 802.1D-2004 standard. Interworking with these older systems is improved with the comp-dot1w mode. The differences between the RSTP mode and the comp-dot1w mode are as follows:

- The RSTP mode implements the improved convergence over shared media feature, for example, RSTP will transition from discarding to forwarding in 4 seconds when operating over shared media. The comp-

dot1w mode does not implement this 802.1D-2004 improvement and transitions conform to 802.1w in 30 seconds (both modes implement fast convergence over point-to-point links).

- In the RSTP mode, the transmitted BPDUs contain the port's designated priority vector (DPV) (conforms to 802.1D-2004). Older implementations may be confused by the DPV in a BPDU and may fail to recognize an agreement BPDU correctly. This would result in a slow transition to a forwarding state (30 seconds). For this reason, in the comp-dot1w mode, these BPDUs contain the port's port priority vector (conforms to 802.1w).

The 7210 SAS supports two BPDU encapsulation formats, and can dynamically switch between the following supported formats (on a per-SAP basis):

- IEEE 802.1D STP
- Cisco PVST

5.2.10.2 Multiple Spanning Tree

The Multiple Spanning Tree Protocol (MSTP) extends the concept of the IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) by allowing grouping and associating VLANs to Multiple Spanning Tree Instances (MSTI). Each MSTI can have its own topology, which provides architecture enabling load balancing by providing multiple forwarding paths. At the same time, the number of STP instances running in the network is significantly reduced as compared to Per VLAN STP (PVST) mode of operation. Network fault tolerance is also improved because a failure in one instance (forwarding path) does not affect other instances.

The 7210 SAS implementation of Management VPLS (mVPLS) is used to group different VPLS instances under single RSTP instance. Introducing MSTP into the mVPLS allows the following:

- interoperation with traditional Layer 2 switches in access network
- an effective solution for dual homing of many business Layer 2 VPNs into a provider network

5.2.10.2.1 Redundancy access to VPLS

The GigE MAN portion of the network is implemented with traditional switches. Using MSTP running on individual switches facilitates redundancy in this part of the network. To provide dual homing of all VPLS services accessing from this part of the network, the VPLS PEs must participate in MSTP.

This can be achieved by the following:


- configuring mVPLS on VPLS-PEs (only PEs directly connected to GigE MAN network)
- assign different managed-vlan ranges to different MSTP instances

Typically, the mVPLS would have SAPs with null encapsulations (to receive, send, and transmit MSTP BPDUs) and a mesh SDP to interconnect a pair of VPLS PEs.

Different access scenarios are displayed in the following figure as example network diagrams dually connected to the PBB PEs:

- Access type A - source devices connected by null or Dot1q SAPs
- Access type B - one QinQ switch connected by QinQ/801ad SAPs
- Access type C - two or more ES devices connected by QinQ/802.1ad SAPs

Figure 65: Access resiliency



The following mechanisms are supported for the I-VPLS:

- STP/RSTP - can be used for all access types
- M-VPLS with MSTP - can be used as is just for access Type A. MSTP is required for access type B and C.
- LAG and MC-LAG - can be used for access Type A and B
- Split-horizon-group - does not require residential

5.2.10.3 MSTP for QinQ SAPs

MSTP runs in a MVPLS context and can control SAPs from source VPLS instances. QinQ SAPs are supported. The outer tag is considered by MSTP as part of VLAN range control.

5.2.10.4 Provider MSTP



Note:

Provider MSTP is only supported on platforms that support PBB, and therefore is supported only on 7210 SAS-T operating in the network mode.

Provider MSTP is specified in (IEEE-802.1ad-2005). It uses a provider bridge group address instead of a regular bridge group address used by STP, RSTP, MSTP BPDUs. This allows for implicit separation of source and provider control planes.

The 802.1ad access network sends PBB PE P-MSTP BPDUs using the specified MAC address and also works over QinQ interfaces. P-MSTP mode is used in PBBN for core resiliency and loop avoidance.

Similar to regular MSTP, the STP mode (for example, PMSTP) is only supported in VPLS services where the m-VPLS flag is configured.

5.2.10.4.1 MSTP general principles

MSTP represents modification of RSTP which allows the grouping of different VLANs into multiple MSTIs. To enable different devices to participate in MSTIs, they must be consistently configured. A collection of interconnected devices that have the same MST configuration (region-name, revision and VLAN-to-instance assignment) comprises an MST region.

There is no limit to the number of regions in the network, but every region can support a maximum of 16 MSTIs. Instance 0 is a special instance for a region, known as the Internal Spanning Tree (IST) instance. All other instances are numbered from 1 to 4094. IST is the only spanning-tree instance that sends and receives BPDUs (typically BPDUs are untagged). All other spanning-tree instance information is included in MSTP records (M-records), which are encapsulated within MSTP BPDUs. This means that single BPDU carries information for multiple MSTI which reduces overhead of the protocol.

Any specific MSTI is local to an MSTP region and completely independent from an MSTI in other MST regions. Two redundantly connected MST regions will use only a single path for all traffic flows (no load balancing between MST regions or between MST and SST region).

Traditional Layer 2 switches running MSTP protocol assign all VLANs to the IST instance per default. The operator may then "re-assign" individual VLANs to a specific MSTI by configuring per VLAN assignment. This means that an SR-series PE can be considered as the part of the same MST region only if the VLAN assignment to IST and MSTIs is identical to the one of Layer 2 switches in access network.

5.2.10.4.2 MSTP in the 7210 SAS platform

The 7210 SAS platform uses a concept of mVPLS to group different SAPs under a single STP instance. The VLAN range covering SAPs to be managed by a specific mVPLS is declared under a specific mVPLS SAP definition. MSTP mode-of-operation is only supported in an mVPLS.

When running MSTP, by default, all VLANs are mapped to the CIST. On the VPLS level VLANs can be assigned to specific MSTIs. When running RSTP, the operator must explicitly indicate, per SAP, which VLANs are managed by that SAP.

5.2.10.5 Enhancements to the Spanning Tree Protocol

To interconnect 7210 SAS devices (PE devices) across the backbone, service tunnels (SDPs) are used. These service tunnels are shared among multiple VPLS instances. The Nokia implementation of the Spanning Tree Protocol (STP) incorporates some enhancements to make the operational characteristics of VPLS more effective. The implementation of STP on the router is modified to guarantee that service tunnels will not be blocked in any circumstance without imposing artificial restrictions on the placement of the root bridge within the network. The modifications are fully compliant with the 802.1D-2004 STP specification.

When running MSTP, spoke-SDPs cannot be configured. Also, ensure that all bridges connected by mesh SDPs are in the same region. If not, the mesh will be prevented from becoming active (trap is generated).

To achieve this, all mesh SDPs are dynamically configured as either root ports or designated ports. The PE devices participating in each VPLS mesh determine (using the root path cost learned as part of the normal protocol exchange) which of the 7210 SAS devices is closest to the root of the network. This PE device is internally designated as the primary bridge for the VPLS mesh. As a result of this, all network ports on the primary bridges are assigned the designated port role and therefore remain in the forwarding state.

The second part of the solution ensures that the remaining PE devices participating in the STP instance see the SDP ports as a lower cost path to the root instead of a path that is external to the mesh. Internal to the PE nodes participating in the mesh, the SDPs are treated as zero cost paths toward the primary bridge. As a consequence, the path through the mesh are seen as lower cost than any alternative and the PE node will designate the network port as the root port. This ensures that network ports always remain in forwarding state.

A combination of the previously mentioned features ensure that network ports are never blocked and maintain interoperability with bridges external to the mesh that are running STP instances.

5.2.10.5.1 L2PT termination

L2PT is used to transparently transport protocol data units (PDUs) of Layer 2 protocols, such as Cisco Discovery Protocol (CDP), Dynamic Trunking Protocol (DTP), Port Aggregation Protocol (PAGP), Spanning Tree Protocol (STP), Unidirectional Link Detection (UDLD), VLAN trunking protocol (VTP), and Link Layer Discovery Protocol (LLDP). This allows users to run these protocols between customer CPEs without involving backbone infrastructure.

The 7210 SAS routers support the transparent tunneling of PDUs across the VPLS core; however, in some network designs the VPLS PE is connected to CPEs through a legacy Layer 2 network, instead of via direct connections. In this type of environment, the termination of tunnels through the infrastructure is required.

L2PT tunnels transport protocol PDUs by overwriting MAC destination addresses at the ingress of the tunnel to a proprietary MAC address, such as 01-00-0c-cd-cd-d0. On egress of the tunnel, the MAC address is overwritten back to the MAC address of the respective Layer 2 protocol.

The 7210 SAS nodes support L2PT termination for STP BPDUs as follows:

- On ingress of every SAP or spoke-SDP that is configured as an L2PT termination, all PDUs with a MAC destination address of 01-00-0c-cd-cd-d0 are intercepted, and their MAC destination address is overwritten to the MAC destination address used for the corresponding protocol. The type of protocol can be derived from LLC and SNAP encapsulation.
- In the egress direction, PDUs of the corresponding protocol received on all VPLS ports are intercepted, and L2PT encapsulation is performed for SAP and spoke-SDPs configured as L2PT termination points. For implementation reasons, PDU interception and redirection to CPM can be performed only on ingress. Therefore, to comply with the preceding requirement, as soon as at least one port of a specific VPLS service is configured as an L2PT termination port, redirection of PDUs to CPM are set on all other ports (SAPs, spoke-SDPs) of the VPLS service.

L2PT termination can be enabled only if STP is disabled in the context of the specific VPLS service.

5.2.10.5.2 BPDU translation

VPLS networks are typically used to interconnect different customer sites using different access technologies such as Ethernet and bridged-encapsulated ATM PVCs. Typically, different Layer 2 devices can support different types of STP and even if they are from the same vendor. In some cases, it is necessary to provide BPDU translation to provide an inter-operable e2e solution.

To address these network designs, BPDU format translation is supported on 7210 SAS devices. If enabled on a specific SAP or spoke-SDP, the system will intercept all BPDUs destined for that interface and perform required format translation such as STP-to-PVST or the other way around.

Similarly, BPDU interception and redirection to the CPM is performed only at ingress meaning that as soon as at least 1 port within a specific VPLS service has BPDU translation enabled, all BPDUs received on any of the VPLS ports will be redirected to the CPM.

BPDU translation involves all encapsulation actions that the datapath would perform for a specific outgoing port (such as adding VLAN tags depending on the outer SAP and the SDP encapsulation type) and adding or removing all the required VLAN information in a BPDU payload.

This feature can be enabled on a SAP/spoke only if STP is disabled in the context of the specific VPLS service.

5.2.10.5.3 L2PT and BPDU translation

L2PT termination for only STP (Spanning Tree Protocol) and PVST (Per VLAN Spanning Tree Protocol), Cisco Discovery Protocol (CDP), Digital Trunking Protocol (DTP), Port Aggregation Protocol (PAgP), Unidirectional Link Detection (UDLD), Virtual Trunk Protocol (VTP), STP (Spanning Tree Protocol) and PVST (per-VLAN Spanning Tree protocol) are supported on 7210 SAS devices.

These protocols automatically pass the other protocols tunneled by L2PT toward the CPM and all carry the same specific Cisco MAC.

The existing L2PT limitations apply:

- The protocols apply only to VPLS.
- The protocols are mutually exclusive with running STP on the same VPLS as soon as one SAP/spoke has L2PT/BPDU translation enabled.
- Forwarding occurs on the CPM and uses CPU processing cycles.

5.2.11 VPLS redundancy

The VPLS standard (RFC 4762, *Virtual Private LAN Services Using LDP Signalling*) includes provisions for hierarchical VPLS, using point-to-point spoke-SDPs. Two applications have been identified for spoke-SDPs:

- to connect to Multi-Tenant Units (MTUs) to PEs in a metro area network
- to interconnect the VPLS nodes of two networks

In both applications the spoke-SDPs serve to improve the scalability of VPLS. While node redundancy is implicit in non-hierarchical VPLS services (using a full mesh of SDPs between PEs), node redundancy for spoke-SDPs needs to be provided separately. In VPLS services, only two spoke-SDPs are allowed in an endpoint.

The 7210 SAS routers have implemented special features for improving the resilience of hierarchical VPLS instances, in both MTU and inter-metro applications.

5.2.11.1 Spoke-SDP redundancy for metro interconnection


When two or more meshed VPLS instances are interconnected by redundant spoke-SDPs (as shown in [Figure 66: H-VPLS with spoke redundancy](#)), a loop in the topology results. To remove such a loop from the topology, Spanning Tree Protocol (STP) can be run over the SDPs (links) which form the loop such that one of the SDPs is blocked. As running STP in each and every VPLS in this topology is not efficient, the node includes functionality which can associate a number of VPLSes to a single STP instance running over

the redundant SDPs. Node redundancy is therefore achieved by running STP in one VPLS, and applying the conclusions of this STP to the other VPLS services. The VPLS instance running STP is referred to as the "management VPLS" or mVPLS.

In the case of a failure of the active node, STP on the management VPLS in the standby node will change the link states from disabled to active. The standby node will then broadcast a MAC flush LDP control message in each of the protected VPLS instances, so that the address of the newly active node can be relearned by all PEs in the VPLS.

It is possible to configure two management VPLS services, where both VPLS services have different active spokes (this is achieved by changing the path-cost in STP). By associating different user VPLSes with the two management VPLS services, load balancing across the spokes can be achieved.

Figure 66: H-VPLS with spoke redundancy



OSSG045

5.2.11.2 Spoke-SDP-based redundant access

This feature provides the ability to have a node deployed as MTUs (Multi-Tenant Unit Switches) to be multi-homed for VPLS to multiple routers deployed as PEs without requiring the use of mVPLS.

In the configuration example displayed in [Figure 66: H-VPLS with spoke 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. If the primary and secondary spoke-SDPs have the same precedence value, the spoke-SDP with lower ID functions as the primary SDP.

The secondary spoke is in a blocking state (both on receive and transmit) as long as the primary spoke is available. When the primary spoke becomes unavailable (because of link failure, PEs failure, and so on), the MTU immediately switches traffic to the backup spoke and starts receiving/sending traffic to/from the standby spoke. Optional revertive operation (with configurable switch-back delay) is applicable only when one of the spokes is configured with precedence of primary. If not, this action does not take place. Forced manual switchover is also supported.

To speed up the convergence time during a switchover, MAC flush is configured. The MTUs generates a MAC flush message over the newly unblocked spoke when a spoke change occurs. As a result, the PEs receiving the MAC flush will flush all MACs associated with the impacted VPLS service instance and forward the MAC flush to the other PEs in the VPLS network if "propagate-mac-flush" is enabled.

5.2.11.3 Inter-domain VPLS resiliency using multi-chassis endpoints



Note:


MC-EP is not supported on 7210 SAS platforms. This section provides an example of how 7210 SAS platforms can be used as MTU devices in an MC-EP solution. In this solution, 7750 SR routers provide the MC-EP functionality.

Inter-domain VPLS refers to a VPLS deployment where sites may be located in different domains. An example of inter-domain deployment can be where different Metro domains are interconnected over a Wide Area Network (Metro1-WAN-Metro2) or where sites are located in different autonomous systems (AS1-ASBRs-AS2).

Multi-chassis endpoint (MC-EP) provides an alternate solution that does not require RSTP at the gateway VPLS PEs while still using pseudowires to interconnect the VPLS instances located in the two domains.

MC-EP expands the single chassis endpoint based on active/standby pseudowires for VPLS shown in the following figure. In the solution depicted by the following figure, 7210 SAS devices are used as MTUs.

Figure 67: H-VPLS resiliency based on AS pseudowires



The active/standby pseudowire solution is appropriate for the scenario when only one VPLS PE (MTU-s) needs to be dual-homed to two core PEs (PE1 and PE2).

5.2.12 VPLS access redundancy


A second application of hierarchical VPLS is using MTUs that are MPLS-enabled which must have spoke-SDPs to the closest PE node. To protect against failure of the PE node, an MTU can be dual-homed.

The following are several mechanisms that can be used to resolve a loop in an access network where 7210 SAS devices are used:

- STP-based access, with or without mVPLS
- Ethernet APS using G.8032

5.2.12.1 STP-based redundant access to VPLS

Figure 68: Dual-homed MTUs in two-tier hierarchy H-VPLS



In configuration shown in the preceding figure, STP is activated on the MTU and two PEs to resolve a potential loop.

To remove such a loop from the topology, Spanning Tree Protocol (STP) can be run over the SDPs (links) which form the loop such that one of the SDPs is blocked. Running STP in every VPLS in this topology is not efficient as the node includes functionality which can associate a number of VPLSes to a single STP instance running over the redundant SDPs. Node redundancy is therefore achieved by running STP in one VPLS. Therefore, this applies the conclusions of this STP to the other VPLS services.

The VPLS instance running STP is referred to as the "management VPLS" or mVPLS. In the case of a failure of the active node, STP on the management VPLS in the standby node will change the link states from disabled to active. The standby node will then broadcast a MAC flush LDP control message in each of the protected VPLS instances, so that the address of the newly active node can be relearned by all PEs in the VPLS. It is possible to configure two management VPLS services, where both VPLS services have different active spokes (this is achieved by changing the path-cost in STP). By associating different user VPLSes with the two management VPLS services, load balancing across the spokes can be achieved.

In this configuration the scope of STP domain is limited to MTU and PEs, while any topology change needs to be propagated in the whole VPLS domain.

This is done by using "MAC-flush" messages defined by RFC 4762, *Virtual Private LAN Services Using LDP Signaling*. In the case where STP acts as a loop resolution mechanism, every Topology Change Notification (TCN) received in a context of STP instance is translated into an LDP-MAC address withdrawal message (also referred to as a MAC-flush message) requesting to clear all FDB entries except the ones learned from the originating PE. Such messages are sent to all PE peers connected through SDPs (mesh and spoke) in the context of VPLS services which are managed by the specific STP instance.

5.2.12.2 Redundant access to VPLS without STP

The Nokia implementation also provides alternative methods for providing redundant access to Layer 2 services, such as MC-LAG. Also in this case, the topology change event needs to be propagated into VPLS topology to provide fast convergence.

Figure 66: H-VPLS with spoke redundancy shows a dual-homed connection to VPLS service (PE-A, PE-B, PE-C, PE-D) and operation in case of link failure (between PE-C and Layer 2-B). Upon detection of

a link failure PE-C will send MAC-Address-Withdraw messages, which will indicate to all LDP peers that they should flush all MAC addresses learned from PE-C. This will lead that to a broadcasting of packets addressing affected hosts and relearning process in case an alternative route exists.

Note that the message described here is different from the message described in previous section and in RFC 4762, *Virtual Private LAN Services Using LDP Signaling*. The difference is in the interpretation and action performed in the receiving PE. According to the standard definition, upon receipt of a MAC withdraw message, all MAC addresses, except the ones learned from the source PE, are flushed,

This section specifies that all MAC addresses learned from the source are flushed. This message has been implemented as an LDP address message with vendor-specific type, length, value (TLV), and is called the flush-mine message.

The advantage of this approach (as compared to RSTP based methods) is that only MAC-affected addresses are flushed and not the full forwarding database. While this method does not provide a mechanism to secure alternative loop-free topology, the convergence time is dependent on the speed of the specific CE device will open alternative link (Layer 2-B switch in Figure 57) as well as on the speed PE routers will flush their FDB.

In addition, this mechanism is effective only if PE and CE are directly connected (no hub or bridge) as it reacts to physical failure of the link.

5.2.13 MAC flush message processing

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

As described on respective sections, the 7210 SAS supports two types of MAC flush message, flush-all-but-mine and flush-mine. The main difference between these messages is the type of action they signal. Flush-all-but-mine requests clearing of all FDB entries which were learned from all other LDP peers except the originating PE. This type is also defined by RFC 4762 as an LDP MAC address withdrawal with an empty MAC address list.

Flush-all-mine message requests clearing all FDB entries learned from originating PE. This means that this message has exactly other effect than flush-all-but-mine message. This type is not included in RFC 4762 definition and it is implemented using vendor specific TLV.

The advantages and disadvantages of the individual types should be apparent from examples in the previous section. The description here focuses on summarizing actions taken on reception and conditions individual messages are generated.

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

- clears FDB entries of all indicated VPLS services conforming the definition
- propagates the message (preserving the type) to all LDP peers, if "propagate-mac-flush" flag is enabled at corresponding VPLS level

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

- The flush-all-but-mine message is received from LDP peer and propagate-mac-flush flag is enabled. The message is sent to all LDP peers in the context of VPLS service it was received in.
- TCN message in a context of STP instance is received. The flush-all-but-mine message is sent to all LDP-peers connected with spoke and mesh SDPs in a context of VPLS service controlled by the

specific STP instance (based on mVPLS definition). The message is sent only to LDP peers which are not part of STP domain, which means corresponding spoke and mesh SDPs are not part of mVPLS.

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

The flush-mine message is generated under following conditions:


- The flush-mine message is received from LDP peer and "propagate-mac-flush" flag is enabled. The message is sent to all LDP peers in the context of VPLS service it was received.
- The flush-mine message is generated when on a SAP or SDP transition from operationally up to an operationally down state and send-flush-on-failure flag is enabled in the context of the specific VPLS service. The message is sent to all LDP peers connected in the context of the specific VPLS service. When enabling "send-flush-on-failure" the flag is blocked in VPLS service managed by mVPLS. This is to prevent both messages being sent at the same time.
- The flush-mine message is generated on an MC-LAG SAP transition from an operationally up state to an operationally down state. The message is sent to all LDP peers connected in the context of the specific VPLS service.

5.2.13.1 MAC Flush with STP

A second application of Hierarchical VPLS is in the use of Multi Tenant Units (MTU). MTUs are typically not MPLS-enabled, and therefore have Ethernet links to the closest PE node (see the following figure). To protect against failure of the PE node, an MTU could be dual-homed and therefore have two SAPs on two PE nodes. To resolve the potential loop, STP is activated on the MTU and the two PEs.

Like in the previous scenario, STP only needs to run in a single VPLS instance, and the results of the STP calculations are applied to all VPLSes on the link. Equally, the standby node will broadcast MAC flush LDP messages in the protected VPLS instances when it detects that the active node has failed.

Figure 69: H-VPLS with SAP redundancy



OSSG046

5.2.13.2 Selective MAC flush


When using STP as described previously is not appropriate, the "Selective MAC flush" feature can be used instead.

In this scenario, the 7210 SAS that detects a port failure will send out a flush-all-from-ME LDP message to all PEs in the VPLS. The PEs receiving this LDP message will remove all MAC entries originated by the sender from the indicated VPLS.

A drawback of this approach is that the selective MAC flush does not signal that a backup path was found, only that the previous path is no longer available. In addition, the selective MAC Flush mechanism is effective only if the CE and PE are directly connected (no intermediate hubs or bridges) as it reacts only to a physical failure of the link. Consequently, Nokia recommends using the MAC flush with STP method previously described where possible.

5.2.13.3 Dual homing to a VPLS service

Figure 70: Dual homed CE connection to VPLS



The preceding figure shows a dual-homed connection to VPLS service (PE-A, PE-B, PE-C, PE-D) and operation in case of link failure (between PE-C and Layer 2-B). Upon detection of a link failure PE-C will send MAC-Address-Withdraw messages, which will indicate to all LDP peers that they should flush all MAC addresses learned from PE-C. This will lead that to a broadcasting of packets addressing affected hosts and relearning process in case an alternative route exists.

Note that the message described here is different from the message described in draft-ietf-l2vpn-vpls-ldp-xx.txt, *Virtual Private LAN Services over MPLS*. The difference is in the interpretation and action performed

in the receiving PE. According the draft definition, upon receipt of a MAC-withdraw message, all MAC addresses, except the ones learned from the source PE, are flushed. This section specifies that all MAC addresses learned from the source are flushed. This message has been implemented as an LDP address message with vendor-specific type, length, value (TLV), and is called the flush-all-from-ME message.

The draft definition message is currently used in management VPLS which is using RSTP for recovering from failures in Layer 2 topologies. The mechanism described in this document represent an alternative solution.

The advantage of this approach (as compared to RSTP based methods) is that only MAC-affected addresses are flushed and not the full forwarding database. While this method does not provide a mechanism to secure alternative loop-free topology, the convergence time is dependent on the speed of the specific CE device will open alternative link (Layer 2-B switch in the preceding figure) as well as on the speed PE routers will flush their FDB.

In addition, this mechanism is effective only if PE and CE are directly connected (no hub or bridge) as it reacts to physical failure of the link.

5.2.14 VPLS service considerations

This section describes various 7210 SAS service features and special capabilities or considerations as they relate to VPLS services.

5.2.14.1 SAP encapsulations

VPLS services are designed to carry Ethernet frame payloads, so it can provide connectivity between any SAPs that pass Ethernet frames. The following SAP encapsulations are supported on the VPLS service:

- Ethernet null
- Ethernet Dot1q
- Ethernet Dot1q Default
- Ethernet Dot1q Explicit Null

5.2.14.2 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:

1. null encapsulation defined on ingress

VLAN tags are ignored and the packet goes to a default service for the SAP.

2. dot1q encapsulation defined on ingress

Only first VLAN tag is considered.

3. dot1q Default encapsulation defined on ingress

Accepted tagged packets not matching any of the configured VLAN encapsulations. This is like a default SAP for tagged packets.

4. dot1q Explicit Null encapsulation defined on ingress

Any untagged or priority-tagged packets will be accepted.

5.3 BGP Auto-Discovery for LDP VPLS

BGP Auto Discovery (BGP AD) for LDP VPLS is a framework for automatically discovering the endpoints of a Layer 2 VPN offering an operational model similar to that of an IP VPN. This model allows carriers to leverage existing network elements and functions, including but not limited to, route reflectors and BGP policies to control the VPLS topology.

BGP AD is an excellent complement to an already established and well deployed Layer 2 VPN signaling mechanism target LDP providing one touch provisioning for LDP VPLS where all the related PEs are discovered automatically. The service provider may make use of existing BGP policies to regulate the exchanges between PEs in the same, or in different, autonomous system (AS) domains. The addition of BGP AD procedures does not require carriers to uproot their existing VPLS deployments and to change the signaling protocol.

5.3.1 BGP AD overview

The BGP protocol establishes neighbor relationships between configured peers. An open message is sent after the completion of the three-way TCP handshake. This open message contains information about the BGP peer sending the message. This message contains Autonomous System Number (ASN), BGP version, timer information and operational parameters, including capabilities. The capabilities of a peer are exchanged using two numerical values: the Address Family Identifier (AFI) and Subsequent Address Family Identifier (SAFI). These numbers are allocated by the Internet Assigned Numbers Authority (IANA). BGP AD uses AFI 65 (L2VPN) and SAFI 25 (BGP VPLS).

5.3.2 Information model

Following is the establishment of the peer relationship, the discovery process begins as soon as a new VPLS service instance is provisioned on the PE.

Two VPLS identifiers are used to indicate the VPLS membership and the individual VPLS instance:

- **VPLS-ID**

Membership information, unique network wide identifier; same value assigned for all VPLS switch instances (VSIs) belonging to the same VPLS; encodable and carried as a BGP extended community in one of the following formats:

- a two-octet AS specific extended community
- an IPv4 address specific extended community


- **VSI-ID**

The unique identifier for each individual VSI, built by concatenating a route distinguisher (RD) with a 4 bytes identifier (usually the system IP of the VPLS PE); encoded and carried in the corresponding BGP NLRI.

To advertise this information, BGP AD employs a simplified version of the BGP VPLS NLRI where just the RD and the next 4 bytes are used to identify the VPLS instance. There is no need for Label Block and Label Size fields as T-LDP will take care of signaling the service labels later on.

The format of the BGP AD NLRI is very similar with the one used for IP VPN as shown in the following figure. The system IP may be used for the last 4 bytes of the VSI-ID further simplifying the addressing and the provisioning process.

Figure 71: BGP AD NLRI versus IP VPN NLRI



Network Layer Reachability Information (NLRI) is exchanged between BGP peers indicating how to reach prefixes. The NLRI is used in the Layer 2 VPN case to tell PE peers how to reach the VSI instead of specific prefixes. The advertisement includes the BGP next hop and a route target (RT). The BGP next hop indicates the VSI location and is used in the next step to determine which signaling session is used for pseudowire signaling. The RT, also coded as an extended community, can be used to build a VPLS full mesh or a H-VPLS hierarchy through the use of BGP import or export policies.

BGP is only used to discover VPN endpoints and the corresponding far end PEs. It is not used to signal the pseudowire labels. This task remains the responsibility of targeted-LDP (T-LDP).


5.3.3 FEC element for T-LDP signaling

Two LDP FEC elements are defined in RFC 4447, *PW Setup & Maintenance Using LDP*. The original pseudowire-ID FEC element 128 (0x80) employs a 32-bit field to identify the virtual circuit ID and it was used extensively in the initial VPWS and VPLS deployments. The simple format is easy to understand but it does not provide the required information model for BGP Auto-Discovery function. To support BGP AD and other new applications a new Layer 2 FEC element, the generalized FEC (0x81) is required.

The generalized pseudowire-ID FEC element has been designed for auto discovery applications. It provides a field, the address group identifier (AGI), that is used to signal the membership information from the VPLS-ID. Separate address fields are provided for the source and target address associated with the VPLS endpoints called the Source Attachment Individual Identifier (SAII) and respectively, Target Attachment Individual Identifier (TAII). These fields carry the VSI-ID values for the two instances that are to be connected through the signaled pseudowire.

The detailed format for FEC 129 is shown in the following figure.

Figure 72: Generalized pseudowire-ID FEC element



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


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

5.3.4 BGP-AD and Target LDP (T-LDP) interaction

BGP is responsible for discovering the location of VSIs that share the same VPLS membership. LDP protocol is responsible for setting up the pseudowire infrastructure between the related VSIs by exchanging service-specific labels between them.

When the local VPLS information is provisioned in the local PE, the related PEs participating in the same VPLS are identified through BGP AD exchanges. A list of far-end PEs is generated and triggers the creation, if required, of the necessary T-LDP sessions to these PEs and the exchange of the service-specific VPN labels. The steps for the BGP AD discovery process and LDP session establishment and label exchange are shown in the following figure.

Figure 73: BGP-AD and T-LDP interaction



sw0708

Key:

1. Establish I-BGP connectivity RR.
2. Configure VPN (10) on edge node (PE3).
3. Announce VPN to RR using BGP-AD.
4. Send membership update to each client of the cluster.
5. LDP exchange or inbound FEC filtering (IFF) of non-match or VPLS down.
6. Configure VPN (10) on edge node (PE2).
7. Announce VPN to RR using BGP-AD.
8. Send membership update to each client of the cluster.
9. LDP exchange or inbound FEC filtering (IFF) of non-match or VPLS down.
10. Complete LDP bidirectional pseudowire establishment FEC 129.

5.3.5 SDP usage

Service Access Points (SAP) are linked to transport tunnels using Service Distribution Points (SDP). The service architecture of the 7210 platform allows services to be abstracted from the transport network.

MPLS transport tunnels are signaled using the Resource Reservation Protocol (RSVP-TE) or by the Label Distribution Protocol (LDP). The capability to automatically create an SDP only exists for LDP based transport tunnels. Using a manually provisioned SDP is available for both RSVP-TE and LDP transport tunnels. See the *7210 SAS-Mxp, R6, R12, S, Sx, T MPLS Guide* for more information about MPLS, LDP, and RSVP.

5.3.6 Automatic creation of SDPs

When BGP AD is used for LDP VPLS and LDP is used as the transport tunnel there is no requirement to manually create an SDP. The LDP SDP can be automatically instantiated using the information advertised by BGP AD. This simplifies the configuration on the service node.

Enabling LDP on the IP interfaces connecting all nodes between the ingress and the egress, builds transport tunnels based on the best IGP path. LDP bindings are automatically built and stored in the hardware. These entries contain an MPLS label pointing to the best next hop along the best path toward the destination.

When two endpoints need to connect and no SDP exists, a new SDP will automatically be constructed. New services added between two endpoints that already have an automatically created SDP will be immediately used. No new SDP will be constructed. The far-end information is gleaned from the BGP next hop information in the NLRI. When services are withdrawn with a BGP_Unreach_NLRI, the automatically established SDP will remain up as long as at least one service is connected between those endpoints. An automatically created SDP will be removed and the resources released when the only or last service is removed.

5.3.7 Manually provisioned SDP

The carrier is required to manually provision the SDP if they create transport tunnels using RSVP-TE. Operators have the option to choose a manually configured SDP, if they use LDP as the tunnel signaling protocol. The functionality is the same regardless of the signaling protocol.

Creating a BGP-AD enabled VPLS service on an ingress node with the manually provisioned SDP option causes the Tunnel Manager to search for an existing SDP that connects to the far-end PE. The far-end IP information is gleaned from the BGP next hop information in the NLRI. If a single SDP exists to that PE, it is used. If no SDP is established between the two endpoints, the service remains down until a manually configured SDP becomes active.

When multiple SDPs exist between two endpoints, the tunnel manager selects the appropriate SDP. The algorithm preferred SDPs with the best (lower) metric. Should there be multiple SDPs with equal metrics, the operational state of the SDPs with the best metric is considered. If the operational state is the same, the SDP with the higher sdp-id is used. If an SDP with a preferred metric is found with an operational state that is not active, the tunnel manager flags it as ineligible and restarts the algorithm.

5.3.8 Automatic instantiation of pseudowires (SDP bindings)

The choice of manual or auto provisioned SDPs has limited impact on the amount of required provisioning. Most of the savings are achieved through the automatic instantiation of the pseudowire infrastructure (SDP bindings). This is achieved for every auto-discovered VSLs through the use of the pseudowire template concept. Each VPLS service that uses BGP AD contains the "pw-template-binding" option defining specific Layer 2 VPN parameters. This command references a "pw-template" which defines the pseudowire parameters. The same "pwtemplate" may be referenced by multiple VPLS services. As a result, changes to these pseudowire templates have to be treated with great care as they may impact many customers at once.

The Nokia implementation provides for safe handling of pseudowire templates. Changes to the pseudowire templates are not automatically propagated. Tools are provided to evaluate and distribute the changes. The following command is used to distribute changes to a "pw-template" at the service level to one or all services that use that template.

```
PERs-4# tools perform service id 300 eval-pw-template 1 allow-service-impact
```

If the service ID is omitted, then all services are updated. The type of change made to the "pwtemplate" influences how the service is impacted in the following ways:

1. Adding or removing a split-horizon-group will cause the router to destroy the original object and recreate using the new value.
2. Changing parameters in the **vc-type {ether | vlan}** command requires LDP to re-signal the labels.

Both of these changes affect the services. Other changes are not service affected.

5.3.9 Mixing statically configured and auto-discovered pseudowires in a VPLS service

The services implementation allows for manually provisioned and auto-discovered pseudowire (SDP bindings) to coexist in the same VPLS instance (for example, both FEC128 and FEC 129 are supported). This allows for gradual introduction of auto discovery into an existing VPLS deployment.

As FEC 128 and 129 represent different addressing schemes, it is important to make sure that only one is used at any point in time between the same two VPLS instances. Otherwise, both pseudowires may become active causing a loop that might adversely impact the correct functioning of the service. It is recommended that FEC128 pseudowire be disabled as soon as the FEC129 addressing scheme is introduced in a portion of the network. Alternatively, RSTP may be used during the migration as a safety mechanism to provide additional protection against operational errors.

5.3.10 Resiliency schemes

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

BGP-AD may coexist with Hierarchical-VPLS (H-VPLS) resiliency schemes (for example, dual homed MTU-s devices to different PE-rs nodes) using existing methods (M-VPLS and statically configured Active or Standby pseudowire endpoint).

If provisioned SDPs are used by BGP AD, M-VPLS may be employed to provide loop avoidance. However, it is currently not possible to auto-discover active or standby pseudowires and to instantiate the related endpoint.

5.4 Routed VPLS



Note:

R-VPLS with IPv6 interfaces is supported only on 7210 SAS-Mxp operating in the network mode.

Routed VPLS (R-VPLS) allows a VPLS instance to be associated with an IP interface.



Note:

- R-VPLS is supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.
- For 7210 SAS platforms operating in network mode, R-VPLS can provide both customer service and in-band management of the node.
- For 7210 SAS platforms operating in access-uplink mode, R-VPLS can only provide in-band management of the node.

Within an R-VPLS service, traffic with a destination MAC matching that of the associated IP interface is routed based on the IP forwarding table; all other traffic is forwarded based on the VPLS forwarding table.

In access-uplink mode, an R-VPLS service can be associated with an IPv4 interface and supports only static routing. It is primarily designed for use of in-band management of the node. It allows for in-band management of the 7210 SAS nodes in a ring deployment using a single IPv4 subnet, reducing the number of IP subnets needed.

In network mode, an R-VPLS service can be associated with an IPv4 or IPv6 interface and supports static routing and other routing protocols. It can be used to provide a service to the customer or for in-band management of the node.

5.4.1 IES or VPRN IP interface binding

A standard IP interface within an existing IES or VPRN service context may be bound to a service name. A VPLS service only supports binding for a single IP interface.

While an IP interface may only be bound to a single VPLS service, the routing context containing the IP interface (IES or VPRN) may have other IP interfaces bound to other VPLS service contexts. That is, R-VPLS allows the binding of IP interfaces in IES and VPRN services to be bound to VPLS services.

5.4.2 Assigning a service name to a VPLS service

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

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

VPLS service associated with the name. Only one interface can be bound to the service with the specified name.

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

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

5.4.3 Service binding requirements

If a defined service name is created on the system, the system checks to ensure that the service type is VPLS. If the created service type is VPLS, the IP interface is eligible to enter the operationally upstate.

5.4.3.1 Bound Service Name Assignment

If a bound service name is assigned to a service within the system, the system first checks to ensure that the service type is VPLS. Secondly, the system ensures that the service is not already bound to another IP interface through the service name. If the service type is not VPLS or the service is already bound to another IP interface through the service ID, the service name assignment fails.

A single VPLS instance cannot be bound to two separate IP interfaces.

5.4.3.2 Binding a service name to an IP interface

An IP interface within an IES or VPRN service context may be bound to a service name at anytime. Only one interface can be bound to a service. When an IP interface is bound to a service name and the IP interface is administratively up, the system scans for a VPLS service context using the name and takes the following actions:

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

5.4.3.3 IP interface attached VPLS service constraints

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

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

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

5.4.3.4 IP interface and VPLS operational state coordination

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

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

5.4.3.5 IP interface MTU and fragmentation



Note:

VPLS service MTU is supported on all 7210 SAS platforms as described in this document, except the 7210 SAS-T operating in access-uplink mode.

The user must ensure that the port MTU is configured appropriately so that the largest packet traversing through any of the SAPs (virtual ports) of the VPLS service can be forwarded out of any of the SAPs. VPLS services do not support fragmentation and can discard packets larger than the configured port MTU.

When an IP interface is associated with a VPLS service, the IP-MTU is based on either the administrative value configured for the IP interface or an operational value derived from port MTU of all the SAPs configured in the service. The port MTU excluding the Layer 2 Header and tags for all the ports which have SAPs configured in this VPLS service are considered and the minimum value among those are computed (which is called computed MTU). The operational value of the IP interface is set as follows:

- If the configured (administrative) value of IP MTU is greater than the computed MTU, then the operational IP MTU is set to the computed MTU.
- If the configured (administrative) value of IP MTU is lesser than or equal to the computed MTU, then operational IP MTU is set to the configured (administrative) value of IP MTU.

5.4.4 ARP and VPLS FIB interactions

Two address-oriented table entries are used when routing into a VPLS service. On the routing side, an ARP entry is used to determine the destination MAC address used by an IP next-hop. In the case where the destination IP address in the routed packet is a host on the local subnet represented by the VPLS instance, the destination IP address is used as the next-hop IP address in the ARP cache lookup. If the destination IP address is in a remote subnet that is reached by another router attached to the VPLS service, the routing lookup returns the local IP address on the VPLS service of the remote router is returned. If the next-hop is not currently in the ARP cache, the system generates an ARP request to determine the destination MAC address associated with the next-hop IP address. IP routing to all destination hosts associated with the next-hop IP address stops until the ARP cache is populated with an entry for the next-hop. The dynamically populated ARP entries age out according to the ARP aging timer.

The second address table entry that affects VPLS routed packets is the MAC destination lookup in the VPLS service context. The MAC associated with the ARP table entry for the IP next-hop may or may not currently be populated in the VPLS Layer 2 FIB table. While the destination MAC is unknown (not

populated in the VPLS FIB), the system is flooded with all packets destined for that MAC (routed or bridged) to all SAPs within the VPLS service context. When the MAC is known (populated in the VPLS FIB), all packets destined to the MAC (routed or bridged) is targeted to the specific SAP where the MAC has been learned. As with ARP entries, static MAC entries may be created in the VPLS FIB. Dynamically learned MAC addresses are allowed to age out or be flushed from the VPLS FIB while static MAC entries always remain associated with a specific virtual port. Dynamic MACs may also be relearned on another VPLS SAP than the current SAP in the FIB. In this case, the system automatically moves the MAC FIB entry to the new VPLS SAP.



Note:

- In 7210 SAS, whenever a MAC entry is removed from the VPLS FIB (either explicitly by the user or because of MAC aging or mac-move), ARP entries which match this MAC address is removed from the ARP cache. Though the VPLS FIB entries are not removed; an ARP entry ages out and is removed from the ARP cache.
- If the VPLS FIB limit is reached and the node is no longer able to learn new MAC address, ARP will also not be learned.

5.4.5 R-VPLS specific ARP cache behavior

In typical routing behavior, the system uses the IP route table to select the egress interface, an ARP entry is used forward the packet to the appropriate Ethernet MAC. With R-VPLS, the egress IP interface may be represented by multiple egress (VPLS service SAPs).

The following table describes how the ARP cache and MAC FIB entry states interact.

Table 48: Routing behavior in R-VPLS and interaction ARP cache and MAC FIB

ARP Cache entry	MAC FIB entry	Routing or system behavior
ARP Cache Miss (No Entry)	Known or Unknown	Triggers a request to control plane ARP processing module, to send out an ARP request, out of all the SAPs. (also known as virtual ports) of the VPLS instance.
ARP Cache Hit	Known	Forward to specific VPLS virtual port or SAP.
	Unknown	This behavior cannot happen typically on 7210 SAS-D, as and when an Layer 2 entry is removed from the FDB, the matching MAC address is also removed from the ARP cache. On 7210 SAS-K, the packet is sent out of all the SAPs of the VPLS instance.

5.4.5.1 The allow-ip-int-binding VPLS flag

The allow-ip-int-binding flag on a VPLS service context is used to inform the system that the VPLS service is enabled for routing support. The system uses the setting of the flag as a key to determine what type of ports the VPLS service may span.

The system also uses the flag state to define which VPLS features are configurable on the VPLS service to prevent enabling a feature that is not supported when routing support is enabled.

5.4.6 R-VPLS SAP support on standard Ethernet ports

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

5.4.6.1 LAG port membership constraints

If a LAG has a unsupported port type as a member, a SAP for the routing-enabled VPLS service cannot be created on the LAG. When one or more routing-enabled VPLS SAPs are associated with a LAG, a unsupported Ethernet port type cannot be added to the LAG membership.

5.4.6.2 VPLS feature support and restrictions

When the **allow-ip-int-binding** flag is set on a VPLS service, the following features cannot be enabled. Additionally, the flag cannot be enabled while any of these features are applied to the VPLS service. The following restrictions apply to both network mode and access-uplink mode, unless otherwise noted:

- Spoke SDPs and mesh SDPs cannot be configured in an R-VPLS service.
- In access-uplink mode, the VPLS service type cannot be MVPLS.
- In network mode, the VPLS service type must be R-VPLS; no other VPLS service is allowed.
- In access-uplink mode, MVR from an R-VPLS SAP to another SAP is not supported.
- In access-uplink mode, default QinQ SAPs are not supported in an R-VPLS service.
- In network mode, MVR from an R-VPLS SAP to another R-VPLS SAP is supported. See [R-VPLS and IGMPv3 snooping](#) for more information.
- The **allow-ip-int-binding** command cannot be used in a VPLS service that is acting as the G.8032 control instance.
- IP (IPv4 and IPv6) filters (ingress and egress) can be used with R-VPLS SAPs. Additionally, IP ingress override filters are supported, which affects the behavior of the IP filters attached to the R-VPLS SAPs.
- MAC filters (ingress and egress) are supported with R-VPLS SAPs only on the 7210 SAS-Mxp.
- In access-uplink mode, a VPLS IP interface is not allowed in an R-VPLS service, and an R-VPLS service/SAP cannot be configured with a VPLS IP interface.
- In access-uplink mode, the VPLS service can be configured with either access SAPs or access-uplink SAPs. In network mode, the VPLS service can be configured only with access SAPs or with SAPs on hybrid ports.
- In access-uplink mode, the VPLS service can use the following **svc-sap-type** values: any, dot1q-preserve, and null-star. Only specific SAP combinations are allowed for a specific **svc-sap-type**. Allowed SAP combinations are similar to those for plain VPLS service, except that default QinQ SAPs are not supported.
- In network mode, the VPLS service can use the following **svc-sap-type** values: any, null-star, and dot1q-preserve.

- G.8032 or MVPLS/STP based protection mechanisms can be used with an R-VPLS service. A separate G.8032 control instance or a separate MVPLS/STP instance must be used and the R-VPLS SAPs must be associated with these control instances such that the R-VPLS SAP forwarding state is driven by the control instance protocols.
- IP multicast is not supported in an R-VPLS service.
- IGMP snooping is supported in an R-VPLS service for 7210 SAS-T (network mode), 7210 SAS-Mxp, 7210 SAS-S, and 7210 SAS-Sx.
- In access-uplink mode, DHCP snooping is not supported for the SAPs configured in an R-VPLS service. Instead, DHCP relay can be enabled on the IES service associated with the R-VPLS service.
- In network mode, only on 7210 SAS-Mxp, DHCP IPv6 relay can be enabled on the IES service and VPRN service associated with the R-VPLS service. DHCPv6 snooping is not supported.
- In network mode, DHCP snooping is not supported for the SAPs configured in an R-VPLS service. Instead, DHCP relay (IPv4) can be enabled on the IES service associated with the R-VPLS service.
- In network mode, R-VPLS SAPs are allowed on access ports and hybrid ports.
- In network mode, an R-VPLS SAP drops packets received with extra tags. That is, if a packet is received on an R-VPLS SAP, with the number of tags greater than the SAP tags to which it is mapped, then it is dropped. This is true for all supported encapsulations (that is, null, dot1q, and QinQ encapsulations) of the port. For example, double-tagged packets received on a Dot1q SAP configured in an R-VPLS service is dropped on ingress.
- In the saved configuration file, for the R-VPLS service, the R-VPLS service instance appears twice, once for service creation and once with all the other configuration parameters. This is required to resolve references to the R-VPLS service and to execute the configuration without any errors.

5.4.7 VPLS SAP ingress IP filter override

When an IP Interface is attached to a VPLS service context, the VPLS SAP provisioned IP filter for ingress routed packets may be optionally overridden to provide special ingress filtering for routed packets. This allows different filtering for routed packets and non-routed packets. The filter override is defined on the IP interface bound to the VPLS service name. A separate override filter may be specified for IPv4 packet types.

If a filter for a specific packet type (IP) is not overridden, the SAP-specified filter is applied to the packet (if defined).

The following tables list ACL lookup behavior with and without Ingress Override filter attached to an IES interface in an R-VPLS service.

Table 49: ACL lookup behavior with ingress override filter attached to an IES interface in an R-VPLS service

Type of traffic	SAP ingress IPv4 filter	SAP egress IPv4 filter	Ingress override IPv4 filter
Destination MAC != IES IP interface MAC	Yes	Yes	No
Destination MAC = IES IP interface MAC and Destination IP on same subnet as IES interface	No	No	Yes

Type of traffic	SAP ingress IPv4 filter	SAP egress IPv4 filter	Ingress override IPv4 filter
Destination MAC = IES IP interface MAC and destination IP not on same subnet as IES IP interface and route to destination IP does not exist	No	No	No
Destination MAC = IES IP interface MAC and destination IP not on same subnet as IES IP interface and route to destination IP exists	No	No	Yes
Destination MAC = IES IP interface MAC and IP TTL = 1	No	No	No
Destination MAC = IES IP interface MAC and IPv4 packet with Options	No	No	No
Destination MAC = IES IP interface MAC and IPv4 Multicast packet	No	No	No

Table 50: ACL lookup behavior without ingress override filter attached to an IES interface in an R-VPLS service

Type of traffic	SAP ingress IPv4 filter	SAP egress IPv4 filter
Destination MAC != IES IP interface MAC	Yes	Yes
Destination MAC = IES IP interface MAC and Destination IP on same subnet as IES IP interface	Yes	No
Destination MAC = IES IP interface MAC and destination IP not on same subnet as IES IP interface and route to destination IP does not exist	No	No
Destination MAC = IES IP interface MAC and destination IP not on same subnet as IES IP interface and route to destination IP exists	Yes	No
Destination MAC = IES IP interface MAC and IP TTL = 1	No	No
Destination MAC = IES IP interface MAC and IPv4 packet with Options	No	No
Destination MAC = IES IP interface MAC and IPv4 Multicast packet	No	No

5.4.7.1 QoS support for VPLS SAPs and IP interface in a R-VPLS service

The follows information describes QoS support for VPLS SAPs and IP interface in an R-VPLS service:

- SAP ingress classification (IPv4, IPv6, and MAC criteria) is supported for SAPs configured in the service. SAP ingress policies cannot be associated with IES IP interface.
- On 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE, egress port-based queuing and shaping are available. It is shared among all the SAPs on the port.
- On 7210 SAS-Mxp, when the node is operating in SAP based queuing mode, unicast traffic sent out of R-VPLS SAPs uses SAP based egress queues while BUM traffic sent out of R-VPLS SAPs uses per port egress queues. When the node is operating in port based queuing mode, both unicast and BUM traffic sent out of R-VPLS SAPs uses per port egress queues. For more information, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Quality of Service Guide*.
- Port based Egress Marking is supported for both routed packets and bridged packets. The existing access egress QoS policy can be used for Dot1p marking and DSCP marking.
- In access-uplink mode, IES IP interface bound to R-VPLS services, IES IP interface on access SAPs, and IES IP interface on access-uplink SAPs are designed for use with inband management of the node. Consequently, they share a common set of queues for CPU-bound management traffic. All CPU bound traffic is policed to predefined rates before being queued into CPU queues for application processing. The system uses meters per application or a set of applications. It does not allocate meters per IP interface. The possibility of CPU overloading has been reduced by use of these mechanisms. Users must use appropriate security policies either on the node or in the network to ensure that this does not happen.

5.4.7.2 R-VPLS and routing protocols support

In access-uplink mode, R-VPLS is supported only in the base routing instance. Only IPv4 addressing support is available for IES interfaces associated with a R-VPLS service.

In network mode, R-VPLS is supported in both base routing instances (IES) and VPRN services. IPv4 and IPv6 addressing support is available for IES and VPRN IP interfaces associated with a R-VPLS service.

The following table lists the support available for routing protocols on IP interfaces bound to a VPLS service in access-uplink mode and network mode.

Table 51: Routing protocols on IP interfaces bound to a VPLS service

Services	Access-uplink	Network
Static-routing	Supported	Supported
BGP	Not Supported	Supported
OSPF	Not Supported	Supported
ISIS	Not Supported	Supported
BFD	Not Supported	Supported
VRRP	Not Supported	Supported
ARP and Proxy-Arp	ARP is supported	Both are supported

Services	Access-uplink	Network
DHCP Relay ¹⁵	Supported	Supported
DHCPv6 Relay	Not Supported	Supported only on the 7210 SAS-Mxp operating in network mode

5.4.7.3 Spanning tree and split horizon

The R-VPLS context supports all spanning tree capabilities that a non R-VPLS service supports. Service-based SHGs are not supported in an R-VPLS context.

5.4.8 R-VPLS MAC ACLs



Note:

This feature is supported only on the 7210 SAS-Mxp.


Operators deploy R-VPLS in enterprise networks on 7210 SAS nodes to provide Layer 2 and Layer 3 connectivity to end devices. To restrict access to the network, operators can use MAC ACLs to selectively allow traffic for specific MAC addresses and drop all other traffic, or the other way around.

Configure **filter mac** *mac-filter-id* in the **config>service>vpls>sap>egress** and **config>service>vpls>sap>ingress** contexts to associate an existing MAC filter policy with an R-VPLS service.

The following figure shows an example of an R-VPLS service that provides routed and bridging services to end devices connected to the enterprise LAN and that uses MAC ACLs to restrict traffic.

¹⁵ In access-uplink mode, DHCP relay can be configured for the IES interface associated with the R-VPLS service. DHCP snooping cannot be configured on the VPLS SAPs in the R-VPLS service. In network mode, DHCP relay can be configured for the IES interface associated with the R-VPLS service. DHCP snooping cannot be configured on the VPLS SAPs in the R-VPLS Service.

Figure 74: Secure access to network using MAC filter



sw1268

5.4.8.1 R-VPLS features supported with MAC ACLs

This following information describes R-VPLS feature support in conjunction with MAC ACLs:

- MAC criteria are supported with the following features in an R-VPLS service:
 - SAP ingress and egress for both bridged and routed packets with drop and forward action
 - matching for source MAC addresses, destination MAC addresses, Ethertype, and dot1p
 - filter entry counters and mirroring (ingress filters only)
 - either MAC or IP filters can be associated with SAP ingress or SAP egress; that is, they are mutually exclusive
- MAC filter policies are supported in R-VPLS services associated with either an IES or VPRN service.
- MAC filter policies are supported in **sap-scale-mode high** and **low**.
- MAC criteria filter policies associated with R-VPLS SAP ingress and SAP egress use the resources for the ingress ACL and egress ACL resource pools configured using the system resource profile. If there are no resources allocated for MAC criteria, the association of the MAC criteria filter policy fails. That is, the user must first allocate the required resources for MAC criteria using the system resource profile before associating a MAC criteria policy with SAPs.
- For routed packets, when the SAP ingress MAC filters and IP override ingress filters are configured, if the entries are matched in both filters and the drop action is configured, the drop action in either the MAC filter policy or IP override filter policy takes precedence. That is, if both the MAC filter policy and IP filter (IP override) policy match a routed packet, the packet is forwarded only if the action configured for all matched entries in both the MAC filter policy and IP filter policy is forward. Otherwise, the packet is dropped if there is a matched entry in either the MAC filter policy or IP filter policy with a drop action. If a counter is associated with both the MAC filter and IP override filter, the counter is incremented for matched entries in both the MAC filter and IP override filter. See [VPLS SAP ingress IP filter override](#) for more information about ACL behavior for IP override filters.

- When an IP filter is associated with an R-VPLS SAP on ingress and an IP override filter is associated with the R-VPLS IP interface, and the R-VPLS SAP and IP interface belong to the same R-VPLS service, routed traffic is matched only with IP override filter entries. It is not matched with R-VPLS SAP ingress filter entries. See [VPLS SAP ingress IP filter override](#) for more information about ACL behavior for IP override filters.

5.4.9 R-VPLS and IGMPv3 snooping



Note:

This feature is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

This feature (IGMPv3 snooping in R-VPLS) extends IGMP snooping to an R-VPLS service. On 7210 SAS, VPLS services that use MPLS uplinks (network mode) support IGMP snooping with IGMP v1 and v2 only. That is, IGMP v3 is not supported, which means that only Layer 2 multicast is supported. To support source-based IP multicast, support for IGMPv3 is needed. To provide source-based IP multicast, support for IGMP snooping v1, v2, and v3 is added to R-VPLS service. The IGMPv3 snooping in R-VPLS feature gives customers an option to use an R-VPLS service without a configured IP interface association to deliver IP multicast traffic in access Layer 2 networks. Users also have an option to configure MVR service.

IGMPv3 snooping in R-VPLS is supported only for IES (not for VPRNs).

For information about IGMP snooping in the context of VPLS, see [IGMP snooping in a VPLS service](#).

5.4.9.1 Configuration guidelines and restrictions for IGMP snooping in R-VPLS

The following items apply to IGMP snooping in R-VPLS and should be included with the regular VPLS multicast configuration guidelines (see [Configuration guidelines for IGMP snooping in VPLS service](#) and [R-VPLS supported functionality and restrictions](#)):

- R-VPLS without an IP interface association can be used to emulate VPLS service with support for IGMPv3 snooping.
- R-VPLS with or without IP interface association can be used for IGMPv3 snooping. If enabling MVR on the service then the service should not have an IP interface association.
- IGMPv3 snooping can be enabled in the context of the R-VPLS (both with and without MVR). It cannot be enabled in regular VPLS service. Regular VPLS service supports IGMP v1 and v2 only.
- MVR can be configured in an R-VPLS without an IP interface association. It can be used to leak multicast traffic to a user R-VPLS service with an IP interface configuration. Therefore, a user R-VPLS can be used to forward both unicast and multicast services.

In addition, the following list of guidelines and restrictions pertain to IGMP snooping in an R-VPLS service:

- R-VPLS service can only have a single SAP per port configured in a service. That is, two SAPs on the same port cannot be configured in the same service.
- Spoke-SDPs and mesh SDPs cannot be configured in an R-VPLS service.
- On 7210 SAS devices, on ingress of a port, multicast traffic can be processed in the context of either **igmp-snooping** (Layer 2 Ethernet multicast with IGMP v1 or v2 snooping) or **I3-multicast** (either multicast in an Layer 3 service or IGMP snooping in an R-VPLS), but not both. That is, it is not possible to configure SAPs on the port such that one SAP is a receiver for multicast traffic to be processed by IGMP snooping, and another SAP is a receiver for multicast traffic to be processed by IP multicast in

the context of Layer 3 service or R-VPLS. An option per port is available using the **configure> port> ethernet> multicast-ingress {l2-mc | ip-mc}** command to enable one or the other. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide* for more information about this command. By default, IGMP snooping is enabled to be backward compatible. Users need to explicitly change the IGMP snooping configuration to allow processing of received multicast traffic as IP multicast in the context of Layer 3 service or R-VPLS.

- If a VPLS SAP is configured on the same port as the port on which IP multicast is enabled, then multicast traffic received on the SAP is dropped. Unicast, broadcast, and unknown-unicast packets received on the SAP are forwarded appropriately. This behavior is true only for VPLS SAPs and does not apply to VPLS SDPs, Epipe SAPs, and Epipe SDPs.
- With R-VPLS multicast and when using MVR capability, a port on which receivers are present can be configured to perform either Layer 2 multicast replication (that is, no IP TTL decrement and no source MAC replacement) or Layer 3 multicast replication (that is, IP TTL is decremented and source MAC is replaced with 7210 SAS chassis MAC or IP interface MAC). An option to use either Layer 2 or Layer 3 multicast replication is available using the **configure port ethernet multicast-egress {l2-switch | l3-forward}** command. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide* for more information about this command. All SAPs on the port have the same behavior.
- An MVR R-VPLS must be configured without an IP interface and will support Layer 2 forwarding of both unicast and multicast traffic (that is, no IP forwarding).
- A user RPVLS can be configured with an IP interface and will support Layer 2 forwarding of both unicast and multicast (with (S,G) IP multicast replications) and support Layer 3 forwarding of unicast traffic.
- On 7210 SAS-Mxp and 7210 SAS-R6, when using SAP-based egress queues and scheduler, R-VPLS BUM traffic uses per port egress queues—not per SAP egress queues.
- In an MVR configuration, the **svc-sap-type** of the R-VPLS service that is the source (also known as MVR R-VPLS service) and the **svc-sap-type** of the R-VPLS service that is the sink (also known as user R-VPLS service) should match.
- On 7210 SAS-Mxp and 7210 SAS-T, the MVR R-VPLS service configured with IGMPv3 snooping shares resources with TWAMP. An increase in one decreases the amount of resources available for the other. Contact your Nokia representative for more information about scaling of these features.

5.4.10 R-VPLS supported functionality and restrictions

R-VPLS supported functionality and restrictions for both access-uplink and network modes are listed as follows. The following items are applicable to both the modes, unless specified explicitly:

- Static ARP cannot be configured with an IES IP interface that is associated with an R-VPLS, though static MAC can be configured in an R-VPLS service.
- In access-uplink mode, only static routes are supported. No dynamic routing protocols are supported.
- In network mode, both static routing and dynamic routing protocols are supported.
- Whenever a VPLS FIB entry is removed either because of user action, aging or mac-move, the corresponding ARP entry whose MAC address matches that of the MAC in the FIB is removed from the ARP cache.
- In access-uplink mode, R-VPLS is supported only in the base routing instance. Only IPv4 addressing support is available for IES interfaces associated with an R-VPLS service.

- In network mode, R-VPLS is supported in both the base routing instance (IES) and VPRN services. IPv4 addressing support is available for IES and VPRN IP interfaces associated with an R-VPLS service.
- In access-uplink mode, IPv6 addressing support is not available for IES interfaces associated with an R-VPLS service.
- In network mode, only on the 7210 SAS-Mxp, IPv6 addressing support is available for IES and VPRN interfaces associated with an R-VPLS service.
- In both network mode and access-uplink mode, multiple SAPs configured on the same port cannot be part of the same R-VPLS service. That is, a single service can only be configured with a single SAP on a specific port.
- Service MTU configuration is not supported in the R-VPLS service.
- In network mode, in **any** service (that is, **svc-sap-type** set to **any**), null sap accepts only untagged packets. Received tagged packets are dropped.
- In network mode, MPLS protocols (for example, RSVP and LDP) cannot be enabled on an R-VPLS IP interface.
- In network mode, MPLS-TP cannot use an R-VPLS IP interface.
- In network mode, R-VPLS SAPs can be configured on an MC-LAG LAG.
- Service-based SHGs are not supported in an R-VPLS service.
- Spoke SDPs and mesh SDPs cannot be configured in an R-VPLS service.

5.5 Epipe emulation using dot1q VLAN range SAP in VPLS with G.8032




Note:

This feature is only supported on 7210 SAS-T operating in access-uplink mode.

On the node where the service originates, in addition to the access dot1q range SAP, the service needs to be configured with access-uplink SAPs on the two G.8032 ring ports. G.8032 mechanism is used to for breaking the loop in the ring and VPLS service protection. The intermediate nodes on the ring needs to use VPLS service with access-uplink SAPs on the ring ports and use the same G.8032 instance for protection, as one is used for service protection on the originating node.

The following figure shows how two business offices, served by an operator are connected in a ring network deployment using Dot1q range SAPs and a VPLS service with G.8032 for protection.

Figure 75: Epipe Emulation in a Ring using VPLS with G.8032



The following are the requirements to provide for an Epipe service connectivity between two business sites:

- Transport all the VLANs used by the internal enterprise network of the businesses.
- Support high availability for the service between the business sites by protecting against failure of the links or nodes in the ring.

To achieve connectivity between two business sites in access-uplink/Layer 2 mode is to configure SAPs for each of the individual VLANs used in the enterprise network in a VPLS service and use G.8032 for protection. The number of VLANs that was supported is limited by the number of SAPs supported on the platform.

The 7210 SAS platforms, currently support the use of Dot1q range SAPs with only Epipe services in either network/MPLS mode or access-uplink/Layer 2 mode. Dot1q range SAPs allows operators to transport a range of VLANs by providing similar service treatment (service treatment refers to forwarding decision along with encapsulation used, QoS and ACL processing, accounting, and so on) to all the VLANs configured in the range. It simplifies service configuration and allows operators to scale the number of VLANs that can be handled by the node. This took care of the need to support hundreds of VLANs using a single SAP or a small number of SAPs. When MPLS the mode is deployed in ring topology, operators have the option of using different redundancy mechanisms such as FRR, primary/secondary LSPs, Active/Standby PWs, to improve Epipe service availability. No such option is available to protect Epipe service in Layer 2 mode when deployed in a ring topology. Additionally many operators prefer G.8032 based ring protection mechanism, because a single control instance on the ring can potentially protect all the VPLS services on the ring.

This feature allows operators to deploy Epipe services in a ring topology when using Layer 2 mode, by emulating an Epipe service using a VPLS service with G.8032 protection and at the same time provides the benefits of using dot1q range SAPs. The user should ensure that the VPLS service is a point-to-point service. This is achieved by configuring a VPLS service with an access dot1q range SAP used at the customer handoff on one node in the ring and an access dot1q range SAP in a customer handoff of a VPLS service on another node (that is, at the other end of the Epipe), such that there are only two endpoints for the service in the network.

On the node where the service originates, in addition to the access dot1q range SAP, the service needs to be configured with access-uplink SAPs on the two G.8032 ring ports. G.8032 mechanism is used to for

breaking the loop in the ring and VPLS service protection. The intermediate nodes on the ring needs to use VPLS service with access-uplink SAPs on the ring ports and use the same G.8032 instance for protection, as one is used for service protection on the originating node.

5.5.1 Epipe emulation configuration guidelines and restrictions

The VPLS service with dot1q-range SAPs uses **svc-sap-type** of dot1q-range and supports limited functionality in comparison to a normal VPLS service. The following list provides more information about the feature, configuration guidelines, and restrictions:

- The user can define access dot1q range SAPs, which specifies a group of VLANs which receive similar service treatment; that is, forwarding behavior, SAP ingress QoS treatment and SAP (behavior similar to that available in Epipe service) and allows it to be configured in a VPLS service:
 - On the node, where the service originates, in addition to the access dot1q range SAP, the service should be configured with Q1.* SAPs on the two G.8032 ring ports. The access or access-uplink Q1.*SAPs can be used, but the access-uplink SAPs are recommended for use. The user cannot configure any other SAPs in the same VPLS service.
 - There is no special configuration required on intermediate nodes, that is, the ring nodes which do not terminate or originate the service. The nodes should be configured for providing transit VPLS service and the VPLS service must use the same G.8032 instance for protection as is used by the service on originating and terminating node.
 - The Epipe service on 7210, currently does not check if the inner tag received on a Q1.* SAP is within the range of the configured VLANs. VPLS service too has the same behavior.
- Support for SAP Ingress QoS, Ingress and Egress ACLs, accounting, and other services, for dot1q range SAP configured in a VPLS service matches the support available in Epipe service.
- G.8032 mechanism is used for loop detection in ring network and service protection. A separate VPLS service representing the G.8032 control instance must be configured and the state should be associated with this service:
 - Use of dot1q range SAPs to provide service on the interconnection node, in a G.8032 major-ring/sub-ring deployment, when using the virtual channel, is not supported. This restriction is not applicable when the interconnection node in a G.8032 major-ring/sub-ring is configured without a virtual channel.
- mVPLS/xSTP support is available for use with Q1.* SAP on the ring ports to break the loop. This is a add-on to the G.8032 support.
- Broadcast, Unknown Unicast and Multicast (BUM) traffic is flooded in the service.
- Learning is enabled on the service by default, to avoid the need to flood the service traffic out of one of the ring ports, after network MAC addresses are learned. The user has an option to disable learning per service. Learning enable/disable per SAP is not supported.
- MAC limiting is available per service. MAC limiting per SAP is not supported.
- CFM OAM is supported. The support for UP MEPs on the dot1q range SAP in the service to be used for fault management and performance management using the CFM/Y.1731 OAM tools is available:
 - Only UP MEP is allowed to be configured only on the dot1q VLAN range SAPs. CFM/Y.1731 tools can be used for trouble shooting and performance measurements. User must pick a VLAN value from the range of VLANs configured for the dot1-range SAP using the CLI command **config>eth-cfm>domain>association>bridge-identifier VLAN** and enable the use of using the CLI command

primary-vlan-enable under the MEP CLI context. It is used as the VLAN tag in the packet header for all the CFM/Y.1731 messages sent out in the context of the UP MEP.

- Down MEPs and MIPs are not allowed to be configured.
- Fault propagation is not supported with UP MEPs for dot1q range SAP in access-uplink mode.
- CFM support is not available for SAPs on the ring ports.
- IGMP snooping and MVR is not supported.

5.6 Configuring a VPLS service with CLI

This section provides information to configure VPLS services using the command line interface.

5.6.1 Basic configuration

The following fields require specific input (there are no defaults) to configure a basic VPLS service:

- Customer ID (see [Configuring customers 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 (only for 7210 SAS devices in network mode) for each far-end node.

Example: Local VPLS service on ALA-1

For 7210 SAS devices configured in access-uplink mode:

```
*A:SAS>config>service>vpls# info
-----
      stp
        shutdown
      exit
      sap 1/1/1:10.* create
        ingress
          filter mac 1
        exit
      exit
      sap 1/1/2:10.* create
      exit
      no shutdown
-----
*A:SAS>config>service>vpls#
```

```
*A:ALA-1>config>service>vpls# info
-----
...
      vpls 9001 customer 6 create
        description "Local VPLS"
        stp
          shutdown
        exit
        sap 1/2/2:0 create
          description "SAP for local service"
        exit
        sap 1/1/5:0 create
```



```

        description "SAP for local service"
        exit
        no shutdown
-----
*A:ALA-1>config>service>vpls#
*A:ALA-1>config>service# info
-----
...
    vpls 7 customer 7 create
        stp
            shutdown
        exit
        sap 1/1/21 create
        exit
        sap lag-1:700 create
        exit
        no shutdown
    exit
...
-----
*A:ALA-1>config>service#

```

Example: Distributed VPLS service between ALA-1, ALA-2, and ALA-3

```

*A:ALA-1>config>service# info
-----
...
    vpls 9000 customer 6 create
        shutdown
        description "This is a distributed VPLS."

        stp
            shutdown
        exit
        sap 1/1/5:16 create
            description "VPLS SAP"
        exit
        spoke-sdp 2:22 create
        exit

    exit
...
-----
*A:ALA-1>config>service#
*A:ALA-2>config>service# info
-----
...
    vpls 9000 customer 6 create
        description "This is a distributed VPLS."

        stp
            shutdown
        exit
        sap 1/1/5:16 create
            description "VPLS SAP"
        exit
        spoke-sdp 2:22 create
        exit

```



```

        no shutdown
        exit
    ...
-----
*A:ALA-2>config>service#

*A:ALA-3>config>service# info
-----
...
    vpls 9000 customer 6 create
        description "This is a distributed VPLS."

        stp
            shutdown
        exit
        sap 1/1/3:33 create
            description "VPLS SAP"
        exit
        spoke-sdp 2:22 create
        exit

        no shutdown
        exit
    ...
-----
*A:ALA-3>config>service#

```

5.6.2 Common configuration tasks

About this task

This section provides a brief overview of the tasks that must be performed to configure both local VPLS services and provides the syntax commands.

For VPLS services:

Procedure

- Step 1.** Associate VPLS service with a customer ID
- Step 2.** Define SAPs:
 - Select nodes and ports
 - Optional - select QoS policies other than the default (configured in **config>qos** context)
 - Optional - select filter policies (configured in **config>filter** context)
 - Optional - select accounting policy (configured in **config>log** context)
- Step 3.** Modify STP default parameters (optional) (see [VPLS and spanning tree protocol](#))
- Step 4.** Enable service

5.6.3 Configuring VPLS components

5.6.3.1 Creating a VPLS service

Use the following syntax to create a VPLS service.

```
For 7210 SAS (other than access-uplink mode):
config>service# vpls service-id [customer customer-id] [create][vpn vpn-id] [m-vpls]
description description-string
no shutdown
For 7210 SAS in Access uplink mode:
config>service# vpls service-id [customer customer-id] [create][vpn vpn-id] [m-vpls]
<service-id> [customer <customer-id>] [create] [vpn <vpn-id>] [m-vpls] [svc-sap-type {null-
star|dot1q-preserve|any}] [customer-vid <vlan-id>]
description description-string
no shutdown
```

The following is a sample VPLS configuration output.

```
*A:ALA-1>config>service>vpls# info
-----
...
    vpls 1000 customer 1 create
        description "This is a VPLS with NULL SAP"
        stp
            shutdown
        exit
        no shutdown
    exit
    vpls 2000 customer 6 create
        description "This is a Distributed VPLS with DOT1Q SAP"
        stp
            shutdown
        exit
        no shutdown
    exit
...
-----
*A:ALA-1>config>service>vpls#
```

5.6.3.1.1 Enabling MAC move

The **mac-move** feature is useful to protect against undetected loops in your VPLS topology as well as the presence of duplicate MACs in a VPLS service. For example, if two clients in the VPLS have the same MAC address, the VPLS will experience a high relearn rate for the MAC and will shut down the SAP when the threshold is exceeded.

Use the following syntax to configure **mac-move** parameters.

```
config>service# vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls]
mac-move
move-frequency frequency
retry-timeout timeout
no shutdown
```

Example

The following is a sample of **mac-move** information.

```
*A:ALA-1# show service id 6 all
....
*A:ALA-1#
-----
Forwarding Database specifics
-----
Service Id       : 1150           Mac Move         : Disabled
Mac Move Rate    : 2             Mac Move Timeout  : 10
Table Size       : 1000          Total Count      : 1000
Learned Count    : 1000          Static Count     : 0
Remote Age       : 900           Local Age        : 300
High WaterMark   : 95%          Low Watermark    : 90%
Mac Learning     : Enabl         Discard Unknown   : Dsabl
Mac Aging        : Enabl         Relearn Only     : True
=====
....
*A:ALA-1#
```

5.6.3.1.2 Configuring STP bridge parameters in a VPLS

Modifying some of the Spanning Tree Protocol parameters allows the operator to balance STP between resiliency and speed of convergence extremes. Modifying particular parameters must be done in the constraints of the following two formulas:

$$2 \times (\text{Bridge\_Forward\_Delay} - 1.0 \text{ seconds}) \geq \text{Bridge\_Max\_Age}$$

$$\text{Bridge\_Max\_Age} \geq 2 \times (\text{Bridge\_Hello0\_Time} + 1.0 \text{ seconds})$$

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

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

5.6.3.1.2.1 Bridge STP admin state

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

When STP on the VPLS is administratively disabled, any BPDUs are forwarded transparently through the 7210 SAS. When STP on the VPLS is administratively enabled, but the administrative state of a SAP is down, BPDUs received on such a SAP are discarded.

```
config>service>vpls service-id# stp
no shutdown
```

5.6.3.1.2.2 Mode

To be compatible with the different iterations of the IEEE 802.1D standard, the 7210 SAS supports several variants of the Spanning Tree protocol:

- **rstp** - Rapid Spanning Tree Protocol (RSTP) compliant with IEEE 802.1D-2004 - default mode
- **dot1w** - compliant with IEEE 802.1w
- **comp-dot1w** - operation as in RSTP but backwards compatible with IEEE 802.1w (this mode was introduced for interoperability with some MTU types)
- **mstp** - compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q REV/D5.0-09/2005 (this mode of operation is only supported in an mVPLS)
- **pmstp** - compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q REV/D3.0-04/2005 but with some changes to make it backwards compatible to 802.1Q 2003 edition and IEEE 802.1w

See section [Spanning tree operating modes](#) for more information about these modes.

```
config>service>vpls service-id# stp
mode {rstp | comp-dot1w | dot1w | mstp|pmstp}
```

Default: rstp

5.6.3.1.2.3 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.

```
config>service>vpls service-id# stp
priority bridge-priority
```

Range:	1 to 65535
Default:	32768
Restore Default:	no priority

5.6.3.1.2.4 Max age

The **max-age** command indicates how many hops a BPDU can traverse the network starting from the root bridge. The message age field in a BPDU transmitted by the root bridge is initialized to 0. Each other bridge will take the message_age value from BPDUs received on their root port and increment this value by 1. The message_age therefore reflects the distance from the root bridge. BPDUs with a message age exceeding max-age are ignored.

STP uses the max-age value configured in the root bridge. This value is propagated to the other bridges by the BPDUs.

The default value of **max-age** is 20. This parameter can be modified within a range of 6 to 40, limited by the standard STP parameter interaction formulas.

```
config>service>vpls service-id# stp
max-age max-info-age
```

Range: 6 to 40 seconds

Default: 20 seconds

Restore Default: no max-age

5.6.3.1.2.5 Forward delay

RSTP, as defined in the IEEE 802.1D-2004 standards, will transition to the forwarding state by a handshaking mechanism (rapid transition), without any waiting times. If handshaking fails (for example on shared links), the system falls back to the timer-based mechanism defined in the original STP (802.1D-1998) standard.

A shared link is a link with more than two Ethernet bridges (for example, a shared 10/100BaseT segment). The **port-type** command is used to configure a link as point-to-point or shared (see section [SAP link type](#)).

For timer-based transitions, the 802.1D-2004 standard defines an internal variable forward-delay, which is used in calculating the default number of seconds that a SAP spends in the discarding and learning states when transitioning to the forwarding state. The value of the forward-delay variable depends on the STP operating mode of the VPLS instance:

- In **rstp** mode, but only when the SAP has not fallen back to legacy STP operation, the value configured by the **hello-time** command is used.
- In all other situations, the value configured by the **forward-delay** command is used.

```
config>service>vpls service-id# stp
forward-delay seconds
```

Range: 4 to 30 seconds

Default: 15 seconds

Restore Default: no forward-delay

5.6.3.1.2.6 Hello time

The **hello-time** command configures the Spanning Tree Protocol (STP) hello time for the Virtual Private LAN Service (VPLS) STP instance.

The *seconds* parameter defines the default timer value that controls the sending interval between BPDU configuration messages by this bridge, on ports where this bridge assumes the designated role.

The active hello time for the spanning tree is determined by the root bridge (except when the STP is running in RSTP mode, then the hello time is always taken from the locally configured parameter).

The configured hello-time value can also be used to calculate the bridge forward delay, see [Forward delay](#).

```
config>service>vpls service-id# stp
hello-time hello-time
```

Range:	1 to 10 seconds
Default:	2 seconds
Restore Default:	no hello-time

5.6.3.1.2.7 Hold count

The **hold-count** command configures the peak number of BPDUs that can be transmitted in a period of one second.

```
config>service>vpls service-id# stp
hold-count count-value
```

Range:	1 to 10
Default:	6
Restore Default:	no hold-count

5.6.3.1.2.8 MST instances

You can create up to 15 MST-instances. They can range from 1 to 4094. By changing path-cost and priorities, you can make sure that each instance will form it's own tree within the region, therefore making sure different VLANs follow different paths.

You can assign non overlapping VLAN ranges to each instance. VLANs that are not assigned to an instance are implicitly assumed to be in instance 0, which is also called the CIST. This CIST cannot be deleted or created.

The parameter that can be defined per instance are **mst-priority** and **vlan-range**:

- **mst-priority**

The bridge-priority for this specific mst-instance. It follows the same rules as bridge-priority. For the CIST, the bridge-priority is used.

- **vlan-range**

The VLANs are mapped to this specific mst-instance. If no VLAN-ranges are defined in any mst-instances, then all VLANs are mapped to the CIST.

5.6.3.1.2.9 MST max hops

The mst-max-hops command defines the maximum number of hops the BPDU can traverse inside the region. Outside the region max-age is used.

5.6.3.1.2.10 MST name

The MST name defines the name that the operator gives to a region. Together with MST revision and the VLAN to MST-instance mapping, it forms the MST configuration identifier. Two bridges that have the same MST configuration identifier form a region if they exchange BPDUs.

5.6.3.1.2.11 MST revision

The MST revision together with MST-name and VLAN to MST-instance mapping define the MST configuration identifier. Two bridges that have the same MST configuration identifier form a region if they exchange BPDUs.

5.6.3.2 Configuring a VPLS SAP

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

5.6.3.2.1 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.

```
*A:ALA-1>config>service# info
-----
vpls 1150 customer 1 create
  fdb-table-size 1000
  fdb-table-low-wmark 5
  fdb-table-high-wmark 80
  local-age 60
  stp
    shutdown
  exit
  sap 1/1/1:1155 create
  exit
  sap 1/1/2:1150 create
  exit
  no shutdown
exit
-----
*A:ALA-1>config>service#
```

5.6.3.2.2 Distributed VPLS SAPs



Note:

Distributed VPLS is supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

To configure a distributed VPLS service, you must configure service entities on originating and far-end nodes. You must use the same service ID on all ends (for example, create a VPLS service ID 9000 on

ALA-1, ALA-2, and ALA-3). A distributed VPLS consists of a SAP on each participating node and an SDP bound to each participating node.

For SDP configuration information, see [Configuring an SDP](#). For SDP binding information, see [Configuring SDP bindings](#).

Example

The following is a sample configuration output of VPLS SAPs configured for ALA-1, ALA-2, and ALA-3.

```
*A:ALA-3>config>service# info
-----
      vpls 1150 customer 1 create
        fdb-table-size 1000
        fdb-table-low-wmark 5
        fdb-table-high-wmark 80
        local-age 60
        stp
          shutdown
        exit
        sap 1/1/1:1155 create
        exit
        sap 1/1/2:1150 create
        exit
        no shutdown
      exit
-----
*A:ALA-3>config>service#
```

5.6.3.2.3 Configuring default QinQ SAPs to pass all traffic from access to access-uplink Port without any tag modifications



Note:

Default QinQ SAPs are only supported on 7210 SAS platforms operating in access-uplink mode.

Example

The following is a sample VPLS SAP configuration output of Default QinQ SAPs.

```
ALA-1>config>service# vpls 9 customer 1 svc-sap-type null-star create
      shutdown
      stp
        shutdown
      exit
      sap 1/1/5:*. * create
        statistics
          ingress
        exit
      exit
      exit
      sap 1/1/6:*. * create
        statistics
          ingress
        exit
      exit
      exit
      exit
```


5.6.3.2.4 Configuring SAP-specific STP parameters

When a VPLS has STP enabled, each SAP within the VPLS has STP enabled by default.

5.6.3.2.4.1 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 VPLS level, but disabled on the SAP, received BPDUs are discarded. Discarding the incoming BPDUs allows STP to continue to operate 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.

```
config>service>vpls>sap>stp#
[no] shutdown
```

Range: shutdown or no shutdown

Default: no shutdown (SAP admin up)

5.6.3.2.4.2 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 it's own virtual port number that is unique to every other SAP defined on the VPLS. The virtual port number is assigned at the time that the SAP is added to the VPLS.

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.

```
config>service>vpls>sap# stp
port-num number
```

Range: 1 — 2047

Default: (automatically generated)

Restore Default: no port-num

5.6.3.2.4.3 SAP priority

SAP priority allows a configurable “tie breaking” parameter to be associated with a SAP. When configuration BPDUs are being received, the configured SAP priority will be used in some circumstances to determine whether a SAP will be designated or blocked.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16 bit value. In the latest STP standard (802.1D-2004) only the upper 4 bits of the port priority field are used to encode the SAP priority. The remaining 4 bits are used to extend the port ID field into a 12 bit virtual port number field. The virtual port number uniquely references a SAP within the STP instance. See [SAP virtual port number](#) for more information about 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.

```
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

5.6.3.2.4.4 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 incremental with the configured path cost for that SAP. When BPDUs are sent out other egress SAPs, the newly calculated root path cost is used.

STP suggests that the path cost is defined as a function of the link bandwidth. Because SAPs are controlled by complex queuing dynamics, in the 7210 SAS the STP path cost is a purely static configuration.

The default value for SAP path cost is 10. This parameter can be modified within a range of 1 to 65535, 1 being the lowest cost.

```
config>service>vpls>sap>stp#
path-cost sap-path-cost
```

Range: 1 to 200000000

Default:	10
Restore Default:	no path-cost

5.6.3.2.4.5 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 (see [Forward delay](#)). 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.

```
config>service>vpls>sap>stp#  
[no] edge-port
```

Default: no edge-port

5.6.3.2.4.6 SAP auto edge

The SAP **edge-port** command is used to instruct STP to dynamically decide whether the SAP is connected to another bridge.

If auto-edge is enabled, and STP concludes there is no bridge behind the SAP, the OPER_EDGE variable will dynamically be set to true. If auto-edge is enabled, and a BPDU is received, the OPER_EDGE variable will dynamically be set to true (see [SAP edge port](#)).

Valid values for SAP auto-edge are enabled and disabled with enabled being the default.

```
config>service>vpls>sap>stp#  
[no] auto-edge
```

Default: auto-edge

5.6.3.2.4.7 SAP link type

The SAP **link-type** parameter instructs STP on the maximum number of bridges behind this SAP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions). If more than two bridges are connected by a shared media, their SAPs should all be configured as shared, and timer-based transitions are used.

Valid values for SAP link-type are shared and pt-pt with pt-pt being the default.

```
config>service>vpls>sap>stp#
link-type {pt-pt|shared}
```

Default: link-type pt-pt

Restore Default: no link-type

5.6.3.2.4.8 MST instances

The SAP mst-instance command is used to create MST instances at the SAP level. MST instance at a SAP level can be created only if MST instances are defined at the service level.

The parameters that can be defined per instance are **mst-path-cost** and **mst-port-priority**:

- **mst-path-cost**
Specifies path-cost within a specific MST instance. The path-cost is proportional to link speed.
- **mst-port-priority**
Specifies the port priority within a specific MST instance.

5.6.3.2.5 STP SAP operational states

The operational state of STP within a SAP controls how BPDUs are transmitted and handled when received.

5.6.3.2.5.1 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 disabled state for the configured forward-delay duration.

5.6.3.2.5.2 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. The duration of the discarding state is described in section [Forward delay](#).

**Note:**

In previous versions of the STP standard, the discarding state was called a blocked state.

5.6.3.2.5.3 Operationally learning

The learning state allows population of the MAC forwarding table before entering the forwarding state. In this state, no user traffic is forwarded.

5.6.3.2.5.4 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.

5.6.3.2.5.5 SAP BPDUs encapsulation state

IEEE 802.1d (referred as dot1d) and the Cisco per VLAN Spanning Tree (PVST) BPDUs encapsulations are supported on a per SAP basis. The STP is associated with a VPLS service like PVST is per VLAN. The difference between the two encapsulations is in the Ethernet and LLC framing and a type-length-value (TLV) field trailing the BPDUs. The encapsulation format cannot be configured by the user, the system automatically determines the encapsulation format based on the BPDUs received on the port.

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

Table 52: SAP BPDUs encapsulation states

Field	dot1d encap-type null	dot1d encap-type dot1q	PVST encap-type null	PVST encap-type dot1q
Destination MAC	01:80:c2:00:00:00	01:80:c2:00:00:00	N/A	01:00:0c:cc:cc:cd
Source MAC	Sending Port MAC	Sending Port MAC	N/A	Sending Port MAC
EtherType	N/A	0x81 00	N/A	0x81 00
Dot1p and CFI	N/A	0xe	N/A	0xe
Dot1q	N/A	VPLS SAP ID	N/A	VPLS SAP encap value
Length	LLC Length	LLC Length	N/A	LLC Length
LLC DSAP SSAP	0x4242	0x4242	N/A	0xaaaa (SNAP)
LLC CNTL	0x03	0x03	N/A	0x03
SNAP OUI	N/A	N/A	N/A	00 00 0c (Cisco OUI)

Field	dot1d encap-type null	dot1d encap-type dot1q	PVST encap-type null	PVST encap-type dot1q
SNAP PID	N/A	N/A	N/A	01 0b
CONFIG	Standard 802.1d	Standard 802.1d	N/A	Standard 802.1d
TLV: Type & Len	N/A	N/A	N/A	58 00 00 00 02
TLV: VLAN	N/A	N/A	N/A	VPLS SAP encap value
Padding	As Required	As Required	N/A	As Required

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

- **Dot1d**

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

- **PVST**

This state specifies that the switch is currently sending proprietary encapsulated BPDUs. PVST BPDUs are only supported on Ethernet interfaces with the encapsulation type set to dot1q. The SAP continues in the PVST BPDUs encapsulation state until a dot1d encapsulated BPDU is received, in which case, the SAP reverts to the dot1d encapsulation state. Each received BPDU must be correctly IEEE 802.1q tagged with the encapsulation value defined for the SAP. PVST BPDUs are silently discarded if received when the SAP is on an interface defined with a null encapsulation type.

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

5.6.3.2.6 Configuring VPLS SAPs with per service split horizon



Note:

Per-service split horizon groups are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

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.

Example

The following is a sample VPLS configuration output with split horizon enabled.

```
*A:ALA-1>config>service# info
-----
...
  vpls 800 customer 6001 vpn 700 create
    description "VPLS with split horizon for DSL"
    stp
      shutdown
    exit
    sap 1/1/3:100 split-horizon-group DSL-group1 create
      description "SAP for residential bridging"
    exit
    sap 1/1/3:200 split-horizon-group DSL-group1 create
      description "SAP for residential bridging"
    exit
    split-horizon-group DSL-group1
      description "Split horizon group for DSL"
    exit
    no shutdown
  exit
...
-----
*A:ALA-1>config>service#
```

5.6.3.3 Configuring SDP bindings



Note:

SDPs are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

VPLS provides scaling and operational advantages. A hierarchical configuration eliminates the need for a full mesh of VCs between participating devices. Hierarchy is achieved by enhancing the base VPLS core mesh of VCs with access VCs (spoke) to form two tiers. Spoke-SDPs are generally created between Layer 2 switches and placed at the Multi-Tenant Unit (MTU). The PE routers are placed at the service provider's Point of Presence (POP). Signaling and replication overhead on all devices is considerably reduced.

A spoke-SDP is treated like the equivalent of a traditional bridge port where flooded traffic received on the spoke-SDP is replicated on all other "ports" (other spoke-SDPs or SAPs) and not transmitted on the port it was received (unless a split horizon group was defined on the spoke-SDP, see section [Configuring VPLS spoke-SDPs with split horizon](#)).

A spoke-SDP connects a VPLS service between two sites and, in its simplest form, could be a single tunnel LSP. A set of ingress and egress VC labels are exchanged for each VPLS service instance to be transported over this LSP. The PE routers at each end treat this as a virtual spoke connection for the VPLS service in the same way as the PE-MTU connections. This architecture minimizes the signaling overhead and avoids a full mesh of VCs and LSPs between the two metro networks.

A VC-ID can be specified with the SDP-ID. The VC-ID is used instead of a label to identify a virtual circuit. The VC-ID is significant between peer SRs on the same hierarchical level. The value of a VC-ID is conceptually independent from the value of the label or any other datalink specific information of the VC.

5.6.3.3.1 Configuring VPLS spoke-SDPs with split horizon



Note:

Split horizon groups are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

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.

Example

The following is a sample VPLS configuration output with split horizon enabled.

```
*A:ALA-1>config>service# info
-----
...
vpls 800 customer 6001 vpn 700 create
description "VPLS with split horizon for DSL"
stp
shutdown
exit
spoke-sdp 51:15 split-horizon-group DSL-group1 create
exit
split-horizon-group DSL-group1
description "Split horizon group for DSL"
exit
no shutdown
exit
...
-----
*A:ALA-1>config>service#
```

5.6.4 Configuring VPLS redundancy

This section describes the service management tasks.

5.6.4.1 Creating a management VPLS for SAP protection

About this task

This section provides a brief overview of the tasks that must be performed to configure a management VPLS for SAP protection and provides the CLI commands, see [Figure 76: Example configuration for protected VPLS SAP](#). The following tasks should be performed on both nodes providing the protected VPLS service.

Before configuring a management VPLS, first read [VPLS redundancy](#) for an introduction to the concept of management VPLS and SAP redundancy:


Procedure

Step 1. Create an SDP to the peer node.

Step 2. Create a management VPLS.

- Step 3.** Define a SAP in the m-vpls on the port toward the 7210 SAS. Note that the port must be dot1q. The SAP corresponds to the (stacked) VLAN on the 7210 SAS in which STP is active.
- Step 4.** Optionally modify STP parameters for load balancing (see [Configuring load balancing with management VPLS](#)).
- Step 5.** Create an SDP in the m-vpls using the SDP defined in Step 1. Ensure that this SDP runs over a protected LSP.
- Step 6.** Enable the management VPLS service and verify that it is operationally up.
- Step 7.** Create a list of VLANs on the port that are to be managed by this management VPLS.
- Step 8.** Create one or more user VPLS services with SAPs on VLANs in the range defined by Step 6.

Figure 76: Example configuration for protected VPLS SAP



Example: Creating a management VPLS for SAP protection

Use the following commands to create a management VPLS for SAP protection.

```
config>service# vpls service-id [customer customer-id] [create] [m-vpls]
description description-string
sap sap-id create
managed-vlan-list
range vlan-range
stp
no shutdown
```

The following example shows output for a configured management VPLS.

```
*A:ALA-1>config>service# info
-----
      vpls 2000 customer 6 m-vpls create
      stp
        no shutdown
      exit
      sap 1/1/1:100 create
      exit
      sap 1/1/2:200 create
      exit
      sap 1/1/3:300 create
      managed-vlan-list
        range 1-50
      exit
      no shutdown
    exit
-----
*A:ALA-1>config>service#
```

5.6.4.2 Creating a management VPLS for spoke-SDP protection



Note:

SDPs are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode. For 7210 SAS platforms operating in access-uplink mode, management VPLS can be used for protection of QinQ uplinks. Refer to the following example for more information.

This section provides a brief overview of the tasks that must be performed to configure a management VPLS for spoke-SDP protection and provides the CLI commands, see [Figure 77: Example configuration for protected VPLS spoke-SDP](#). The following tasks should be performed on all four nodes providing the protected VPLS service.

Before configuring a management VPLS, please first read [Configuring a VPLS SAP](#) for an introduction to the concept of management VPLS and spoke-SDP redundancy:


1. Create an SDP to the local peer node (node ALA-A2 in [Figure 77: Example configuration for protected VPLS spoke-SDP](#)).
2. Create an SDP to the remote peer node (node ALA-B1 in [Figure 77: Example configuration for protected VPLS spoke-SDP](#)).
3. Create a management VPLS.
4. Create a spoke-SDP in the m-vpls using the SDP defined in Step 1. Ensure that this spoke-SDP runs over a protected LSP.
5. Enable the management VPLS service and verify that it is operationally up.
6. Create a spoke-SDP in the m-vpls using the SDP defined in Step 2.

Optionally, modify STP parameters for load balancing.

7. Create one or more user VPLS services with spoke-SDPs on the tunnel SDP defined by Step 2.

As long as the user spoke-SDPs created in step 7 are in this same tunnel SDP with the management spoke-SDP created in step 6, the management VPLS will protect them.

Figure 77: Example configuration for protected VPLS spoke-SDP



Use the following syntax to create a management VPLS for spoke-SDP protection.

```
config>service# sdp sdp-id mpls create
far-end ip-address
lsp lsp-name
no shutdown
```

```
vpls service-id customer customer-id [m-vpls] create
description description-string
spoke-sdp sdp-id:vc-id create
stp
no shutdown
```

Example: VPLS configuration output

```
*A:ALA-A1>config>service# info
-----
...
sdp 100 mpls create
far-end 10.0.0.30
lsp "toALA-B1"
no shutdown
exit
sdp 300 mpls create
far-end 10.0.0.20
lsp "toALA-A2"
no shutdown
exit
vpls 101 customer 1 m-vpls create
spoke-sdp 100:1 create
exit
```

```

    spoke-sdp 300:1 create
    exit
    stp
    exit
    no shutdown
    exit
...
-----
*A:ALA-A1>config>service#


```

5.6.4.3 Configuring load Balancing with management VPLS

With the concept of management VPLS, it is possible to load balance the user VPLS services across the two protecting nodes. This is done by creating two management VPLS instances, where both instances have different active QinQ spokes (by changing the STP path-cost). When different user VPLS services are associated with either the two management VPLS services, the traffic will be split across the two QinQ spokes. Load balancing can be achieved in SAP protection scenarios.

The following figure shows an example of a configuration for load balancing with management VPLS.

Figure 78: Example configuration for load balancing with management VPLS



Note: the STP path costs in each peer node should be reversed.

```

config>service# vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-
star | dot1q | dot1q-preserve}] [customer-vid vlan-id]
description description-string
sap sap-id create
managed-vlan-list
range vlan-range

```

```
stp
no shutdown
```

Example: VPLS configuration output

```
*A:ALA-1>config>service# info
-----
vpls 100 customer 1 m-vpls svc-sap-type dot1q create
  stp
    no shutdown
  exit
  sap 1/1/2:100.* create
    managed-vlan-list
      range 1-10
    exit
  stp
    path-cost 1
  exit
exit
sap 1/1/3:500.* create
  shutdown
  managed-vlan-list
    range 1-10
  exit
exit
no shutdown
exit
vpls 200 customer 6 m-vpls svc-sap-type dot1q create
  stp
    no shutdown
  exit
  sap 1/1/2:1000.* create
    managed-vlan-list
      range 110-200
    exit
  exit
  sap 1/1/3:2000.* create
    managed-vlan-list
      range 110-200
    exit
  stp
    path-cost 1
  exit
exit
no shutdown
exit
vpls 101 customer 1 svc-sap-type dot1q create
  stp
    shutdown
  exit
  sap 1/1/1:100 create
  exit
  sap 1/1/2:1.* create
  exit
  sap 1/1/3:1.* create
  exit
  no shutdown
exit
vpls 201 customer 1 svc-sap-type dot1q create
  stp
    shutdown
  exit
  sap 1/1/1:200 create
```

```

        exit
        sap 1/1/2:110.* create
        exit
        sap 1/1/3:110.* create
        exit
        no shutdown
    exit
-----
*A:ALA-1>config>service#

```

5.6.4.4 Configuring a BGP-auto-discovery

```
config>service# sdp-template sdp-template-id
```

```
config>service# l2-auto-bind policy-id [use-provisioned-sdp]
```

BGP-AD automatically creates SDP-bindings using a template to configure SDP-binding configuration parameters. Layer 2-auto-bind is a command used to initiate a template that is used by BGP-AD for PW instantiation under related VPLS instances.

The template may be referenced in the "service vpls bgp-ad" object and used subsequently to instantiate PWs to a remote PE and VSI instance advertised through BGP Auto-Discovery. Changes to these dynamically created objects cannot be performed directly through CLI or SNMP. There are two possible methods to initiate the change:

- Configure a new "Layer 2-auto-bind" association under **service>vpls>bgp-ad**. This method is used when the existing policy is used by multiple VPLS services and only one or a few require the change.
- Change the parameters of the current template. This method is used when a change in parameter is required for the majority of VPLS services that use the template.

Changes are not automatically propagated to the instantiated objects and must be done through one of two tool commands:

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

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

This command forces evaluation of changes that were made in the Layer 2-auto-bind template indicated in the command. This command can be applied to an individual VPLS service or all VPLS services that reference the template if no service is specified.

The parameters are divided into three classes:

- class 1 - modified at create time only
- class 2 - modified only when the object is administratively shutdown
- class 3 - no restrictions

Parameters that fall into class 1 will destroy existing objects and recreate objects with the new values. Parameters in class 2 will momentarily shutdown the object, change the parameter, then re-enable the object. Class 3 can be changed without affecting the operational status of the objects of service.

For the Layer 2-auto-bind template, the parameters are treated as follows:

- class 1 - adding or removing a split-horizon-group, switching between a manual and auto SDP

- class 2 - changing the **vc-type** {ether | vlan}
- class 3 - all other changes

The keyword `allow-service-impact` enables service impacting changes. If this keyword is not configured, an error message is generated if the parameter changes are service impacting.

5.6.4.5 Configuring load balancing with management VPLS

With the concept of management VPLS, it is possible to load balance the user VPLS services across the two protecting nodes. This is done by creating two management VPLS instances, where both instances have different active spokes (by changing the STP path-cost). When different user VPLS services are associated with either the two management VPLS services, the traffic will be split across the two spokes.

Load balancing can be achieved in both the SAP protection and spoke-SDP protection scenarios. The following figure shows an example with the following configuration.

- Dut-C - Spoke-SDP

```
mvpls 100
MVPLS M1
Dut-A - Spoke SDP 1201:100 (STP blocked); 1401:100
Dut-B - Spoke SDP 1201:100; 2301:100
```

- Dut-C - Spoke-SDP 1401:100; 2301:100

```
uvpls 101
UVPLS U1
Dut-A - Spoke SDP 1201:101; 1401:101
Dut-B - Spoke SDP 1201:101; 2301:101
```


- Dut-C - Spoke-SDP 1401:101; 2301:101

```
mvpls 200
MVPLSM2
Dut-A - Spoke SDP 1202:200; 1402:200 (STP blocked)
Dut-B - Spoke SDP 1202:200; 2302:200
```

- Dut-C - Spoke-SDP 1402:200; 2302:200

```
uvpls 201
UVPLS U2
Dut-A - Spoke SDP 1202:201; 1402:201
Dut-B - Spoke SDP 1202:201; 2302:201
Dut-C - Spoke SDP 1402:201; 2302:201
```

Figure 79: Example Configuration for Load Balancing Across Two Protected VPLS Spoke-SDPs



Use the following syntax to create a load balancing across two management VPLS instances.

```
config>service# sdp sdp-id mpls create
    far-end ip-address
    lsp lsp-name
    no shutdown
```

```
vpls service-id customer customer-id [m-vpls] create
    description description-string
    spoke-sdp sdp-id:vc-id create
    stp
        path-cost
    stp
    no shutdown
```

This following output shows example configurations for load balancing across two protected VPLS spoke-SDPs:

Example: ALA-A configuration

The following is a sample configuration output on ALA-A.

```
# MVPLS 100 configs

*A:ALA-A# configure service vpls 100
*A:ALA-A>config>service>vpls# info
-----
    description "Default tls description for service id 100"
    stp
        no shutdown
    exit
    sap lag-3:100 create
        description "Default sap description for service id 100"
        managed-vlan-list
            range 101-110
        exit
    exit
    spoke-sdp 1201:100 create
        stp
            path-cost 100
        exit
    exit
    spoke-sdp 1401:100 create
    exit
```



```

no shutdown
-----
*A:ALA-A>config>service>vpls#

# UVPLS 101 configs
*A:ALA-A>config>service# vpls 101
*A:ALA-A>config>service>vpls# info
-----
description "Default tls description for service id 101"
sap lag-3:101 create
description "Default sap description for service id 101"
exit
spoke-sdp 1201:101 create
exit
spoke-sdp 1401:101 create
exit
no shutdown
-----
*A:ALA-A>config>service>vpls#

# MVPLS 200 configs
*A:ALA-A# configure service vpls 200
*A:ALA-A>config>service>vpls# info
-----
description "Default tls description for service id 200"
stp
no shutdown
exit
sap lag-3:200 create
description "Default sap description for service id 200"
managed-vlan-list
range 201-210
exit
exit
spoke-sdp 1202:200 create
exit
spoke-sdp 1402:200 create
stp
path-cost 100
exit
exit
no shutdown
-----
*A:ALA-A>config>service>vpls#

# UVPLS 201 configs
*A:ALA-A>config>service# vpls 201
*A:ALA-A>config>service>vpls# info
-----
description "Default tls description for service id 201"
sap lag-3:201 create
description "Default sap description for service id 201"
exit
spoke-sdp 1202:201 create
exit
spoke-sdp 1402:201 create
exit

```

```

no shutdown
-----
*A:ALA-A>config>service>vpls# exit all

```

Example: ALA-B configuration

The following is a sample configuration output on ALA-B (7210), the upper left node. It is configured such that it becomes the root bridge for MVPLS 100 and MVPLS 200.

```

# MVPLS 100 configs

*A:ALA-B# configure service vpls 100
*A:ALA-B>config>service>vpls# info
-----
description "Default tls description for service id 100"
stp
    priority 0
    no shutdown
exit
spoke-sdp 1201:100 create
exit
spoke-sdp 2301:100 create
exit
no shutdown
-----
*A:ALA-B>config>service>vpls#

# UVPLS 101 configs

*A:ALA-B>config>service# vpls 101
*A:ALA-B>config>service>vpls# info
-----
description "Default tls description for service id 101"
spoke-sdp 1201:101 create
exit
spoke-sdp 2301:101 create
exit
no shutdown
-----
*A:ALA-B>config>service>vpls#

# MVPLS 200 configs

*A:ALA-B# configure service vpls 200
*A:ALA-B>config>service>vpls# info
-----
description "Default tls description for service id 200"
stp
    priority 0
    no shutdown
exit
spoke-sdp 1202:200 create
exit
spoke-sdp 2302:200 create
exit
no shutdown
-----
*A:ALA-B>config>service>vpls#

```

```
# UVPLS 201 configs

*A:ALA-B>config>service# vpls 201
*A:ALA-B>config>service>vpls# info
-----

description "Default tls description for service id 201"
spoke-sdp 1202:201 create
exit
spoke-sdp 2302:201 create
exit
no shutdown
-----

*A:ALA-B>config>service>vpls#
```

Example: ALA-C configuration

The following is a sample configuration output on ALA-C (7210), the upper right node.

```
# MVPLS 100 configs

*A:ALA-C# configure service vpls 100
*A:ALA-C>config>service>vpls# info
-----

description "Default tls description for service id 100"
stp
  priority 4096
  no shutdown
exit
spoke-sdp 1401:100 create
exit
spoke-sdp 2301:100 create
exit
no shutdown
-----

*A:ALA-C>config>service>vpls#

# UVPLS 101 configs

*A:ALA-C>config>service# vpls 101
*A:ALA-C>config>service>vpls# info
-----

description "Default tls description for service id 101"
spoke-sdp 1401:101 create
exit
spoke-sdp 2301:101 create
exit
no shutdown
-----

*A:ALA-C>config>service>vpls#
```

```
# MVPLS 200 configs

*A:ALA-C# configure service vpls 200
*A:ALA-C>config>service>vpls# info
-----
```

```

        description "Default tls description for service id 200"
        stp
            priority 4096
            no shutdown
        exit
        spoke-sdp 1402:200 create
        exit
        spoke-sdp 2302:200 create
        exit
        no shutdown
-----
*A:ALA-C>config>service>vpls#

# UVPLS 201 configs

*A:ALA-C>config>service# vpls 201
*A:ALA-C>config>service>vpls# info
-----

        description "Default tls description for service id 201"
        spoke-sdp 1402:201 create
        exit
        spoke-sdp 2302:201 create
        exit
        no shutdown
-----
*A:ALA-C>config>service>vpls#

```

5.6.4.6 Configuring selective MAC flush

Use the following syntax to enable selective MAC Flush in a VPLS.

```

config>service# vpls service-id
send-flush-on-failure

```

Use the following syntax to disable selective MAC Flush in a VPLS.

```

config>service# vpls service-id
no send-flush-on-failure


```

5.6.4.7 Configuring load balancing with management VPLS

With the concept of management VPLS, it is possible to load balance the user VPLS services across the two protecting nodes. This is done by creating two management VPLS instances, where both instances have different active QinQ spokes (by changing the STP path-cost). When different user VPLS services are associated with either the two management VPLS services, the traffic will be split across the two QinQ spokes. Load balancing can be achieved in SAP protection scenarios.

The following figure shows an example of a configuration for load balancing with management VPLS.

Figure 80: Example configuration for load balancing with management VPLS

**Note:**

The STP path costs in each peer node should be reversed.

```
config>service# vpls service-id [customer customer-id] [create][m-vpls] [svc-sap-type {null-
star | any | dot1q-preserve}] [customer-vid vlan-id]
description description-string
sap sap-id create
managed-vlan-list
range vlan-range
stp
no shutdown
```

Example

The following is a sample VPLS configuration output.

```
*A:ALA-1>config>service# info
-----
vpls 100 customer 1 m-vpls svc-sap-type any create
stp
no shutdown
exit
sap 1/1/2:100.* create
managed-vlan-list
range 1-10
exit
stp
path-cost 1
exit
exit
sap 1/1/3:500.* create
```

```

        shutdown
        managed-vlan-list
            range 1-10
        exit
    exit
    no shutdown
exit
vpls 200 customer 6 m-vpls svc-sap-type any create
    stp
        no shutdown
    exit
    sap 1/1/2:1000.* create
        managed-vlan-list
            range 110-200
        exit
    exit
    sap 1/1/3:2000.* create
        managed-vlan-list
            range 110-200
        exit
    stp
        path-cost 1
    exit
    exit
    no shutdown
exit
vpls 101 customer 1 svc-sap-type any create
    stp
        shutdown
    exit
    sap 1/1/1:100 create
    exit
    sap 1/1/2:1.* create
    exit
    sap 1/1/3:1.* create
    exit
    no shutdown
exit
vpls 201 customer 1 svc-sap-type any create
    stp
        shutdown
    exit
    sap 1/1/1:200 create
    exit
    sap 1/1/2:110.* create
    exit
    sap 1/1/3:110.* create
    exit
    no shutdown
exit
-----
*A:ALA-1>config>service#

```

5.6.5 Configuring IGMPv3 snooping in RVPLS

IGMPv3 snooping in RVPLS is supported only for IES (not for VPRNs).

Use the following syntax to configure IGMPv3 snooping in routed VPLS bound to an IES.

```

config>service# vpls service-id customer customer-id [svc-sap-type {any}] [b-vpls | i-vpls | r-
vpls] create

```

```

config>service>vpls# service-name service-name
config>service>vpls# allow-ip-int-bind
config>service>vpls>allow-ip-int-bind# exit
config>service>vpls# igmp-snooping
config>service>vpls>igmp-snooping# no shutdown
config>service>vpls# exit
config>service>vpls# sap sap-id create
config>service>vpls>sap# igmp-snooping
config>service>vpls>sap>igmp-snooping# mrouter-port
config>service>vpls>sap>igmp-snooping# exit
config>service>vpls>sap># exit
config>service>vpls># exit
config>service# ies service-id customer customer-id create
config>service>ies# interface ip-int-name create
config>service>ies>interface# address ip-address/mask
config>service>ies>interface# vpls service-name

```

Example

The following is a sample RVPLS configuration output that uses IGMPv3 snooping.

```

#-----
echo "Port Configuration"
#-----

...snip...

port 1/1/5
  ethernet
    mode hybrid
    access
    exit
    encap-type dot1q
    multicast-ingress ip-mc
  exit
  no shutdown
exit

#-----

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

service
  customer 1 create
    description "Default customer"
  exit
  ies 6 customer 1 create
    interface "IGMP-test" create
    exit
  exit

...snip

vpls 3 customer 1 r-vpls svc-sap-type any create
  allow-ip-int-bind
  exit
  stp
    shutdown
  exit
  igmp-snooping
    no shutdown
  exit

```

```

service-name "GS-IGMP-Snooping"
sap 1/1/5:333 create
  igmp-snooping
    mrouter-port
  exit
  ingress
  exit
  egress
  exit
exit
....snip

```

```

ies 6 customer 1 create
  interface "IGMP-test" create
    address 192.168.x.x/24
    vpls "GS-IGMP-Snooping"
  exit
  exit
  no shutdown
  exit
exit
#-----

```


5.6.6 Configuring BGP Auto-Discovery

This section provides important information to describe the different configuration options used to populate the required BGP AD and generate the LDP generalized pseudowire-ID FEC fields. There are a large number of configuration options that are available with this feature. Not all these configurations option are required to start using BGP AD. At the end of this section, it will be apparent that a very simple configuration will automatically generate the required values used by BGP and LDP. In most cases, deployments will provide full mesh connectivity between all nodes across a VPLS instance. However, capabilities are available to influence the topology and build hierarchies or hub and spoke models.

5.6.6.1 Configuration steps

Using the following figure, assume PE6 was previously configured with VPLS 100 as indicated by the configurations lines in the upper right. The BGP AD process will commence after PE134 is configured with the VPLS 100 instance as shown in the upper left. This shows a very basic and simple BGP AD configuration. The minimum requirement for enabling BGP AD on a VPLS instance is configuring the VPLS-ID and point to a pseudowire template.

Figure 81: BGP AD configuration example



In many cases, VPLS connectivity is based on a pseudowire mesh. To reduce the configuration requirement, the BGP values can be automatically generated using the VPLS-ID and the MPLS router-ID. By default, the lower six bytes of the VPLS-ID are used to generate the RD and the RT values. The VSI-ID value is generated from the MPLS router-ID. All of these parameters are configurable and can be coded to suit requirements and build different topologies

A helpful command displays the service information, the BGP parameters and the SDP bindings in use. When the discovery process is completed successfully each endpoint will have an entry for the service.

```
PE134># show service l2-route-table
```

When only one of the endpoints has an entry for the service in the Layer 2-routing-table, it is most likely a problem with the RT values used for import and export. This would most likely happen when different import and export RT values are configured using a router policy or the route-target command.


Service-specific commands continue to be available to display service-specific information, including status.

```
PERs6# show service sdp-using
```

BGP AD advertises the VPLS-ID in the extended community attribute, VSI-ID in the NLRI and the local PE ID in the BGP next hop. At the receiving PE, the VPLS-ID is compared against locally provisioned information to determine whether the two PEs share a common VPLS. If it is found that they do, the BGP information is used in the signaling phase.

5.6.7 Configuring AS pseudowire in VPLS

Figure 82: Sample topology-AS pseudowire in VPLS



In the preceding figure, pseudowire is configured on MTU.

Example

The following is a sample configuration output on the MTU.

```

*A:MTU>config>service>vpls>endpoint# back
*A:MTU>config>service>vpls# info
-----
    send-flush-on-failure
    stp
        shutdown
    exit
    endpoint "vpls1" create
        description "vpls1_endpoint"
        revert-time 60
        ignore-standby-signaling
        no suppress-standby-signaling
        block-on-mesh-failure
    exit
    sap 1/1/3 create
    exit
    spoke-sdp 301:1 endpoint "vpls1" create
        stp
            shutdown
        exit
        block-on-mesh-failure
    exit
    spoke-sdp 302:1 endpoint "vpls1" create
        stp
            shutdown
        exit
        block-on-mesh-failure
    exit
    no shutdown
-----
*A:MTU>config>service>vpls#
  
```

5.7 Service management tasks

This section describes the service management tasks.

5.7.1 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 parameter such as description SAP and then enter the new information.

Example

The following is a sample modified VPLS configuration output.

```
*A:ALA-1>config>service>vpls# info
-----
description "This is a different description."
disable-learning
disable-aging
discard-unknown
local-age 500

stp
  shutdown
exit
sap 1/1/5:22 create
  description "VPLS SAP"
exit

exit
no shutdown
-----
*A:ALA-1>config>service>vpls#
```

5.7.2 Modifying management VPLS parameters

To modify the range of VLANs on an access port that is to be managed by an existing management VPLS, enter the new range first and then remove the old range. 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.

```
config>service# vpls service-id
  sap sap-id
    managed-vlan-list
      [no] range vlan-range
```

5.7.3 Deleting a management VPLS

As with normal VPLS service, a management VPLS cannot be deleted until SAPs and SDPs are unbound (deleted), interfaces are shutdown, and the service is shutdown on the service level.

Use the following syntax to delete a management VPLS service.

```
config>service
[no] vpls service-id
shutdown
[no] spoke-sdp sdp-id
[no] sap sap-id
shutdown
```

5.7.4 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 needed, first un-manage the user VPLS service by removing them from the managed-vlan-list or moving the spoke-SDPs on to another tunnel SDP.

```
config>service
vpls service-id
shutdown
```

Example:

```
config>service# vpls 1
config>service>vpls# shutdown
config>service>vpls# exit
```

5.7.5 Deleting a VPLS service

A VPLS service cannot be deleted until SAPs and SDPs (not applicable for 7210 SAS-T configured in access-uplink mode) are unbound (deleted), interfaces are shutdown, and the service is shutdown on the service level.

Use the following syntax to delete a VPLS service.

```
config>service
[no] vpls service-id
shutdown
[no] spoke-sdp sdp-id
shutdown
sap sap-id
no sap sap-id
shutdown
```

5.7.6 Disabling a VPLS service

Use the following syntax to shut down a VPLS service without deleting the service parameters.

```
config>service> vpls service-id
[no] shutdown
```

Example:

```
config>service# vpls 1
```

```
config>service>vpls# shutdown
config>service>vpls# exit
```

5.7.7 Re-enabling a VPLS service

Use the following syntax to re-enable a VPLS service that was shut down.

```
config>service> vpls service-id
[no] shutdown
```

Example:

```
config>service# vpls 1
config>service>vpls# no shutdown
config>service>vpls# exit
```

5.8 VPLS services command reference

5.8.1 Command hierarchies

- [VPLS service configuration commands](#)
 - [Global commands](#)
 - [VPLS service xSTP commands](#)
 - [VPLS service SAP DHCP snooping commands](#)
 - [VPLS DHCPv6 snooping commands for SAP and SDP bindings](#)
 - [SAP commands](#)
 - [VPLS SAP QoS and filter commands \(for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE \(standalone and standalone-VC\), and 7210 SAS-Sx 10/100GE\)](#)
 - [VPLS SAP QoS and filter Commands \(for 7210 SAS-Mxp\)](#)
 - [VPLS service and SAP IGMP snooping and MVR commands](#)
 - [VPLS SAP meter override commands](#)
 - [VPLS service SAP xSTP commands](#)
 - [VPLS SAP statistics commands](#)
 - [Mesh SDP commands](#)
 - [Spoke-SDP commands](#)
 - [Routed VPLS commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

5.8.1.1 VPLS service configuration commands

5.8.1.1.1 Global commands

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
- no vpls service-id
- bgp
- pw-template-binding policy-id [split-horizon-group group-name] [import-rt
{ext-community... (up to 5 max)}}
- no pw-template-binding policy-id
- route-distinguisher rd
- no route-distinguisher
- route-target {ext-community | {[export ext-community] [import ext-
community]}}
- no route-target
- vsi-export policy-name [policy-name... (up to 5 max)]
- no vsi-export
- vsi-import policy-name [policy-name... (up to 5 max)]
- no vsi-import
- [no] bgp-ad
- [no] shutdown
- vpls-id vpls-id
- vsi-id
- prefix low-order-vsi-id
- no prefix
- description description-string
- no description
- [no] disable-aging
- [no] disable-learning
- [no] discard-unknown
- endpoint endpoint-name [create]
- no endpoint
- block-on-mesh-failure
- [no] block-on-mesh-failure
- description description-string
- no description
- [no] ignore-standby-signaling
- [no] mac-pinning
- max-nbr-mac-addr table-size
- no max-nbr-mac-addr
- revert-time revert-time | infinite
- no revert-time
- static-mac ieee-address [create]
- no static-mac
- [no] suppress-standby-signaling
- [no] fdb-table-high-wmark high-water-mark
- [no] fdb-table-low-wmark low-water-mark
- fdb-table-size table-size
- no fdb-table-size [table-size]
- local-age aging-timer
- no local-age
- [no] mac-move
- move-frequency frequency
- no move-frequency
- retry-timeout timeout

```

```

- no retry-timeout
- [no] shutdown
- [no] propagate-mac-flush
- remote-age aging-timer
- no remote-age
- [no] send-flush-on-failure
- service-mtu octets
- no service-mtu
- service-mtu-check octets
- no service-mtu-check
- no service-name
- [no] shutdown
- split-horizon-group group-name [create]
  - description description-string
  - no description

```

5.8.1.1.2 VPLS service xSTP commands

```

config
- service
  - vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
  - vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
  - no vpls service-id
  - stp
    - forward-delay forward-delay
    - no forward-delay
    - hello-time hello-time
    - no hello-time
    - hold-count BDPU tx hold count
    - no hold-count
    - max-age max-age
    - no max-age
    - mode {rstp | comp-dot1w | dot1w | mstp | pmstp}
    - no mode
    - [no] mst-instance mst-inst-number
      - mst-port-priority bridge-priority
      - no mst-port-priority
      - [no] vlan-range vlan-range
    - mst-max-hops hops-count
    - no mst-max-hops
    - mst-name region-name
    - no mst-name
    - mst-revision revision-number
    - no mst-revision
    - priority bridge-priority
    - no priority
    - [no] shutdown

```

5.8.1.1.3 VPLS service SAP DHCP snooping commands

```

config
- service
  - vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
  - vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]

```

```

- no vpls service-id
- sap sap-id [split-horizon-group group-name] [g8032-shg-enable] [eth-ring ring-
index] [create]
- no sap sap-id
- dhcp
- description description-string
- no description
- [no] option
- action {replace | drop | keep}
- no action
- [no] circuit-id [ascii-tuple | vlan-ascii-tuple]
- [no] remote-id [mac | string string]
- [no] vendor-specific-option
- [no] client-mac-address
- [no] sap-id
- [no] service-id
- string text
- no string
- [no] system-id
- [no] shutdown
- [no] snoop

```

5.8.1.1.4 VPLS DHCPv6 snooping commands for SAP and SDP bindings

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
- no vpls service-id
- mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
- no mesh-sdp sdp-id[:vc-id]
- dhcp6
- [no] description description-string
- [no] shutdown
- [no] snoop [network-facing]
- sap sap-id [split-horizon-group group-name] [g8032-shg-enable] [eth-ring ring-
index] [create]
- no sap sap-id
- dhcp6
- [no] description description-string
- [no] option
- interface-id
- interface-id ascii-tuple
- interface-id sap-id
- interface-id string string
- no interface-id
- [no] remote-id [mac | string string]
- [no] shutdown
- [no] snoop [client-facing | network-facing | both]
- [no] trusted
- spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] [split-horizon-group group-name]
[use-evpn-default-shg]
- no spoke-sdp sdp-id[:vc-id]
- dhcp6
- [no] description description-string
- [no] shutdown
- [no] snoop [network-facing]

```


5.8.1.1.5 SAP commands

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
- no vpls service-id
- sap sap-id [split-horizon-group group-name] [g8032-shg-enable] [eth-ring ring-
index] [create]
- no sap sap-id
- accounting-policy acct-policy-id
- no accounting-policy
- bpdu-translation {auto | pvst | stp}
- no bpdu-translation
- [no] cflowd
- [no] collect-stats
- description description-string
- no description
- [no] disable-aging
- [no] disable-learning
- [no] discard-unknown-source
- eth-cfm
- mep mep-id domain md-index association ma-index [direction {up | down}]
primary-vlan-enable
- no mep mep-id domain md-index association ma-index
- [no] ais-enable
- client-meg-level [level [level...]]
- no client-meg-level
- [no] description
- interval {1 | 60}
- no interval
- priority priority-value
- no priority
- [no] 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
- mip [mac mac address]
- mip default-mac
- no mip
- l2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp] [lldp]
- no l2pt-termination
- limit-mac-move [blockable | non-blockable]
- no limit-mac-move
- [no] mac-pinning
- max-nbr-mac-addr table-size
- no max-nbr-mac-addr
- managed-vlan-list
- default-sap

```

```

- no default-sap
- no range vlan-range
- range vlan-range
- [no] shutdown

```

5.8.1.1.6 VPLS SAP QoS and filter commands (for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), and 7210 SAS-Sx 10/100GE)

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
- no vpls service-id
- sap sap-id [create] [g8032-shg-enable] [eth-ring ring-index] [split-horizon-
group group-name]
- sap sap-id [g8032-shg-enable] [eth-ring ring-index] [create]
- no sap sap-id
- egress
- aggregate-meter-rate rate-in-kbps [burst burst-in-kbits] [enable-stats]
- no aggregate-meter-rate
- filter ip ip-filter-id
- filter ipv6 ipv6 -filter-id
- filter mac mac-filter-id
- no filter [ip ip-filter-id] [ipv6 ipv6 -filter-id] [mac mac-filter-id]
- ingress
- aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
- no aggregate-meter-rate
- filter ip ip-filter-id
- filter [ipv6 ipv6-filter-id]
- filter mac mac-filter-id
- no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
- qos policy-id
- no qos

```

5.8.1.1.7 VPLS SAP QoS and filter Commands (for 7210 SAS-Mxp)

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- sap sap-id [create] [g8032-shg-enable] [eth-ring ring-index] [split-horizon-
group group-name]
- sap sap-id [g8032-shg-enable] [eth-ring ring-index] [create]
- no sap sap-id
- egress
- agg-rate-limit [cir cir-rate] [pir pir-rate]
- no agg-rate-limit
- filter ip ip-filter-id
- filter ipv6 ipv6 -filter-id
- filter mac mac-filter-id
- no filter [ip ip-filter-id] [ipv6 ipv6 -filter-id] [mac mac-filter-id]
- qos policy-id
- no qos
- ingress
- aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]

```

```

- no aggregate-meter-rate
- filter ip ip-filter-id
- filter [ipv6 ipv6-filter-id]
- filter mac mac-filter-id
- no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
- qos policy-id [enable-table-classification]
- no qos policy-id

```

5.8.1.1.8 VPLS service and SAP IGMP snooping and MVR commands

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- igmp-snooping
- mvr
- description description-string
- no description
- group-policy policy-name
- no group-policy
- no shutdown
- shutdown
- query-interval interval
- no query-interval
- no query-src-ip
- query-src-ip ip-address
- no report-src-ip
- report-src-ip ip-address
- robust-count count
- no robust-count
- no shutdown
- shutdown
- sap sap-id [split-horizon-group group-name] [g8032-shg-enable] [eth-ring ring-
index] [create]
- no sap sap-id
- igmp-snooping
- [no] disable-router-alert-check
- [no] fast-leave
- import policy-name
- no import
- last-member-query-interval interval
- no last-member-query-interval
- max-num-groups max-num-groups
- no max-num-groups
- max-num-sources max-num-sources
- no max-num-sources
- [no] mrouter-port
- mvr
- from-vpls service-id
- no from-vpls
- to-sap sap-id
- no to-sap
- query-interval seconds
- no query-interval
- query-response-interval interval
- no query-response-interval
- robust-count count
- no robust-count
- [no] send-queries
- static
- [no] group group-address

```

```

- [no] source ip-address
- [no] starg
- version version
- no version
- mfib-table-high-wmark high-water-mark
- no mfib-table-high-wmark
- mfib-table-low-wmark low-water-mark
- no mfib-table-low-wmark
- mfib-table-size table-size
- no mfib-table-size

```

5.8.1.1.9 VPLS SAP meter override commands

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
- no vpls service-id
- sap sap-id [split-horizon-group group-name] [g8032-shg-enable] [eth-ring ring-
index] [create]
- no sap sap-id
- ingress
- meter-override
- meter meter-id [create]
- no meter meter-id
- adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
- cbs size [kbits | bytes | kbytes]
- no cbs
- mbs size [kbits | bytes | kbytes]
- no mbs
- mode mode
- no mode
- rate cir cir-rate [pir pir-rate]

```

5.8.1.1.10 VPLS service SAP xSTP commands

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
- no vpls service-id
- sap sap-id [split-horizon-group group-name] [g8032-shg-enable] [eth-ring ring-
index] [create]
- no sap sap-id
- stp
- [no] auto-edge
- [no] edge-port
- link-type {pt-pt | shared}
- no link-type {pt-pt | shared}
- mst-instance mst-inst-number
- mst-path-cost inst-path-cost
- no mst-path-cost
- mst-port-priority stp-priority
- no mst-port-priority

```

```

- path-cost sap-path-cost
- no path-cost
- [no] port-num virtual-port-number
- priority stp-priority
- no priority
- root-guard
- no root-guard
- [no] shutdown
- tod-suite tod-suite-name
- no tod-suite

```

5.8.1.1.11 VPLS SAP statistics commands

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
- no vpls service-id
- sap sap-id [split-horizon-group group-name] [g8032-shg-enable] [eth-ring ring-
index] [create]
- no sap sap-id
- statistics
- ingress
- counter-mode {in-out-profile-count | forward-drop-count}
- [no] drop-count-extra-vlan-tag-pkts

```

5.8.1.1.12 Mesh SDP commands



Note:

Mesh SDP commands are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- no vpls service-id
- mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
- no mesh-sdp sdp-id[:vc-id]
- accounting-policy acct-policy-id
- no accounting-policy
- [no] collect-stats
- [no] control-word
- description description-string
- no description
- egress
- no vc-label [egress-vc-label]
- eth-cfm
- mep mep-id domain md-index association ma-index [direction {up} {down}]
- no mep mep-id domain md-index association ma-index
- [no] ais-enable
- client-meg-level [[level [level...]]
- no client-meg-level
- interval {1 | 60}
- no interval

```

```

- priority priority-value
- no priority
- [no] ccm-enable
- ccm-ltm-priority priority
- no ccm-ltm-priority
- [no] description description-string
- [no] eth-test-enable
- bit-error-threshold bit-errors
- test-pattern {all-zeros | all-ones} [crc-enable]
- no test-pattern
- 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-vlan-vc-forwarding
- hash-label
- hash-label [signal-capability]
- no hash-label
- igmp-snooping
- [no] disable-router-alert-check
- import policy-name
- no import
- last-member-query-interval interval
- no last-member-query-interval
- max-num-groups max-num-groups
- no max-num-groups
- [no] mrouter-port
- query-interval interval
- no query-interval
- query-response-interval interval
- no query-response-interval
- robust-count count
- no robust-count
- [no] send-queries
- static
- [no] group grp-ip-address
- [no] source
- version version
- no version
- ingress
- vc-label egress-vc-label
- [no] mac-pinning
- [no] static-mac ieee-address
- [no] shutdown
- statistics
- ingress
- [no] drop-count-extra-vlan-tag-pkts
- vlan-vc-tag 0..4094
- no vlan-vc-tag [0..4094]

```

5.8.1.1.13 Spoke-SDP commands



Note:

Spoke-SDP commands are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

```

config
- service

```

```

- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- no vpls service-id
- spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] [split-horizon-
group group-name] [use-evpn-default-shg]
- no spoke-sdp sdp-id[:vc-id]
- accounting-policy acct-policy-id
- no accounting-policy
- [no] block-on-mesh-failure
- bpdu-translation {auto | pvst | stp}
- no bpdu-translation
- [no] collect-stats
- [no] control-word
- description description-string
- no description
- [no] disable-aging
- [no] disable-learning
- [no] discard-unknown-source
- eth-cfm
- mep mep-id domain md-index association ma-index [direction {up} {down}]
- no mep mep-id domain md-index association ma-index[no] ais-enable
- client-meg-level [[level [level...]]
- no client-meg-level
- interval {1 | 60}
- no interval
- priority priority-value
- no priority
- [no] ccm-enable
- ccm-ltm-priority priority
- no ccm-ltm-priority
- [no] description description string[
- no] eth-test-enable
- bit-error-threshold bit-errors
- test-pattern {all-zeros | all-ones} [crc-enable]
- no test-patternlow-priority-defect {allDef | macRemErrXcon | remErr
Xcon | errXcon | xcon | noXcon}
- mac-address mac-address
- no mac-addressone-way-delay-threshold seconds
- [no] shutdown
- mip [mac mac address]
- mip default-mac
- no mip
- egress
- vc-label egress-vc-label
- no vc-label [egress-vc-label]
- [no] force-vlan-vc-forwarding
- hash-label [signal-capability]
- no hash-label
- igmp-snooping
- [no] disable-router-alert-check
- import policy-name
- no import
- last-member-query-interval interval
- no last-member-query-interval
- max-num-groups max-num-groups
- no max-num-groups
- [no] mrouter-port
- query-interval interval
- no query-interval
- query-response-interval interval
- no query-response-interval
- robust-count count
- no robust-count
- [no] send-queries

```

```

- static
  - [no] group group-address
  - [no] source
- version version
- no version
- [no] ignore-standby-signaling
- ingress
  - vc-label egress-vc-label
  - no vc-label [egress-vc-label]
- l2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp] [lldp]
- no l2pt-termination
- no limit-mac-move [blockable | non-blockable]
- no limit-mac-move
- [no] mac-pinning
- max-nbr-mac-addr table-size
- no max-nbr-mac-addr
- precedence precedence-value | primary
- no precedence
- [no] shutdown
- [no] static-mac ieee-address
- statistics
  - ingress
    - [no] drop-count-extra-vlan-tag-pkts
- stp
  - [no] auto-edge
  - [no] edge-port
  - link-type {pt-pt | shared}
  - no link-type [pt-pt | shared]
  - path-cost sap-path-cost
  - no path-cost
  - [no] port-num virtual-port-number
  - priority stp-priority
  - no priority
  - no root-guard
  - root-guard
  - [no] shutdown
- vlan-vc-tag 0..4094
- no vlan-vc-tag [0..4094]

```

5.8.1.1.14 Routed VPLS commands

```

config
- service
  - vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
  - vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
  - no vpls service-id
  - [no] allow-ip-int-bind

```

5.8.1.2 Show commands



Note:

SDP commands are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

```
show
```



```

- service
- egress-label egress-label1 [egress-label2]
- fdb-info
- fdb-mac ieee-address [expiry]
- id service-id
- all
- base [msap] [bfd]
- dhcp
- statistics [sap sap-id] [interface interface-name]
- summary [interface interface-name | saps]
- dhcp6
- statistics [interface interface-name]
- statistics sap sap-id
- statistics sdp sdp-id:vc-id
- endpoint [endpoint-name]
- fdb [sap sap-id] [expiry] | [mac ieee-address [expiry]] | [detail] [expiry]
- igmp-snooping
- all
- base
- mvr
- mroute [detail]
- port-db sap sap-id [detail]
- port-db sap sap-id group grp-address
- port-db sdp sdp-id:vc-id [detail]
- port-db sdp sdp-id:vc-id group grp-address
- proxy-db [detail]
- proxy-db [group grp-ip-address]
- querier
- static [sap sap-id]
- statistics [sap sap-id | sdp sdp-id:vc-id]
- labels
- l2pt disabled
- l2pt [detail]
- mac-move
- mfib [brief ]
- mfib [group grp-address | mstp-configuration]
- sap [sap-id [detail]]
- sdp [sdp-id | far-end ip-addr] [detail]
- split-horizon-group [group-name]
- stp [detail]
- ingress-label start-label [end-label]
- sap-using [sap sap-id]
- sap-using [ingress | egress] filter filter-id
- sap-using [ingress | egress] qos-policy qos-policy-id
- sap-using [ingress | egress]
- sdp [sdp-id | far-end ip-address] [detail | keep-alive-history]
- sdp-using [sdp-id[:vc-id] | far-end ip-address]
- service-using [vpls]

```

5.8.1.3 Clear commands



Note:

SDP commands are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

```

clear
- service
- id service-id
- fdb {all | mac ieee-address | sap sap-id | mesh-sdp sdp-id[:vc-id] | spoke-
sdp sdp-id:vc-id}

```

```

- igmp-snooping
  - port-db sap sap-id [group grp-address]
  - querier
- statistics [all | sap sap-id | sdp sdp-id:vc-id]
- mesh-sdp sdp-id[:vc-id] ingress-vc-label
  - spoke-sdp sdp-id:vc-id ingress-vc-label
- spoke-sdp sdp-id:vc-id
- stp
  - detected-protocols [all | sap sap-id]
- statistics
  - id service-id
    - counters
    - mesh-sdp [sdp-id:vc-id] {all | counters | stp}
    - spoke-sdp sdp-id[:vc-id] {all | counters | stp | l2pt}
    - stp
  - sap sap-id {all | counters | stp}

```

5.8.1.4 Debug commands

```

debug
- service
  - id service-id

```

5.8.2 Command descriptions

- [VPLS configuration commands](#)
- [VPLS show commands](#)
- [VPLS clear commands](#)
- [VPLS debug commands](#)

5.8.2.1 VPLS configuration commands

- [Generic commands](#)
- [VPLS service commands](#)
- [VPLS STP commands](#)
- [VPLS SAP commands](#)
- [VPLS service SAP DHCP snooping commands](#)
- [VPLS DHCPv6 snooping commands for SAP and SDP bindings](#)
- [ETH-CFM service commands](#)
- [VPLS filter and QoS policy commands](#)
- [VPLS SDP commands](#)
- [SAP IGMP-snooping commands](#)
- [Routed VPLS commands](#)

5.8.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>service>vpls

config>service>vpls>split-horizon-group (not supported in access-uplink operating mode)

config>service>vpls>igmp-snooping>mvr

config>service>vpls>sap

config>service>vpls>sap>dhcp

config>service>vpls>sap>dhcp6 (supported on the 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone))

config>service>vpls>mesh-sdp>dhcp6 (supported on the 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone))

config>service>vpls>spoke-sdp>dhcp6 (supported on the 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone))

config>service>vpls>spoke-sdp (not supported in access-uplink operating mode)

config>service>pw-template>split-horizon-group (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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.

Parameters

string

Specifies the description character string. Allowed values are any string up to 80 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax

[no] shutdown

Context

config>service>vpls

config>service>vpls>snooping

config>service>vpls>igmp-snooping

config>service>vpls>sap

config>service>vpls>sap>dhcp

config>service>vpls>sap>dhcp6 (supported on the 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone))

config>service>vpls>mesh-sdp>dhcp6 (supported on the 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone))

config>service>vpls>spoke-sdp>dhcp6 (supported on the 7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone))

config>service>vpls>sap>stp

config>service>vpls>stp

config>service>vpls>spoke-sdp>stp (not supported in access-uplink operating mode)

config>service>vpls>bgp-ad (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

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

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

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

5.8.2.1.2 VPLS service commands

vpls

Syntax

vpls *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**m-vpls**] [**svc-sap-type** {**null-star** | **dot1q-preserve** | **any**}] [**b-vpls** | **i-vpls** | **r-vpls**] [**b-vid** *vid*]

no vpls *service-id*

Context

config>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates or edits a virtual private LAN services (VPLS) instance. The **vpls** command is used to create or maintain a VPLS service. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites, acting like a zero-hop Layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context. When a service is created, the **customer** keyword and *customer-id* parameter must be specified to associate the service with a customer. The *customer-id* value must already exist, having been created using the **customer** command in the service context. When 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.

When a service is created, the use of the **customer** *customer-id* command is optional for navigating into the service configuration context. Editing a service with the incorrect *customer-id* value specified results 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

Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values *service-id*: 1 to 2147483648

customer *customer-id*

Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

m-vpls

Specifies a management VPLS.

b-vpls

Specifies a PBB backbone-VPLS service, which can only be configured with SAPs. This keyword is supported only on 7210 SAS-T operating in network mode.

i-vpls

Specifies a PBB I-VPLS service, which can only be configured with SAPs. This keyword is only supported when the **svc-sap-type** value **any** is configured. This keyword is supported only on 7210 SAS-T operating in network mode.

create

Mandatory keyword while creating a VPLS service. Create the service instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

svc-sap-type

Specifies the type of service and allowed SAPs in the service.

Values **dot1q-preserve** — Specifies that the allowed SAPs in the service are dot1q. The dot1q ID is not stripped after packets match the SAP. This option can be configured in conjunction with the **b-vpls** or **r-vpls** keywords.

null-star — Specifies the allowed SAP in the service, which can be null SAP, dot1q default, Q.* SAP, 0.* SAP, or default QinQ SAP. This option can be configured in conjunction with the **b-vpls** or **r-vpls** keywords.

any — Specifies that all supported SAPs are allowed in the service. This option can be configured in conjunction with the **b-vpls**, **r-vpls**, or **i-vpls** keywords. When these keywords are not configured, **any** can be used with a plain VPLS service, which can be configured with SAPs, spoke-SDPs, and mesh SDPs. See section [QinQ SAP Configuration Restrictions for 7210 SAS platforms in network operating mode](#) for more information about restrictions related to QinQ SAPs.

Default any

b-vid *vid*

Specifies the VLAN ID to use when the **svc-sap-type** value is set to **dot1q-preserve**. This parameter is supported only when the **b-vpls** keyword and **svc-sap-type** value **dot1q-preserve** are configured.

Values 1 to 4094

r-vpls

Specifies the VPLS instance to be associated with an IP interface to provide routed VPLS (R-VPLS) functionality. When configured with the **svc-sap-type** values **null-star**, **dot1q-preserve**, and **any**, this keyword instantiates an R-VPLS service that can be configured only with SAPs.

**Note:**

The **r-vpls** keyword is not supported in access-uplink mode (that is, in access-uplink mode, a routed VPLS service can be configured without using this parameter).

vpls**Syntax**

vpls *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**m-vpls**] [**svc-sap-type** {**null-star** | **dot1q-preserve** | **dot1q-range** | **any**}] [**customer-vid** *vlan-id*]

no vpls *service-id*

Context

config>service

Platforms

Supported only on platforms configured in the access-uplink operating mode

Description

This command creates or maintains a virtual private LAN services (VPLS) instance. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites, acting like a zero-hop, Layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context. When a service is created, the **customer** keyword and *customer-id* parameter must be specified to associate the service with a customer. The *customer-id* value must already exist, having been created using the **customer** command in the service context. When 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.

When a service is created, the use of the **customer** *customer-id* command is optional for navigating into the service configuration context. Editing a service with the incorrect *customer-id* value specified results 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

Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values *service-id*: 1 to 2147483648

customer customer-id

Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

create

Mandatory keyword when creating a VPLS service. This keyword is used to create the service instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

customer-vid vlan-id

Defines the dot1q VLAN ID to be specified while creating the local dot1q SAP for the **svc-sap-type** value **dot1q-preserve**. This parameter is supported only on platforms operating in access-uplink mode.

Values 1 to 4094

svc-sap-type

Specifies the type of service and allowed SAPs in the service.

Values **dot1q-preserve** — Specifies that the allowed SAPs in the service are dot1q. The dot1q ID is not stripped after packets matches the SAP.

dot1q-range - Specifies that the access SAP in the service can use VLAN ranges as the SAP tags. The VLAN ranges are configured using the **configure>connection-profile** CLI command. On ingress of the access dot1q SAP using VLAN ranges, the outermost tag is not removed before forwarding. This option is supported only in the access-uplink operating mode.

null-star — Specifies the allowed SAP in the service, which can be null SAPs, dot1q default, Q.* SAP, 0.* SAP or default QinQ SAP.

any — Specifies that the SAPs allowed in the service are defined as shown in [Table 8: SAP and service combinations for 7210 SAS-T in access-uplink mode](#). See the section [SAP configuration notes for 7210 SAS platforms in access-uplink operating mode](#) for more information about configuring SAPs.

Default any

bgp

Syntax

bgp

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure the parameters related to BGP.

bgp-ad

Syntax

bgp-ad

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure the parameters related to BGP AD.

block-on-mesh-failure

Syntax

[no] block-on-mesh-failure

Context

config>service>vpls>spoke-sdp

config>service>vpls>endpoint

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables blocking (brings the entity to an operationally down state) after all configured mesh-SDPs are in the operationally down state. This event is signaled to corresponding T-LDP peer by withdrawing service label (status-bit-signaling non-capable peer) or by setting "PW not forwarding" status bit in T-LDP message (status-bit-signaling capable peer).

Default

disabled

bpdu-translation

Syntax

bpdu-translation {auto | pvst | stp}

no bpdu-translation

Context

config>service>vpls>spoke-sdp (not supported in access-uplink operating mode)

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables the translation of BPDUs to a specific format, meaning that all BPDUs transmitted on a specific SAP or spoke-SDP have a specified format.

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

Default

no bpdu-translation

Parameters

auto

Specifies that appropriate format is detected automatically, based on type of bpdus received on such port.

pvst

Specifies the BPDU-format as PVST. Note that the correct VLAN tag is included in the payload (depending on encapsulation value of outgoing SAP).

stp

Specifies the BPDU-format as STP.

cflowd**Syntax**

[no] cflowd

Context

config>service>vpls>sap

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE

Description

This command enables cflowd to collect traffic flow samples through a service interface (SAP) for analysis. When cflowd is enabled on an Ethernet service SAP, the Ethernet traffic can be sampled and processed by the system's cflowd engine and exported to IPFIX collectors with the **I2-ip** template enabled.

Cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement. When cflowd is enabled at the SAP level, all packets forwarded by the interface are subjected to analysis according to the cflowd configuration.

**Note:**

See the "Configuration notes" section in the "Cflowd" chapter of the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide* for more information about the sampling of packets.

The **no** form of this command disables cflowd to collect traffic flow samples through a SAP.

Default

no cflowd

l2pt-termination**Syntax**

l2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp] [lldp]

no l2pt-termination

Context

config>service>vpls>sap

config>service>vpls>spoke-sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description



Note:

The **config>service>vpls>spoke-sdp** context is not supported on 7210 SAS platforms configured in the access-uplink operating mode.

This command enables Layer 2 Protocol Tunneling (L2PT) termination on a specific SAP or spoke-SDP. L2PT termination is supported for CDP, DTP, PAGP, STP, UDLD, VTP, and LLDP PDUs.

This feature can be enabled only if STP is disabled in the context of the specific VPLS service.

Default

no l2pt-termination

Parameters

cdp

Specifies the Cisco Discovery Protocol.

dtp

Specifies the Dynamic Trunking Protocol.

pagp

Specifies the Port Aggregation Protocol.

stp

Specifies all spanning tree protocols: **stp**, **rstp**, **mstp**, **pvst** (default) values.

udld

Specifies Unidirectional Link Detection.

vtp

Specifies the VLAN Trunking Protocol.

lldp

Specifies Link Layer Discovery Protocol (LLDP). This keyword is supported only on the 7210 SAS-Mxp.

disable-aging

Syntax

[no] **disable-aging**

Context

config>service>vpls

config>service>vpls>spoke-sdp

```
config>service>vpls>sap
config>service>pw-template
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description



Note:

The **config>service>vpls>spoke-sdp** and **config>service>pw-template** contexts are not supported on platforms configured in the access-uplink operating mode.

This command disables MAC address aging across a VPLS service or on a VPLS service SAP.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In 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 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 is 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>pw-template
config>template>vpls-template
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description



Note:

The **config>service>pw-template** context is not supported on platforms configured in the access-uplink operating mode.

This command disables learning of new MAC addresses in the VPLS forwarding database (FDB) for the service instance.

When **disable-learning** is enabled, new source MAC addresses is not entered in the VPLS service forwarding database.

When **disable-learning** is disabled, new source MAC addresses is 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

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

By default, packets with unknown destination MAC addresses are flooded. If discard-unknown is enabled at the VPLS level, packets with unknown destination MAC address is dropped instead (even when configured FIB size limits for VPLS or SAP are not yet reached).

The **no** form of this command allows flooding of packets with unknown destination MAC addresses in the VPLS.

Default

no discard-unknown

endpoint

Syntax

endpoint *endpoint-name* [**create**]

no endpoint

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures a service endpoint.

Parameters

endpoint-name

Specifies an endpoint name up to 32 characters.

create

Mandatory keyword to create a service endpoint.

description

Syntax

description *description-string*

no description

Context

config>service>vpls>endpoint

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

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

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

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

Default

no description

Parameters***string***

Specifies the description character string. Allowed values are any string up to 80 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

ignore-standby-signaling**Syntax**

[no] ignore-standby-signaling

Context

config>service>vpls>endpoint

config>service>vpls>spoke-sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command causes the node to ignore the standby-bit received from TLDP peers for the specific spoke-SDP and performs internal tasks without taking it into account.

This command is present at endpoint level as well as spoke-SDP level. If the spoke-SDP is part of the explicit-endpoint, it is not possible to change this setting at the spoke-SDP level. The existing spoke-SDP becomes part of the explicit-endpoint only if the setting is not conflicting. The newly created spoke-SDP which is a part of the specific explicit-endpoint inherits this setting from the endpoint configuration.

Default

disabled

revert-time**Syntax**

revert-time *revert-time* | **infinite**

no revert-time

Context

config>service>vpls>endpoint

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the time to wait before reverting to primary spoke-SDP.

In a regular endpoint the revert-time setting affects just the pseudowire defined as primary (precedence 0). For a failure of the primary pseudowire followed by restoration the revert-timer is started. After it expires the primary pseudowire takes the active role in the endpoint. This behavior does not apply for the case when both pseudowires are defined as secondary. For example, if the active secondary pseudowire fails and is restored it stays in standby until a configuration change or a force command occurs.

Parameters

revert-time

Specifies the time to wait, in seconds, before reverting back to the primary spoke-SDP defined on this service endpoint, after having failed over to a backup spoke-SDP.

Values 0 to 600

infinite

Specifies that the endpoint is non-revertive.

static-mac

Syntax

static-mac *ieee-address* [**create**]

no static-mac

Context

config>service>vpls>endpoint

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command assigns a static MAC address to the endpoint. In the FDB, the static MAC is then associated with the active spoke-SDP.

Parameters

ieee-address

Specifies the static MAC address to the endpoint.

Values 6-byte mac-address (xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx). (Note: This value cannot be all zeros.)

create

Mandatory keyword while creating a static MAC.

suppress-standby-signaling**Syntax**

[no] **suppress-standby-signaling**

Context

config>service>vpls>endpoint

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

When this command is enabled, the pseudowire standby bit (with value 0x00000020) is not sent to T-LDP peer when the specific spoke is selected as a standby. This allows faster switchover as the traffic is sent over this SDP and discarded at the blocking side of the connection. This is particularly applicable to multicast traffic.

Default

enabled

propagate-mac-flush**Syntax**

[no] **propagate-mac-flush**

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies whether MAC flush messages received from the specific LDP are propagated to all spoke and mesh SDPs within the context of this VPLS service. The propagation follows the split-horizon principle and any datapath blocking to avoid the looping of these messages.

Default

no propagate-mac-flush

fdb-table-high-wmark

Syntax

[no] **fdb-table-high-wmark** *high-water-mark*

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the value to send logs and traps when the threshold is reached.

Parameters

high-water-mark

Specifies the value to send logs and traps when the threshold is reached.

Values 0 to 100

Default 95%

fdb-table-low-wmark

Syntax

[no] **fdb-table-low-wmark** *low-water-mark*

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the value to send logs and traps when the threshold is reached.

Parameters

low-water-mark

Specifies the value to send logs and traps when the threshold is reached.

Values 0 to 100

Default 90%

fdb-table-size

Syntax

fdb-table-size *table-size*

no fdb-table-size [*table-size*]

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the maximum number of MAC entries in the Forwarding Database (FDB) for the VPLS instance on this node.

The **fdb-table-size** specifies the maximum number of forwarding database entries for both learned and static MAC addresses for the VPLS instance.

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

Default

250

Parameters

table-size

Specifies the maximum number of MAC entries in the FDB.

vsi-export

Syntax

vsi-export *policy-name* [*policy-name...*(up to 5 max)]

no vsi-export

Context

config>service>vpls>bgp-ad

config>service>vpls>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies the name of the VSI export policies to be used for BGP auto-discovery, if this feature is configured in the VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

vsi-import

Syntax

vsi-import *policy-name* [*policy-name...*(up to 5 max)]

no vsi-import

Context

config>service>vpls>bgp-ad>vsi-id

config>service>vpls>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies the name of the VSI import policies to be used for BGP auto-discovery, if this feature is configured in the VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

pw-template-binding

Syntax

pw-template-binding *policy-id* [*split-horizon-group group-name*] [*import-rt {ext-community,...(up to 5 max)}*]

no pw-template-binding *policy-id*

Context

config>service>vpls>bgp-ad

config>service>vpls>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command binds the advertisements received with the route target (RT) that matches the configured list (either the generic or the specified import) to a specific pw-template. If the RT list is not present the pw-template is used for all of them.

The **pw-template-binding** applies to BGP-AD, if this feature is configured in the VPLS service.

The tools perform commands can be used to control the application of changes in pw-template for BGP-AD.

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

Parameters

policy-id

Specifies an existing policy ID.

Values 1 to 2147483647

split-horizon-group group-name

Specifies the group-name that overrides the split horizon group template settings.

import-rt ext-comm

Specifies the communities allowed to be accepted from remote PE neighbors. An extended BGP community in the type:x:y format. The value x can be an integer or IP address.

The type can be the target or origin, and x and y are 16-bit integers.

Values target:{ip-addr:comm-val | 2byte-asnumber:ext-comm-val | 4byte-asnumber:comm-val} ip-addr a.b.c.d
comm-val 0 to 65535
2byte-asnumber 0 to 65535
ext-comm-val 0 to 4294967295
4byte-asnumber 0 to 4294967295

route-distinguisher

Syntax

route-distinguisher [*ip-addr:comm-val* | *as-number:ext-comm-val*]

no route-distinguisher

Context

config>service>vpls>bgp-ad>vsi-id

config>service>vpls>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the Route Distinguisher (RD) component that is signaled in the MPBGP NLRI for L2VPN AFI. This value is used for BGP-AD, if this feature is configured in the VPLS service.

If this command is not configured, the RD is automatically built using the BGP-AD VPLS ID. The following rules apply:

- If BGP AD VPLS-id is configured and no RD is configured under BGP node - RD = VPLS-ID.
- If BGP AD VPLS-id is not configured then an RD value must be configured under BGP node (this is the case when only BGP VPLS is configured).
- If BGP AD VPLS-id is configured and an RD value is also configured under BGP node, the configured RD value prevails Values and format (6 bytes, other 2 bytes of type is automatically generated)

Parameters

ip-addr:comm-val

Specifies the IP address.

Values *ip-addr* a.b.c.d
 comm-val 0 to 65535
 as-number:ext-comm-val — Specifies the ASN and the
 as-number 1 to 65535
 ext-comm-val 0 to 4294967295

local-age

Syntax

local-age *aging-timer*
no local-age

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the aging time for locally learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance. In a VPLS service, MAC addresses are

associated with a Service Access Point (SAP). MACs associated with a SAP are classified as local MACs, and MACs associated with are remote MACs QinQ / access-uplink SAPs.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). The **local-age** timer specifies the aging time for local learned MAC addresses.

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

Default

local age 300 — Local MACs aged after 300 seconds.

Parameters

aging-timer

Specifies the aging time for local MACs expressed in seconds.

Values 60 to 86400

mac-move

Syntax

[no] **mac-move**

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures MAC move attributes. 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. If the rate exceeds the configured maximum allowed limit, it disables the SAP where the source MAC was last seen. The SAP can be disabled permanently (until a **shutdown/no shutdown** command is executed) or for a length of time that grows linearly with the number of times the specific SAP was disabled. You have the option of marking a SAP as non-blockable in the **config>service>vpls>sap>limit-mac-move** context. This means that when the relearn rate has exceeded the limit, another (blockable) SAP is disabled instead.

The **mac-move** command enables the feature at the service level for SAPs, as only those objects can be blocked by this feature.

The operation of this feature is the same on the SAP. For example, if a MAC address moves from SAP to SAP, one is blocked to prevent thrashing.

The **mac-move** command disables a VPLS port when the number of relearns detected has reached the number of relearns needed to reach the move-frequency in the 5-second interval. For example, when the

move-frequency is configured to 1 (relearn per second) mac-move disables one of the VPLS ports when 5 relearns were detected during the 5-second interval because then the average move-frequency of 1 relearn per second has been reached. This can already occur in the first second if the real relearn rate is 5 relearns per second or higher.

The **no** form of this command disables MAC move.

move-frequency

Syntax

move-frequency *frequency*

no move-frequency

Context

config>service>vpls>mac-move

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies 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 to protect the system against undetected loops or duplicate MACs.

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

Default

2 (when mac-move is enabled). For example, 10 relearns in a 5 second period.

Parameters

frequency

Specifies the rate, in 5-second intervals for the maximum number of relearns.

Values 1 to 100

retry-timeout

Syntax

retry-timeout *timeout*

no retry-timeout

Context

config>service>vpls>mac-move

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies 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 retry-timeout value is larger or equal to 5s * cumulative factor of the highest priority port so that the sequential order of port blocking is not disturbed by reinitializing lower priority ports.

A zero value indicates that the SAP is not 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.

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

Default

10 (when mac-move is enabled)

Parameters

timeout

Specifies the time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.

Values 0 to 120

mfib-table-high-wmark

Syntax

[no] **mfib-table-high-wmark** *high-water-mark*

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the multicast FIB high watermark. When the percentage filling level of the multicast FIB exceeds the configured value, a trap is generated and a log entry is added.

Parameters

high-water-mark

Specifies the multicast FIB high watermark as a percentage.

Values	1 to 100
Default	95%

mfib-table-low-wmark

Syntax
[no] **mfib-table-low-wmark** *low-water-mark*

Context
config>service>vpls

Platforms
Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description
This command specifies the multicast FIB low watermark. When the percentage filling level of the Multicast FIB drops below the configured value, the corresponding trap is cleared and a log entry is added.

Parameters
low-water-mark
Specifies the multicast FIB low watermark as a percentage.

Values	1 to 100
Default	90%

mfib-table-size

Syntax
mfib-table-size *size*
no mfib-table-size

Context
config>service>vpls

Platforms
Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the maximum number of (s,g) entries in the multicast forwarding database (MFIB) for this VPLS instance.

The *size* parameter specifies the maximum number of multicast database entries for both learned and static multicast addresses for the VPLS instance. When a table-size limit is set on the mfib of a service which is lower than the current number of dynamic entries present in the mfib then the number of entries remains above the limit.

The **no** form of this command removes the configured maximum MFIB table size.

Parameters

size

Specifies the maximum number of (s,g) entries allowed in the Multicast FIB.

Values 1 to 2047 (M and Mxp)
 1 to 2043(T)

remote-age

Syntax

remote-age *seconds*

no remote-age

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the aging time for remotely learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance. In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The **remote-age** timer specifies the aging time for remote learned MAC addresses. To reduce the amount of signaling required between switches configure this timer larger than the **local-age** timer.

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

Default

remote age 900

Parameters

seconds

Specifies 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

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables sending out "flush-all-from-ME" messages to all LDP peers included in affected VPLS, in the event of physical port failures or "oper-down" events of individual SAPs. This feature provides an LDP-based mechanism for recovering a physical link failure in a dual-homed connection to a VPLS service. This method provides an alternative to RSTP solutions where dual homing redundancy and recovery, in the case of link failure, is resolved by RSTP running between a PE router and CE devices. If the endpoint is configured within the VPLS and send-flush-on-failure is enabled, flush-all-from-me messages is sent out only when all spoke-SDPs associated with the endpoint go down.

This feature cannot be enabled on management VPLS.

Default

no send-flush-on-failure

service-mtu

Syntax

service-mtu *octets*

no service-mtu

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the service payload (Maximum Transmission Unit – MTU), in bytes, for the service. This MTU value overrides the service-type default MTU. The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP binding operational state within the service.

The service MTU and a SAP service delineation encapsulation overhead (that is, 4 bytes for a dot1q tag) is used to derive the required MTU of the physical port or channel on which the SAP was created. If the required payload is larger than the port or channel MTU, then the SAP is placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP is able to transition to the operative 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 reevaluated.

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



Note:

To disable service MTU check execute the command **no service-mtu-check**. Disabling service MTU check allows the packets to pass to the egress if the packet length is lesser than or equal to the MTU configured on the port.

Default

VPLS: 1514

The following table displays MTU values for specific VC types.

Table 53: MTU Values for VC Types (VPLS)

VC-type	Example service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
VPLS	1514	1500
VPLS (with preserved dot1q)	1518	1504
VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (QinQ with preserved bottom Qtag)	1518	1504

Parameters

octets

Specifies the size of the MTU in octets, expressed as a decimal integer.

Values 1 to 9194

service-mtu-check

Syntax

[no] **service-mtu-check**

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

Disabling service MTU check allows the packets to pass to the egress if the packet length is lesser than or equal to the MTU configured on the port. The length of the packet sent from a SAP is limited only by the access port MTU. In case of a pseudowire the length of a packet is limited by the network port MTU (including the MPLS encapsulation).



Note:

If TLDP is used for signaling, the configured value for **service-mtu** is used during pseudowire setup.

The **no** form of this command disables the service MTU check.

Default

enabled

service-name

Syntax

service-name *service-name*

no service-name

Context

config>service>epipe

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures an optional service name, up to 64 characters, which adds a name identifier to a specific service to then use that service name in configuration references as well as display and use

service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7210 SAS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a specific service when 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).

split-horizon-group

Syntax

[no] split-horizon-group [*group-name*] [**create**]

Context

config>service>vpls

config>service>pw-template

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates a new split-horizon group (SHG) for the VPLS instance. Traffic arriving on a SAP or spoke-SDP within this SHG is not copied to other SAPs or spoke-SDPs in the same SHG.

The SHG must be created before SAPs and spoke-SDPs can be assigned to the group.

The SHG is defined within the context of a single VPLS. The same group name can be reused in different VPLS instances.



Note:

- The **split-horizon-group** command is only supported on 7210 SAS platforms operating in the network mode.
- Service-based SHGs are only supported on the 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE operating in the network mode. On these platforms, service-based SHGs and mesh-SDPs are mutually exclusive in a VPLS service.
- On the 7210 SAS-Mxp, an SHG can be used with spoke-SDPs or mesh SDPs configured in the service.
- Service-based SHGs are not supported in an R-VPLS service.

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

Parameters

group-name

Specifies the name of the SHG to which the SAP or spoke-SDP belongs.

create

Mandatory keyword to create an SHG.

root-guard

Syntax

[no] root-guard

Context

config>service>vpls>sap>stp

config>service>vpls>spoke-sdp>stp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies whether this port is allowed to become an STP root port. It corresponds to the restrictedRole parameter in 802.1Q. If set, it can cause lack of spanning tree connectivity.

Default

no root-guard

tod-suite

Syntax

tod-suite *tod-suite-name*

no tod-suite

Context

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command applies a time-based policy (filter or QoS policy) to the service SAP. The suite name must already exist in the **config>cron** context.

Default

no tod-suite

Parameters***tod-suite-name***

Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP. The suite can be applied to more than one SAP.

vpls-id**Syntax****vpls-id** *vpls-id***Context**

config>service>vpls>bgp-ad

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the VPLS ID component that is signaled in one of the extended community attributes (ext-comm).

Values and format (6 bytes, other 2 bytes of type-subtype are automatically generated).

Parameters***vpls-id***

Specifies a globally unique VPLS ID for BGP auto-discovery in this VPLS service.

Values *vpls-id* : *ip-addr:comm-val*>|<*as-number:ext-comm-val*
 ip-addr: a.b.c.d
 comm-val 0 to 65535
 as-number: 1 to 65535
 ext-comm-val 0 to 4294967295

vsi-id**Syntax****vsi-id**

Context

config>service>vpls>bgp-ad

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure the Virtual Switch Instance Identifier (VSI-ID).

prefix

Syntax

prefix low-order-vsi-id
no prefix

Context

config>service>vpls>bgp-ad>vsi-id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies the low-order 4 bytes used to compose the Virtual Switch Instance Identifier (VSI-ID) to use for NLRI in BGP auto-discovery in this VPLS service.
If no value is set, the system IP address is used.

Default

no prefix

Parameters

low-order-vsi-id
Specifies a unique VSI-ID.
Values 0 to 4294967295

service-name

Syntax

service-name *service-name*

no service-name**Context**

```
config>service>vpls
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures an optional service name, up to 64 characters, which adds a name identifier to a specific service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7210 SAS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a specific service when 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 to 9).

5.8.2.1.3 VPLS STP commands

```
stp
```

Syntax

```
stp
```

Context

```
config>service>vpls
```

```
config>service>vpls>sap
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure the Spanning Tree Protocol (STP) parameters.

The Nokia STP has a few modifications to better suit the operational characteristics of VPLS services. The most evident change is to the root bridge election. Because the core network operating between the Nokia service routers should not be blocked, the root path is calculated from the core perspective.

auto-edge

Syntax

auto-edge

no auto-edge

Context

config>service>vpls>sap>stp

config>service>vpls>spoke-sdp>stp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures automatic detection of the edge port characteristics of the SAP or spoke-SDP.

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

Default

auto-edge

edge-port

Syntax

[no] edge-port

Context

config>service>vpls>sap>stp

config>service>vpls>spoke-sdp>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description



Note:

The **config>service>vpls>spoke-sdp>stp** context is not supported on platforms configured in the access-uplink operating mode.

This command configures the SAP or SDP as an edge or non-edge port. If **auto-edge** is enabled for the SAP, this value is used only as the initial value.

RSTP, however, can detect that the actual situation is different from what **edge-port** may indicate.

Initially, the value of the SAP or spoke-SDP parameter is set to edge-port. This value changes if:

- A BPDU is received on that port. This means that after all there is another bridge connected to this port. Then the edge-port becomes disabled.
- If auto-edge is configured and no BPDU is received within a specific period of time, RSTP concludes that it is on an edge and enables the edge-port.

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

Default

no edge-port

forward-delay

Syntax

forward-delay *seconds*

no forward-delay

Context

config>service>vpls>stp

config>service>vpls>spoke-sdp>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description



Note:

The **config>service>vpls>spoke-sdp>stp** context is not supported on platforms configured in the access-uplink operating mode.

RSTP, as defined in the IEEE 802.1D-2004 standards, transition to the forwarding state via a handshaking mechanism (rapid transition), without any waiting times. If handshaking fails (for example, on shared links), the system falls back to the timer-based mechanism defined in the original STP (802.1D-1998) standard.

A shared link is a link with more than two nodes (for example, a shared 10/100BaseT segment). The **port-type** command is used to configure a link as point-to-point or shared.

For timer-based transitions, the 802.1D-2004 standard defines an internal variable forward-delay, which is used in calculating the default number of seconds that a SAP spends in the discarding and learning states when transitioning to the forwarding state.

The value of the forward-delay variable depends on the STP operating mode of the VPLS instance:

- in **rstp** or **mstp** mode, but only when the SAP has not fallen back to legacy STP operation, the value configured by the **hello-time** command is used.
- in all other situations, the value configured by the **forward-delay** command is used.

Default

15 seconds

Parameters***seconds***

Specifies the forward delay timer for the STP instance in seconds.

Values 4 to 30

hello-time

Syntax

hello-time *hello-time*

no hello-time

Context

config>service>vpls>stp

config>service>vpls>spoke-sdp>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description**Note:**

The **config>service>vpls>spoke-sdp>stp** context is not supported on platforms configured in the access-uplink operating mode.

This command configures the STP hello time for the Virtual Private LAN Service (VPLS) STP instance.

The hello time parameter defines the default timer value that controls the sending interval between BPDU configuration messages by this bridge, on ports where this bridge assumes the designated role.

The active hello time for the spanning tree is determined by the root bridge (except when the STP is running in RSTP mode, then the hello time is always taken from the locally configured parameter).

The configured hello-time can also be used to calculate the forward delay. See [auto-edge](#).

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

Default

2

Parameters***hello-time***

Specifies the hello time for the STP instance in seconds.

Values 1 to 10

hold-count

Syntax

hold-count *BDPU tx hold count*

no hold-count

Context

config>service>vpls>stp

config>service>vpls>spoke-sdp>stp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the peak number of BPDUs that can be transmitted in a period of one second.

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

Default

6

Parameters

BDPU tx hold count

Specifies the hold count for the STP instance in seconds.

Values 1 to 10

link-type

Syntax

link-type {pt-pt | shared}

no link-type

Context

config>service>vpls>sap>stp

config>service>vpls>spoke-sdp>stp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command instructs STP on the maximum number of bridges behind this SAP. If there is only a single bridge, transitioning to forwarding state is based on handshaking (fast transitions). If more than two bridges are connected via a shared media, their SAP should all be configured as shared, and timer-based transitions are used.

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

Default

pt-pt

mst-instance

Syntax

mst-instance *mst-inst-number*

Context

config>service>vpls>sap>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures MSTI related parameters at SAP level. This context can be open only for existing mst-instances defined at the service level.

Parameters

mst-inst-number

Specifies an existing MSTI number.

Values 1 to 4094

mst-path-cost

Syntax

mst-path-cost *inst-path-cost*

no mst-path-cost

Context

config>service>vpls>sap>stp>mst-instance

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This commands specifies path-cost within a specific instance. If a loop occurs, this parameter indicates the probability of a specific port being assigned a forwarding state. (The highest value expresses lowest priority).

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

Default

The path-cost is proportional to link speed.

Parameters

inst-path-cost

Specifies the contribution of this port to the MSTI path cost.

Values 1 to 200000000

mst-port-priority

Syntax

mst-port-priority *stp-priority*

no mst-port-priority

Context

config>service>vpls>sap>stp>mst-instance

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This commands specifies the port priority within a specific instance. If a loop occurs, this parameter indicates the probability of a specific port being assigned a forwarding state.

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

Default

128

Parameters

stp-priority

Specifies the value of the port priority field.

max-age

Syntax

max-age *seconds*

no max-age

Context

config>service>vpls>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command indicates how many hops a BPDU can traverse the network starting from the root bridge. The message age field in a BPDU transmitted by the root bridge is initialized to 0. Each other bridge takes the message_age value from BPDUs received on their root port and increment this value by 1. The message_age therefore reflects the distance from the root bridge. BPDUs with a message age exceeding max-age are ignored.

STP uses the max-age value configured in the root bridge. This value is propagated to the other bridges via the BPDUs.

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

Default

20

Parameters

seconds

Specifies the max info age for the STP instance in seconds. Allowed values are integers in the range 6 to 40.

mode

Syntax

mode {*rstp* | *comp-dot1w* | *dot1w* | *mstp* | *pmstp*}

no mode

Context

config>service>vpls>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the version of STP the bridge is currently running.

See [Spanning tree operating modes](#) for more information about these modes.

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

Default

rstp

Parameters

rstp

Corresponds to the Rapid Spanning Tree Protocol specified in IEEE 802.1D/D4-2003.

dot1w

Corresponds to the mode where the Rapid Spanning Tree is backward compatible with IEEE 802.1w.

compdot1w

Corresponds to the Rapid Spanning Tree Protocol fully conformant to IEEE 802.1w.

mstp

Sets MSTP as the STP mode of operation. Corresponds to the Multiple Spanning Tree Protocol specified in 802.1Q REV/D5.0-09/2005

pmstp

Specifies the PMSTP mode, which is only supported in VPLS services where the mVPLS flag is configured.

mst-instance

Syntax

[no] **mst-instance** *mst-inst-number*

Context

config>service>vpls>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures Multiple Spanning Tree Instance (MSTI) related parameters. MSTP supports 16 instances. The instance 0 is mandatory (by protocol) and cannot be created by the CLI. The software automatically maintains this instance.

Parameters

mst-inst-number

Specifies the MST instance.

Values 1 to 4094

mst-priority

Syntax

mst-priority *bridge-priority*

no mst-priority

Context

config>service>vpls>stp>mst-instance

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the bridge priority for this specific Multiple Spanning Tree Instance for this service. The *bridge-priority* value reflects likelihood that the switch is chosen as the regional root switch (65535 represents the least likely). It is used as the highest 4 bits of the Bridge ID included in the MSTP BPDU's generated by this bridge.

The values of the priority are only multiples of 4096 (4k). If a value is specified that is not a multiple of 4K, the value is replaced by the closest multiple of 4K (lower than the value entered).

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

Default

32768 — All instances that are created by the [vlan-range](#) command do not have explicit definition of bridge-priority and inherits the default value.

Parameters

bridge-priority

Specifies the priority of this specific Multiple Spanning Tree Instance for this service.

Values 0 to 65535

vlan-range

Syntax

[no] vlan-range *[vlan-range]*

Context

config>service>vpls>stp>mst-instance

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies a range of VLANs associated with a specific MST-instance. This range applies to all SAPs of the mVPLS.

Every VLAN range that is not assigned within any of the created [mst-instance](#) is automatically assigned to mst-instance 0. This instance is automatically maintained by the software and cannot be modified. Changing the VLAN range value can be performed only when the specific mst-instance is shutdown.

The **no** form of this command removes the **vlan-range** from a specific mst-instance.

Parameters

vlan-range

The first VLAN range specifies the left-bound (that is, minimum value) of a range of VLANs that are associated with the mVPLS SAP. This value must be smaller than (or equal to) the second VLAN range value. The second VLAN range specifies the right-bound (that is, maximum value) of a range of VLANs that are associated with the mVPLS SAP.

Values 1 to 4094

mst-max-hops

Syntax

mst-max-hops *hops-count*

no mst-max-hops

Context

config>service>vpls>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the number of hops in the region before BPDU is discarded and the information held for the port is aged out. The root bridge of the instance sends a BPDU (or M-record) with remaining-hop-count set to configured *max-hops*. When a bridge receives the BPDU (or M-record), it decrements the received remaining-hop-count by 1 and propagates it in BPDU (or M-record) it generates.

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

Default

20

Parameters

hops-count

Specifies the maximum number of hops.

Values 1 to 40

mst-name

Syntax

mst-name *region-name*

no mst-name

Context

config>service>vpls>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command defines an MST region name. Two bridges are considered part of the same MST region when their configuration of the MST region name, the MST-revision and VLAN-to-instance assignment, is identical.

The **no** form of this command removes *region-name* from the configuration.

Default

no mst-name

Parameters

region-name

Specifies an MST-region name up to 32 characters.

mst-revision

Syntax

mst-revision *revision-number*

Context

config>service>vpls>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command defines the MST configuration revision number. Two bridges are considered as a part of the same MST region if their configured MST-region name, MST-revision, and VLAN-to-instance are identical.

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

Default

0

Parameters

revision-number

Specifies the MSTP region revision number to define the MSTP region.

Values 0 to 65535

path-cost

Syntax

path-cost *sap-path-cost*

no path-cost

Context

config>service>vpls>sap>stp

config>service>vpls>spoke-sdp>stp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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. When BPDUs are sent out other egress SAPs or 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 SAPs are controlled by complex queuing dynamics, in the 7210 SAS the STP path cost is a purely static configuration.

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

Parameters

path-cost

Specifies the path cost for the SAP or spoke-SDP.

Values 1 to 200000000 (1 is the lowest cost)

Default 10

port-num

Syntax

[no] port-num *virtual-port-number*

Context

config>service>vpls>sap>stp

config>service>vpls>spoke-sdp>stp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the virtual port number which uniquely identifies a SAP within configuration bridge protocol data units (BPDUs). The internal representation of a SAP is unique to a system and has a reference space much bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a SAP and identifies it with its own virtual port number that is unique to every other SAP defined on the TLS. The virtual port number is assigned at the time that the SAP is added to the TLS. Because the order that the SAP was added to the TLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance.

The virtual port number cannot be administratively modified.

priority

Syntax

priority *bridge-priority*

no priority

Context

config>service>vpls>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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 are truncated to multiples of 4096, conforming with IEEE 802.1t and 802.1D-2004.

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

Default

4096

Parameters

bridge-priority

Specifies the bridge priority for the STP instance.

Values	Allowed values are integers in the range of 4096 to 65535 with 4096 being the highest priority. The actual bridge priority value stored/used is the number entered with the lowest 12 bits masked off, which means the actual range of values is 4096 to 61440 in increments of 4096.

priority

Syntax

priority *stp-priority*

no priority

Context

config>service>vpls>spoke-sdp (not supported in access-uplink operating mode)

config>service>vpls>sap>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the Nokia STP priority for the SAP or spoke-SDP.

STP priority is a configurable parameter associated with a SAP or spoke-SDP. When configuration BPDUs are received, the priority is used in some circumstances as a tie breaking mechanism to determine whether the SAP or spoke-SDP is designated or blocked.

In traditional STP implementations (802.1D-1998), this field is called the port priority and has a value of 0 to 255. This field is coupled with the port number (0 to 255 also) to create a 16 bit value. In the latest STP standard (802.1D-2004) only the upper 4 bits of the port priority field are used to encode the SAP or spoke-SDP priority. The remaining 4 bits are used to extend the port ID field into a 12 bit virtual port number field. The virtual port number uniquely references a SAP within the STP instance.

STP computes the actual priority by taking the input value and masking out the lower four bits. The result is the value that is stored in the priority parameter. For instance, if a value of 0 is entered, masking out the lower 4 bits results in a parameter value of 0. If a value of 255 is entered, the result is 240.

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

Default

128

Parameters

stp-priority

Specifies the STP priority value for the SAP. Allowed values are integers in the range of 0 to 255, 0 being the highest priority. The actual value used for STP priority (and stored in the configuration) is the result of masking out the lower 4 bits, therefore the actual value range is 0 to 240 in increments of 16.

5.8.2.1.4 VPLS SAP commands

```
sap
```

Syntax

sap *sap-id* [**split-horizon-group** *group-name*] [**create**] [**g8032-shg-enable**] [**eth-ring** *ring-index*] (for 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T (network mode), 7210 SAS Mxp)

sap *sap-id* [**g8032-shg-enable**] [**eth-ring** *ring-index*] [**create**] (for 7210 SAS-T (access-uplink mode))

no sap *sap-id*

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the 7210 SAS. Each SAP must be unique.

A physical port can have only one SAP to be part of one service. Multiple SAPs can be defined over a physical port but each of these SAPs should belong to a different service.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP does not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface port-type port-id mode access** command.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service is 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 is also deleted.

This command is also used to create a Ring APS Control SAP or a Data SAP whose traffic is protected by a Ring APS Instance.

Special Cases

A default SAP has the following format: port-id:*. This type of SAP is supported only on Ethernet MDAs and its creation is allowed only in the scope of Layer 2 services (Epipe and VPLS). The 7210 SAS supports explicit null encapsulation for VPLS service.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

create

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

g8032-shg-enable

Platforms Supported - 7210 SAS-T (access-uplink).

This command must only be used with the SAPs created in the service for the virtual channel on the interconnection nodes in a topology that uses multiple rings. This command creates a split-horizon group to ensure that Sub-Ring control messages from the major ring are only passed to the Sub-Ring control service.

eth-ring

Keyword to create an instance of a Ring APS Control SAP or a Data SAP whose traffic is protected by a Ring APS Instance.

ring-index

Specifies the ring index of the Ethernet ring.

split-horizon-group *group-name*

Specifies the name of the split horizon group to which the SAP belongs.

discard-unknown-source**Syntax**

[no] discard-unknown-source

Context

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

When this command is enabled, packets received on a SAP or a spoke-SDP with an unknown source MAC address is 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 address to be forwarded by destination MAC addresses in VPLS.

Default

no discard-unknown-source

5.8.2.1.5 VPLS service SAP DHCP snooping commands

dhcp**Syntax**

dhcp

Context

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure DHCP snooping parameters.

action

Syntax

action {replace | drop | keep}

no action

Context

config>service>vpls>sap>dhcp>option

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the Relay Agent Information Option (Option 82) processing.

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

Default

no action

Parameters

replace

Specifies that 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

Specifies that a DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.

keep

Specifies that the existing information is kept in the packet and the router does not add any more information. In the downstream direction the Option 82 field is not stripped and is forwarded toward the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router inserts its own VSO into the Option 82 field. This is only done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router does not create the Option 82 field. In this in that case, no VSO is added to the message.

circuit-id

Syntax

circuit-id [ascii-tuple | vlan-ascii-tuple]

no circuit-id

Context

config>service>vpls>sap>dhcp>option

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAPID, separated by "|". If no keyword is configured, then the **circuit-id** sub-option is not part of the information option (Option 82).

When the command is configured without any parameters, it equals to **circuit-id** ascii-tuple.

If disabled, the **circuit-id** sub-option of the DHCP packet is left empty.



Note:

By default, **circuit-id** is enabled if any of the other options, such as remote-id or vso, are configured.

Default

no circuit-id

Parameters

ascii-tuple

Specifies that the ASCII-encoded concatenated tuple is used which consists of the access-node-identifier, service-id, and interface-name is used.

hex

Specifies the circuit-id hex string.

vlan-ascii-tuple

Specifies that the format includes VLAN ID and dot1p bits as well as what is included in ascii-tuple already. The format is supported on dot1q and qinq encapsulated ports only. Therefore, when the Option 82 bits are stripped, dot1p bits are copied to the Ethernet header of an outgoing packet.

option

Syntax

[no] option

Context

config>service>vpls>sap>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.

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

Default

no option

remote-id

Syntax

remote-id [mac | string *string*]

no remote-id

Context

config>service>vpls>sap>dhcp>option

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

When enabled, the router sends the MAC address of the remote end (typically the DHCP client) in the **remote-id** sub-option of the DHCP packet. This command identifies the host at the other end of the circuit.

If disabled, the **remote-id** sub-option of the DHCP packet is left empty.

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

Default

remote-id

Parameters

mac

Specifies the MAC address of the remote end is encoded in the sub-option.

string *string*

Specifies the remote-id.

vendor-specific-option

Syntax

[no] vendor-specific-option

Context

config>service>vpls>sap>dhcp>option

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the vendor specific sub-option of the DHCP relay packet.

client-mac-address

Syntax

[no] client-mac-address

Context

config>service>vpls>sap>dhcp>option>vendor

config>service>ies>if>dhcp>option>vendor

config>service>vprn>if>dhcp>option>vendor

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables sending the MAC address in the vendor specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the MAC address in the vendor specific sub-option of the DHCP relay packet.

sap-id

Syntax

[no] sap-id

Context

config>service>vpls>sap>dhcp>option>vendor

config>service>ies>if>dhcp>option>vendor

config>service>vprn>if>dhcp>option>vendor

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables sending the SAP ID in the vendor specific suboption of the DHCP relay packet.

The **no** form of this 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

config>service>ies>if>dhcp>option>vendor

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables sending the service ID in the vendor specific suboption of the DHCP relay packet.

The **no** form of this command disables the sending of the service ID in the vendor specific suboption of the DHCP relay packet.

string

Syntax

[no] **string** *text*

Context

config>service>vpls>sap>dhcp>option>vendor

config>service>ies>if>dhcp>option>vendor

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the string in the vendor specific suboption of the DHCP relay packet.

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

Parameters

text

The string can be any combination of ASCII characters up to 32 characters. If spaces are used in the string, enclose the entire string in quotation marks (" ").

system-id

Syntax

[no] **system-id**

Context

config>service>vpls>sap>dhcp>option>vendor

config>service>ies>if>dhcp>option>vendor

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies whether the system-id is encoded in the vendor specific sub-option of Option 82.

relay-plain-bootp

Syntax

relay-plain-bootp

no relay-plain-bootp

Context

config>service>vpls>if>dhcp

config>service>ies>if>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables the relaying of plain BOOTP packets.

The **no** form of this command disables the relaying of plain BOOTP packets.

server

Syntax

server *server1* [*server2*...(up to 8 max)]

Context

config>service>ies>if>dhcp

config>service>vprn>if>dhcp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies a list of servers where requests are 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 then the request is forwarded to all of the servers in the list.

There can be a maximum of eight DHCP servers configured.

Default

no server

Parameters

server

Specifies the DHCP server IP address.

trusted

Syntax

[no] trusted

Context

config>service>ies>if>dhcp

config>service>vprn>if>dhcp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables relaying of untrusted packets.

The **no** form of this command disables the relay.

Default

not enabled

snoop

Syntax

[no] snoop

Context

config>service>vpls>sap>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables DHCP snooping of DHCP messages on the SAP. Enabling DHCP snooping on VPLS interfaces (SAPs) is required where DHCP messages where Option 82 information is to be inserted. This includes interfaces that are in the path to receive messages from either DHCP servers or from subscribers.

The **no** form of this command disables DHCP snooping on the specified VPLS SAP.

Default

no snoop

5.8.2.1.6 VPLS DHCPv6 snooping commands for SAP and SDP bindings**dhcp6****Syntax**

dhcp6

Context

config>service>vpls>mesh-sdp
config>service>vpls>sap
config>service>vpls>spoke-sdp

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure DHCPv6 parameters.

option**Syntax**

option
no option

Context

config>service>vpls>sap>dhcp6

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enters the context for configuring option-18 (interface ID option) and option-37 (remote ID option) suboptions.

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

Default

no option

interface-id

Syntax

interface-id

interface-id *ascii-tuple*

interface-id *sap-id*

interface-id *string* *string*

interface-id *vlan-ascii-tuple*

no interface-id

Context

config>service>vpls>sap>dhcp6>option

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the interface ID option to be inserted in the DHCPv6 client messages and sent toward the DHCPv6 server.

The **no** form of this command disables sending the interface ID option.

Default

Nothing is enabled by default, but if the **option** command is enabled but not configured explicitly the default is **interface-id** *ascii-tuple* [*system-name* | *service-id* | *sap-id*].

Parameters

ascii-tuple

Keyword to specify the use of the ASCII-encoded concatenated tuple, which consists of the *system-name*, *service-id*, and *sap-id*, separated by the pipe (" | ") symbol.

sap-id

Keyword to specify the use of the SAP identifier.

string

Specifies a string of up to 32 characters, composed of printable, seven-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

vlan-ascii-tuple

Keyword to specify that the format includes a VLAN ID and dot1p bits, in addition to the information already included in ASCII tuple. The format is supported on dot1q and qinq ports only.

remote-id

Syntax

no remote-id

remote-id

remote-id mac

remote-id string *string*

Context

config>service>vpls>sap>dhcp6>option

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command configures the information that is included in the remote ID suboption that is inserted in the DHCPv6 messages received from the client and forwarded toward the DHCPv6 sever.

If disabled, the **remote-id** suboption of the DHCPv6 packet is left empty. When this command is configured without any parameters, its behavior is the same as the **remote-id mac** option.

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

Default

If the **remote-id** command is enabled but not configured explicitly the default is **remote-id mac**.

Parameters

mac

Keyword to specify that the MAC address of the remote end is encoded in the suboption.

string

Specifies the remote ID, up to 32 characters.

snoop

Syntax

no snoop

snoop [**network-facing**]

snoop [**client-facing** | **network-facing** | **both**]

Context

config>service>vpls>mesh-sdp>dhcp6

config>service>vpls>sap>dhcp6


```
config>service>vpls>spoke-sdp>dhcp6
```

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables DHCPv6 snooping of DHCP messages on the SAP or SDP binding. DHCPv6 snooping is enabled on VPLS interfaces (SAPs or SDP bindings) and indicates where option-18 (interface ID option) and option-37 (remote ID option) information must be inserted or removed in a DHCPv6 message. This includes interfaces that are in the path to receive messages from either DHCPv6 servers or DHCPv6 clients.

The **no** form of this command disables DHCPv6 snooping on the specified VPLS SAP.

Default

snoop network-facing (for SDP bindings)

snoop client-facing (for SAPs)

Parameters

client-facing

Keyword to specify that the service object is client-facing. On client-facing ports, only DHCPv6 client messages are processed. The parameter is only supported under the **config>service>vpls>sap>dhcp6>snoop** context.

network-facing

Keyword to specify that the service object is network-facing. On network-facing ports, only DHCPv6 relay-reply messages are processed.

both

Keyword to specify that the service object is both client-facing and network-facing. Specifying **both** indicates that DHCPv6 messages received from both the client and server/relay must be processed. The parameter is only supported under the **config>service>vpls>sap>dhcp6>snoop** context.

trusted

Syntax

no trusted

trusted

Context

```
config>service>vpls>sap>dhcp6
```

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables the forwarding of relay-forward messages received on service objects configured as **client-facing** or **both**.

The **no** form of this command results in dropping the relay-forward messages received on service objects configured as **client-facing** or **both**.

Default

no trusted

5.8.2.1.7 ETH-CFM service commands

eth-cfm

Syntax

eth-cfm

Context

config>service>vpls

config>service>vpls>mesh-sdp (not supported in access-uplink operating mode)

config>service>vpls>spoke-sdp (not supported in access-uplink operating mode)

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure ETH-CFM parameters.

mep

Syntax

mep *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {**up** | **down**}] **primary-vlan-enable**

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

Context

config>service>vpls>mesh-sdp>eth-cfm (not supported in access-uplink operating mode)

config>service>vpls>sap>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the ETH-CFM maintenance endpoint (MEP).

Parameters

mep-id

Specifies the maintenance association end point identifier.

Values 1 to 8191

md-index

Specifies the maintenance domain (MD) index value.

Values 1 to 4294967295

ma-index

Specifies the MA index value.

Values 1 to 4294967295

direction up | down

Indicates the direction in which the maintenance association (MEP) faces on the bridge port. Direction is not supported when a MEP is created directly under the `vpls>eth-cfm` construct (vMEP).

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

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

primary-vlan-enable

Provides a method for linking the MEP with the primary VLAN configured under the bridge-identifier for the MA. MEPs cannot be changed from or to primary vlan functions. This must be configured as part of the creation step and can only be changed by deleting the MEP and recreating it. Primary VLANs are only supported under Ethernet SAPs. This parameter is only supported on 7210 SAS-T (network mode), 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC).

ais-enable

Syntax

[no] ais-enable

Context

config>service>vpls>mesh-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

config>service>vpls>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables the generation and reception of AIS messages.

client-meg-level

Syntax

client-meg-level *[/level [/level ...]]*
no client-meg-level

Context

config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable (not supported in access-uplink operating mode)
config>service>vpls>sap>eth-cfm>mep>ais-enable

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the client maintenance entity group (MEG) levels to use for AIS message generation. Up to 7 levels can be provisioned with the restriction that the client MEG level must be higher than the local MEG level.

Parameters

<i>level</i>	Specifies the client MEG level.
Values	1 to 7
Default	1

interval

Syntax

interval {1 | 60}
no interval

Context

config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable (not supported in access-uplink operating mode)

```
config>service>vpls>sap>eth-cfm>mep>ais-enable
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the transmission interval of AIS messages in seconds.

Parameters

1 | 60

Specifies the transmission interval of AIS messages in seconds.

Default 1

priority

Syntax

priority *priority-value*

no priority

Context

config>service>vpls>mesh-sdp>eth-cfm>mep>ais-enable (not supported in access-uplink operating mode)

config>service>vpls>sap>eth-cfm>mep>ais-enable

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the priority of AIS messages originated by the node.

Parameters

priority-value

Specifies the priority value of the AIS messages originated by the node.

ccm-enable

Syntax

[no] ccm-enable

Context

config>service>vpls>mep

config>service>vpls>sap>eth-cfm>mep

config>service>vpls>mesh-sdp>mep (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables the generation of CCM messages.

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

ccm-ltm-priority

Syntax

ccm-ltm-priority *priority*

no ccm-ltm-priority

Context

config>service>vpls>sap>eth-cfm>mep

config>service>vpls>mesh-sdp>mep (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

The **no** form of this 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

eth-test-enable

Syntax

[no] eth-test-enable

Context

config>service>vpls>spoke-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

config>service>vpls>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

For ETH-test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test then can be done using the following OAM commands:

oam eth-cfm eth-test mac-address mep mep-id domain md-index association ma-index [priority priority] [data-length data-length]

A check is done for both the provisioning and test to ensure that the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP indicates the problem.

test-pattern

Syntax

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

no test-pattern

Context

config>service>vpls>sap>eth-cfm>mep>eth-test-enable

config>service>vpls>mesh-sdp>eth-cfm>mep>eth-test-enable (not supported in access-uplink operating mode)

config>service>vpls>spoke-sdp>eth-cfm>mep>eth-test-enable (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the test pattern for eth-test frames.

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

Parameters

- all-zeros**
Specifies to use all zeros in the test pattern.
- all-ones**
Specifies to use all ones in the test pattern.
- crc-enable**
Generates a CRC checksum.
- Default** all-zeros

low-priority-defect

Syntax

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

Context

- config>service>vpls>mesh-sdp>eth-cfm>mep (not supported in access-uplink operating mode)
- config>service>vpls>sap>eth-cfm>mep
- config>service>vpls>spoke-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

Default

macRemErrXcon

Values	allDef	DefRDICCM, DefMACstatus, DefRemoteCCM, DefErrorCCM, and DefXconCCM
	macRemErrXcon	Only DefMACstatus, DefRemoteCCM, Def ErrorCCM, and DefXconCCM
	remErrXcon	Only DefRemoteCCM, DefErrorCCM, and Def XconCCM
	errXcon	Only DefErrorCCM and DefXconCCM
	xcon	Only DefXconCCM; or
	noXcon	No defects DefXcon or lower are to be reported

mac-address

Syntax

mac-address *mac-address*

no mac-address

Context

config>service>vpls>mesh-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

config>service>vpls>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the MAC address of the MEP.

The **no** form of this command reverts the MAC address of the MEP back to that of the port (if the MEP is on a SAP) or the bridge (if the MEP is on a spoke).

Parameters

mac-address

Specifies the MAC address of the MEP.

Values 6-byte mac-address in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx of the MEP. Must be unicast. Using the all zeros address is equivalent to the no form of this command.

one-way-delay-threshold

Syntax

one-way-delay-threshold *seconds*

Context

config>service>vpls>mesh-sdp>eth-cfm>mep (not supported in access-uplink operating mode)

config>service>vpls>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables/disables eth-test functionality on MEP.

Parameters

seconds

Specifies the one way delay threshold, in seconds.

Values 0..600

Default 3

limit-mac-move

Syntax

limit-mac-move [blockable | non-blockable]

no limit-mac-move

Context

config>service>vpls>spoke-sdp (not supported in access-uplink operating mode)

config>service>vpls>mesh-sdp (not supported in access-uplink operating mode)

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command indicates whether the mac-move agent, when enabled using the **config service vpls mac-move** or **config service epipe mac-move** command, limits the MAC relearn (move) rate on this SAP.

Default

blockable

Parameters

blockable

Specifies that the agent monitors the MAC relearn rate on the SAP, and it blocks it when the relearn rate is exceeded.

non-blockable

Specifies that this SAP is not blocked and another blockable SAP is blocked instead.

mac-pinning

Syntax

[no] **mac-pinning**

Context

config>service>vpls>sap

config>service>vpls>spoke-sdp (not supported in access-uplink operating mode)

config>service>vpls>mesh-sdp (not supported in access-uplink operating mode)

config>service>pw-template (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command disables relearning of MAC addresses on other mesh SDPs within the VPLS.

The MAC address remains attached to a specific Mesh for duration of its age-timer.

The age of the MAC address entry in the FIB is set by the age timer. If mac-aging is disabled on a specific VPLS service, any MAC address learned on a mesh with mac-pinning enabled remains in the FIB on this mesh forever. Every event that otherwise results in relearning is logged (MAC address; original - mesh SDP; new - mesh SDP).

Default

MAC pinning is not enabled by default.

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 (not supported in access-uplink operating mode)

config>service>vpls>endpoint (not supported in access-uplink operating mode)

config>service>pw-template (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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 are discarded.

The **no** form of this 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 30719 (7210 SAS-T)
	1 to 61439 (7210 SAS-Mxp)
	1 to 30202 (7210 SAS-Sx)

static-mac

Syntax

[no] static-mac *ieee-mac-address* [create]

Context

- config>service>vpls>sap
- config>service>vpls>mesh-sdp (not supported in access-uplink operating mode)
- config>service>vpls>spoke-sdp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates a local static MAC entry in the Virtual Private LAN Service (VPLS) forwarding database (FDB) associated with the Service Access Point (SAP).

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Local static MAC entries create a permanent MAC address to SAP association in the forwarding database for the VPLS instance so that MAC addresses are not learned on the edge device.

Note that static MAC definitions on one edge device are not propagated to other edge devices participating in the VPLS instance, that is, each edge device has an independent forwarding database for the VPLS.

Only one static MAC entry (local or remote) can be defined per MAC address per VPLS instance.

By default, no static MAC address entries are defined for the SAP.

The **no** form of this command deletes the static MAC entry with the specified MAC address associated with the SAP from the VPLS forwarding database.

Parameters

ieee-mac-address

Specifies the 48-bit MAC address for the static ARP in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

create

Mandatory keyword when specifying a static MAC address.

managed-vlan-list

Syntax

managed-vlan-list

Context

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure VLAN ranges to be managed by a management VPLS. The list indicates, for each SAP, the ranges of associated VLANs that are affected when the SAP changes state.

This command is only valid when the VPLS in which it is entered was created as a management VPLS.

default-sap

Syntax

[no] default-sap

Context

```
config>service>vpls>sap>managed-vlan-list
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command adds a default SAP to the managed VLAN list.

The **no** form of this command removes the default SAP to the managed VLAN list.

range

Syntax

```
[no] range vlan-range
```

Context

```
config>service>vpls>sap>managed-vlan-list
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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.

To modify the range of VLANs, first the new range should be entered and afterwards the old range removed. See [Modifying VPLS service parameters](#).

Parameters

vlan-range

Specifies the VLAN start value and VLAN end value. The end-vlan must be greater than start-vlan. The format is <start-vlan>-<end-vlan>

Values start-vlan: 0 to 4094
 end-vlan: 0 to 4094

5.8.2.1.8 VPLS SAP statistics commands

statistics

Syntax

statistics

Context

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure the counters associated with SAP ingress and egress.

ingress

Syntax

ingress

Context

config>service>vpls>sap>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink operating mode

Description

Commands in this context configure the ingress SAP statistics counters.

counter-mode

Syntax

counter-mode {in-out-profile-count | forward-drop-count}

Context

config>service>vpls>sap>statistics>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink operating mode

Description

This command allows the user to set the counter mode for the counters associated with SAP ingress meters (also known as policers). A pair of counters is available with each meter. These counters count different events based on the counter mode value.



Note:

The counter mode can be changed if an accounting policy is associated with a SAP. If the counter mode is changed, the counters associated with the meter are reset and the counts are cleared. If an accounting policy is in use when the counter-mode is changed, a new record is written into the current accounting file.

Perform the following steps on the specified SAP to ensure that the correct statistics are collected when the counter-mode is changed.

1. Disable writing of accounting records for the SAP by executing the **config service vpls sap no collect-stats** command.
2. Change the counter-mode to the needed option by executing the **config service vpls sap counter-mode {in-out-profile-count | forward-drop-count}** command.
3. Enable writing of accounting records for the SAP by executing the **config service vpls sap collect-stats** command.

Default

in-out-profile-count

Parameters

forward-drop-count

Specifies that one counter counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

in-out-profile-count

Specifies that one counter counts the total in-profile packets and octets received on ingress of a SAP and another counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

drop-count-extra-vlan-tag-pkts

Syntax

[no] drop-count-extra-vlan-tag-pkts

Context

config>service>vpls>sap>statistics>ingress

config>service>vpls>mesh-sdp>statistics>ingress

config>service>vpls>spoke-sdp>statistics>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command associates a counter which enables the counting of extra VLAN-tag dropped packets for the SAP, spoke-SDP, or mesh SDP. A limited amount of such counters are available for use.

The **no** form of this command removes the associated counter.

5.8.2.1.9 VPLS filter and QoS policy commands

egress

Syntax

egress

Context

config>service>vpls>sap

config>service>ies>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure egress filter policies.

If **no** egress filter is defined, no filtering is performed.

ingress

Syntax

ingress

Context

config>service>vpls>sap

config>service>ies>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context 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 [*cir cir-rate*] [*pir pir-rate*]

no agg-rate-limit

Context

config>service>vpls>sap>egress

config>service>epipe>sap>egress

Platforms

7210 SAS-Mxp

Description

This command defines a maximum total rate for all egress queues on a service SAP.

The port scheduler mode should be set to "sap-based" scheduling mode before using this command. The egress port scheduler enforces the aggregate queue rate for the SAP as it distributes its bandwidth to all the SAPs configured on the port. The port scheduler stops distributing bandwidth to member queues when it has detected that the aggregate rate limit has been reached.

A SAP aggregate scheduler is created for each instance of the SAP queues created on each of the member ports of the LAG. For a LAG, the port scheduler-mode configured for the primary port is used for all the member ports of the LAG.

The scheduler mode is specified by the **scheduler-mode** command. To implement the aggregate rate limit, the scheduler mode must be specified as "sap-based". See the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide* for more information about the **scheduler-mode** command.

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

Parameters

cir-rate

Specifies the CIR in kilobits per second.

Values 0 to 10000000

pir-rate

Specifies the PIR in kilobits per second.

Values 1 to 10000000, max

filter

Syntax

filter ip *ip-filter-id*

filter ipv6 *ipv6-filter-id*

filter mac *mac-filter-id*

Context

config>service>vpls>sap>egress

config>service>vpls>sap>ingress

config>service>ies>sap>egress

config>service>ies>sap>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command associates an IP filter policy or MAC filter policy with an ingress or egress SAP or IP interface.

Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria. There are two types of filter policies: IP and MAC. Only one type may be applied to a SAP at a time.

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

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID is not removed from the system.

Special Cases

VPLS

Specifies that both MAC and IP filters are supported on a VPLS service SAP.

Parameters

ip *ip-filter-id*

Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

ipv6 *ipv6-filter-id*

Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

mac *mac-filter-id*

Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters.

Values 1 to 65535

qos

Syntax

qos *policy-id*

qos *policy-id* [**enable-table-classification**]

no qos *policy-id*

Context

config>service>vpls>sap>ingress

config>service>vpls>sap>egress (for 7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP).

QoS ingress 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 is used to associate both ingress and egress QoS policies. The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress, and only allows egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second policy of same or different type replaces the earlier one with the new policy.



Note:

SAP egress QoS policies are only supported on the 7210 SAS-Mxp.

On the 7210 SAS-Mxp (ingress), using the **enable-table-classification** keyword enables the use of IP DSCP tables to assign FC and profile on a per-SAP ingress basis. The match-criteria configured in the service ingress policy, which require CAM resources, are ignored. Only meters from the service ingress policy are used (and the meters still require CAM resources). The IP DSCP classification policy configured in the SAP ingress policy is used to assign FC and profile. The default FC is assigned from the SAP ingress policy.

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

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

Parameters

policy-id

The ingress or egress policy ID to associate with SAP on ingress or egress. The policy ID must already exist.

Values 1 to 65535

enable-table-classification

Enables the use of table-based classification instead of CAM-based classification at SAP ingress. The FC and profile are taken from the IP DSCP classification policy configured in the ingress policy, along with the meters from the SAP ingress policy. Match-criteria entries in the SAP ingress policy are ignored.

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

config>service>vpls>spoke-sdp (not supported in access-uplink operating mode)

```
config>service>vpls>mesh-sdp (not supported in access-uplink operating mode)
config>service>vpls>sap
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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

default accounting policy

Parameters

acct-policy-id

Enter 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>spoke-sdp (not supported in access-uplink operating mode)
```

```
config>service>vpls>mesh-sdp (not supported in access-uplink operating mode)
```

```
config>service>vpls>sap
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables accounting and statistical data collection for either the SAP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the cards. However, the CPU does not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default

no collect-stats

5.8.2.1.10 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

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command binds a VPLS service to an existing Service Distribution Point (SDP). Mesh SDPs bound to a service are logically treated like a single bridge "port" for flooded traffic where flooded traffic received on any mesh SDP on the service is replicated to other "ports" (spoke-SDPs and SAPs) and not transmitted on any mesh SDPs.

This command creates a binding between a service and an SDP. The SDP has an operational state, which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service is down.

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

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

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

Default

no *sdp-id* is bound to a service

Special Cases

VPLS

Specifies that 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 router, an error occurs and the second SDP binding is rejected.

Parameters

sdp-id

Specifies the SDP identifier.

Values 1 to 17407

vc-id

Specifies 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

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 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 VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

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

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

ether

Defines the VC type as Ethernet. The **ethernet** and **vlan** keywords are mutually exclusive. When the VC type is not defined, 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, the default is Ethernet for mesh SDP bindings.

spoke-sdp

Syntax

spoke-sdp *sdp-id[:vc-id]* [**vc-type** {**ether** | **vlan**}] [**create**] [**split-horizon-group** *group-name*] [**use-evpn-default-shg**]

no spoke-sdp *sdp-id[:vc-id]*

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command binds a service to an existing Service Distribution Point (SDP). A spoke-SDP is treated like the equivalent of a traditional bridge "port" on which 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 on which it was received.

The operational state of the SDP determines the SDP state within the service. For example, if the SDP is administratively or operationally down, the SDP for the service is down.

SDPs must be explicitly associated and bound to a service to allow far-end devices to participate in the service. The SDP must already exist in the **config>service>sdp** context before it can be associated with a VPLS service. If the *sdp-id* is not already configured, an error message is generated. If the *sdp-id* exists, a binding between the specific *sdp-id* and service is created.

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

Special Cases

VPLS

Several SDPs can be bound to a VPLS service. Each SDP must use a unique *vc-id*. An error message is generated if two SDP bindings with an identical *vc-id* terminate on the same router. Split-horizon groups can only be created in the scope of a VPLS service.

Parameters

sdp-id

Specifies the SDP identifier.

Values 1 to 17407

vc-id

Specifies the virtual circuit identifier.

Values 1 to 4294967295

create

Mandatory keyword to create a spoke-SDP.

ether

Keyword to define 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

Keyword to define 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.

split-horizon-group group-name

Specifies the name of the split horizon group to which the SDP belongs.

vc-type

Keyword to override 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 VC type. 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. If signaling is enabled, a change of the bindings VC type causes the binding to signal the new VC type to the far end.

VC types are derived in accordance with IETF *draft-martini-l2circuit-trans-mps*.

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

use-evpn-default-shg

Keyword to add the spoke SDP to the default SHG, which causes the spoke SDP to behave as a mesh SDP. See [Note](#) for more information. This keyword is supported only on the 7210 SAS-Sx/S 1/10GE operating in standalone mode.

**Note:**

This option is not blocked in a VPLS service, but it can be configured only for an EVPN-VPLS service. The default SHG is created when EVPN is enabled in the service, and all EVPN bindings are added to it by default.

egress

Syntax

egress

Context

config>service>vpls>mesh-sdp

config>service>vpls>spoke-sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure egress SDP.

ingress**Syntax**

ingress

Context

```
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure ingress SDP.

force-vlan-vc-forwarding**Syntax**

[no] force-vlan-vc-forwarding

Context

```
config>service>epipe>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>pw-template
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command forces vc-vlan-type forwarding in the datapath for spoke/mesh SDPs which have either vc-type. This command is not allowed on vlan-vc-type SDPs.

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

Default

disabled

hash-label

Syntax

hash-label [signal-capability]

no hash-label

Context

config>service>vpls>spoke-sdp

config>service>vpls>mesh-sdp

Platforms

7210 SAS-Mxp, 7210 SAS-R6 IMM-b, 7210 SAS-R12 IMM-b, 7210 SAS-R6 IMM-c, 7210 SAS-R12 IMM-c, and 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC)

Description

This command enables the use of a hash label on a VLL or VPLS service bound to LDP or RSVP SDP, using the autobind mode with the **ldp**, **rsvp-te**, or **mpls** options. When this command is enabled, the ingress datapath is modified such that the result of the hash on the packet header is communicated to the egress datapath for use, as the value of the label field of the hash label. Only the hash-2 parameters are used to compute the hash label, even if the SDP is over a lag (with **load-balancing** set as **hash-1** or **hash-2**) or a port. The egress datapath adds the hash label at the bottom of the stack (BoS) and sets the S-bit to one.



Note:

On the 7210 SAS, the hash label is not used on the local node for ECMP hashing and LAG hashing. It is available for use by LSR nodes, through which the traffic flows and which are capable of using the labels for hashing.

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

The TTL of the hash label is set to a value of 0.

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

- The 7210 local PE inserts the Flow Label Interface Parameters sub-TLV with T=1 and R=1 in the PW ID FEC element in the label mapping message for that spoke-sdp or mesh-sdp.
- If remote PE does not send the Flow Label sub-TLV in the PW ID FEC element, or sends a Flow Label sub-TLV in the PW ID FEC element with T=FALSE and R=FALSE, then the local node disables

the hash label capability. Therefore local PE node does not insert a hash label in user and control plane packets it forwards on the spoke-sdp or mesh-sdp. It also drops user and control plane packets received from remote PE if they include a hash label. Note that the latter may be caused by a remote 7210 PE which does not support the hash-label option, or which has the hash-label option enabled but does not support the signal-capability option, or does support both options but the user did not enable them because of a mis-configuration.

- If remote PE sends Flow Label sub-TLV in the PW ID FEC element with T=TRUE and R=TRUE, then the local PE enables the hash label capability. Therefore local PE inserts a hash label in user and control plane packets it forwards on the spoke-sdp or mesh-sdp. It also accepts user and control plane packets remote PE with or without hash label
 - If the hash-label option was enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire packets received by the local PE have the hash label included. These packets must be dropped. The only way to solve this is to disable the signaling capability option on the local node which results in the insertion of the hash label by both PE nodes.
 - If the hash-label option is not supported or was not enabled on the local configuration of the spoke-sdp or mesh-sdp at the remote PE, the pseudowire received by the local PE does not have the hash label included.

The user can enable or disable the signal-capability option in CLI as needed. When doing so, the router must withdraw the label it sent to its peer and send a new label mapping message with the new value of the F bit in the flow label interface parameters sub-TLV of the PW ID FEC element.



Note:

- This feature is supported only for VLL and VPLS services. It is not supported for VPRN services. It is also not supported on multicast packets forwarded using RSVP P2MP LSP or mLDP LSP in both the base router instance and in the multicast VPN (mVPN) instance.
- In 7x50 and possibly other vendor implementations, to allow applications where the egress LER infers the presence of the hash label implicitly from the value of the label, the Most Significant Bit (MSB) of the result of the hash is set before copying into the Hash Label. This means that the value of the hash label is always in the range [524,288 - 1,048,575] and does not overlap with the signaled/static LSP and signaled/static service label ranges. This also guarantees that the hash label does not match a value in the reserved label range. This is not supported on 7210 for service traffic (for MPLS OAM traffic the MSB bit is set). That is, 7210 SAS devices do not set the MSB bit in the hash label value for service traffic. Therefore, user must ensure that both the ends are correctly configured to either process hash labels or disable it.

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

Default

no hash-label

Parameters

signal-capability

Enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes.

vc-label

Syntax

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

Context

config>service>vpls>mesh-sdp>egress

config>service>vpls>spoke-sdp>egress

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the egress VC label.

Parameters

vc-label

Specifies the VC egress value that indicates a specific connection.

Values 16 to 1048575

vc-label

Syntax

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

Context

config>service>vpls>mesh-sdp>ingress

config>service>vpls>spoke-sdp>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the ingress VC label.

Parameters

vc-label

Specifies the VC ingress value that 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>vpls>spoke-sdp

config>service>vpls>mesh-sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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.

5.8.2.1.11 SAP IGMP-snooping commands

igmp-snooping

Syntax

igmp-snooping

Context

```
config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp (not supported in access-uplink operating mode)
config>service>vpls>mesh-sdp (not supported in access-uplink operating mode)
config>service>pw-template (not supported in access-uplink operating mode)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure Internet Group Management Protocol (IGMP) snooping.

disable-router-alert-check

Syntax

[no] disable-router-alert-check

Context

```
config>service>vpls>sap>igmp-snooping
```

Platforms

Supported on 7210 SAS platforms operating in access-uplink mode, or in an R-VPLS on 7210 SAS platforms operating in network mode.

Description

This command enables the IGMP router alert check option.



Note:

The **disable-router-alert-check** command is not supported in a VPLS on 7210 SAS platforms operating in network mode.

The **no** form of this command disables the router alert check.

description

Syntax

description *description-string*
no description

Context

```
config>service>vpls>igmp-snooping>mvr
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates a text description stored in the configuration file for a configuration context. The description command associates a text string with a configuration context to help identify the content in the configuration file.

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

Default

no description

Parameters

string

Specifies the description character string. Allowed values are any string up to 80 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

fast-leave

Syntax

[no] fast-leave

Context

```
config>service>vpls>sap>igmp-snooping
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables fast leave. When IGMP fast leave processing is enabled, the 7210 SAS immediately removes a SAP or SDP from the multicast group when it detects an IGMP "leave" on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a 'leave' from the forwarding table without first sending out group-specific queries to the SAP or SDP, and therefore speeds up the process of changing channels ('zapping').

Fast leave should only be enabled when there is a single receiver present on the SAP or SDP. When fast leave is enabled, the configured last-member-query-interval value is ignored.

Default

no fast-leave

from-vpls**Syntax**

from-vpls *service-id*

no from-vpls

Context

config>service>vpls>sap>igmp-snooping>mvr

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the VPLS and R-VPLS service from which multicast traffic is copied upon receipt of an IGMP join request. IGMP snooping must be enabled on the MVR VPLS and MVR R-VPLS service.

Default

no from-vpls

Parameters

service-id

Specifies the MVR VPLS from which multicast channels should be copied into this SAP.

Values *service-id*: 1 to 2147483648

group**Syntax**

[no] **group** *grp-address*

Context

config>service>vpls>sap>igmp-snooping>static

config>service>vpls>spoke-sdp>igmp-snooping>static (not supported in access-uplink operating mode)

config>service>vpls>mesh-sdp>igmp-snooping>static (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command adds a static multicast group as a (*, g). When a static IGMP group is added, multicast data for that (*,g) is forwarded to the specific SAP or SDP without receiving any membership report from a host.



Note:

Only SAPs are supported in an R-VPLS. SDPs are not supported in an R-VPLS.

Parameters

grp-address

Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

```
group-policy
```

Syntax

```
group-policy policy-name
```

```
no group-policy
```

Context

```
config>service>vpls>igmp-snooping>mvr
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command identifies a filter policy of multicast groups to be applied to this VPLS entity. The sources of the multicast traffic must be a member of the VPLS.

The **no** form of this command removes the policy association from the VPLS configuration.

Default

```
no group policy
```

Parameters

policy-name

Specifies the group policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Routing policies are configured in the **config>router>policy-options** context. The router policy must be defined before it can be imported.

import

Syntax

import *policy-name*

no import

Context

config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (not supported in access-uplink operating mode)

config> service>vpls>mesh-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>pw-template>igmp-snooping (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP or SDP at any time.

The **no** form of this command removes the policy association from the SAP or SDP.

Default

no import

Parameters

policy-name

Specifies the import policy name. Values can be string up to 32 characters of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. These policies are configured in the **config>router> policy-options** context. The router policy must be defined before it can be imported.

last-member-query-interval

Syntax

last-member-query-interval *tenths-of-seconds*

no last-member-query-interval

Context

config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>vpls>mesh-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>pw-template>igmp-snooping (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the maximum response time used in group-specific queries sent in response to 'leave' messages, and is also the amount of time between 2 consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. The configured last-member-query-interval is ignored when fast-leave is enabled on the SAP or SDP.

Default

10

Parameters

seconds

Specifies 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>spoke-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>vpls>mesh-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>pw-template>igmp-snooping (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command defines the maximum number of multicast groups that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

Default

no max-num-groups

Parameters

max-num-groups

Specifies the maximum number of groups that can be joined on this SAP or SDP.

Values	For VPLS (SAP, mesh SDP, and spoke-SDP):
	1 to 2043
	For R-VPLS:
	1 to 1000

max-num-sources

Syntax

max-num-sources *max-num-sources*
no max-num-sources

Context

config>service>vpls>sap>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command defines the maximum number of multicast sources allowed per group that can be joined on this SAP. If the node receives an IGMP join message that would exceed the configured number of sources, the request is ignored.



Note:
The **max-num-sources** command is applicable only in the context of R-VPLS service. It cannot be used in the context of VPLS service.

The **no** form of this command disables checking the number of sources.

Default

no max-num-sources

Parameters

max-num-sources

Specifies the maximum number of multicast sources per group that can be joined on this SAP.

Values 1 to 2043

mrouter-port

Syntax

[no] mrouter-port

Context

config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>vpls>mesh-sdp>igmp-snooping (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies whether a multicast router is attached behind this SAP or SDP.

Configuring a SAP or SDP as an mrouter-port has a double effect. First, all multicast traffic received on another SAP or SDP is copied to this SAP or SDP. Second, IGMP reports generated by the system as a result of someone 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. Even though the other multicast router (non-querier) stops sending IGMP queries, it still receives reports to keep its multicast trees up to date. To support this, the **mrouter-port** command should be enabled on all SAPs or SDPs connecting to a multicast router.



Note: The IGMP version to be used for the reports (v1 or v2) can only be determined after an initial query has been received. Until then, no reports are sent on the SAP or SDP, even if **mrouter-port** is enabled.

If the **send-queries** command is enabled on this SAP, the mrouter-port parameter cannot be set.

Default

no mrouter-port

mvr

Syntax

mvr

Context

config>service>vpls>igmp-snooping

```
config>service>vpls>sap>igmp-snooping
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure Multicast VPLS Registration (MVR) parameters.

query-interval

Syntax

query-interval *seconds*

no query-interval

Context

```
config>service>vpls>igmp-snooping
```

```
config>service>vpls>sap>igmp-snooping
```

```
config>service>vpls>spoke-sdp>igmp-snooping (not supported in access-uplink operating mode)
```

```
config>service>vpls>mesh-sdp>igmp-snooping (not supported in access-uplink operating mode)
```

```
config>service>pw-template>igmp-snooping (not supported in access-uplink operating mode)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the IGMP query interval. If the **send-queries** command is enabled, this parameter specifies the interval between two consecutive general queries sent by the system on this SAP or SDP. The configured query-interval must be greater than the configured query-response-interval. If send-queries is not enabled on this SAP or SDP, the configured query-interval value is ignored.

Default

125

Parameters

seconds

Specifies the time interval, in seconds, that the router transmits general host-query messages.

Values 2 to 1024

query-src-ip

Syntax

query-src-ip *ip-address*

no query-src-ip

Context

config>service>vpls>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the IP source address used in IGMP queries.

query-response-interval

Syntax

query-response-interval *seconds*

Context

config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>vpls>mesh-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>pw-template>igmp-snooping (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the IGMP query response interval. If the **send-queries** command is enabled, this parameter specifies the maximum response time advertised in IGMP queries.

The configured query-response-interval must be smaller than the configured query-interval.

If send-queries is not enabled on this SAP or SDP, the configured query-response-interval value is ignored.

Default

10

Parameters***seconds***

Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 to 1023

robust-count**Syntax**

robust-count *robust-count*

no robust-count

Context

config>service>vpls>igmp-snooping

config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>vpls>mesh-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>pw-template>igmp-snooping (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

If the **send-queries** command is enabled, this parameter allows tuning for the expected packet loss on a SAP or SDP. The robust-count variable allows tuning for the expected packet loss on a subnet and is comparable to a retry count. If this SAP or SDP is expected to be 'lossy', this parameter may be increased. IGMP snooping on this SAP or SDP is robust to (robust-count-1) packet losses.

If send-queries is not enabled, this parameter is ignored.

Default

2

Parameters***robust-count***

Specifies the robust count for the SAP or SDP.

Values config>service>vpls>sap>igmp-snooping: 2 to 7
 config>service>vpls>igmp-snooping: 1 to 255
 config>service>vpls>spoke->sdp>igmp-snooping: 2 to 7
 config>service>vpls>mesh-sdp>igmp-snooping: 2 to 7

report-src-ip

Syntax

report-src-ip-address

no report-src-ip

Context

config>service>vpls>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This parameter specifies the source IP address used when generating IGMP reports. According the IGMPv3 standard, a zero source address is allowed in sending IGMP reports. However, for interoperability with some multicast routers, the source IP address of IGMP group reports can be configured using this command.

Default

0.0.0.0

Parameters

ip-address

Specifies the source IP source address in transmitted IGMP reports.

precedence

Syntax

precedence *precedence-value* | **primary**

no precedence

Context

config>service>vpls>spoke-sdp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the spoke-SDP precedence.

Default

4

Parameters***precedence-value***

Specifies the spoke-SDP precedence.

Values 0 to 4***primary***

Specifies that the precedence is primary.

propagate-mac-flush**Syntax****[no] propagate-mac-flush****Context**

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies whether MAC flush messages received from the specific LDP are propagated to all spoke and mesh SDPs within the context of this VPLS service. The propagation follows the split-horizon principle and any datapath blocking to avoid the looping of these messages.

Default

no propagate-mac-flush

send-queries**Syntax****[no] send-queries****Context**

config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>vpls>mesh-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>pw-template>igmp-snooping (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies whether to send IGMP general query messages on the SAP or SDP.

When **send-queries** is configured, all type of queries generate ourselves are of the configured version. If a report of a version higher than the configured version is received, the report gets dropped and a new wrong version counter gets incremented. If send-queries is not configured, the version command has no effect. The version used is the version of the querier.

Default

no send-queries

static

Syntax

static

Context

config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (not supported in access-uplink operating mode)

config>service>vpls>mesh-sdp>igmp-snooping (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables access to the context to configure static group addresses. Static group addresses can be configured on a SAP or SDP. When present either as a (*, g) entry, multicast packets matching the configuration is forwarded even if no join message was registered for the specific group.

source

Syntax

source *ip-address*

no source *ip-address*

Context

config>service>vpls>sap>igmp-snooping>static>group

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command adds a static (s,g) entry to allow multicast traffic for the corresponding multicast group from the specified source.

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



Note:

The **source** command is supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode. For 7210 SAS platforms operating in the network mode, the **source** command must be executed within the context of an R-VPLS.

The **source** command cannot be used within the context of a VPLS.

Default

no source

starg

Syntax

[no] **starg**

Context

config>service>vpls>sap>igmp-snooping>static>group

config>service>vpls>spoke-sdp>igmp-snooping>static>group (not supported in access-uplink operating mode)

config>service>vpls>mesh-sdp>igmp-snooping>static>group (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command adds a static (*,g) entry to allow multicast traffic for the corresponding multicast group from any source. This command can only be enabled if no existing source addresses for this group are specified.

The **no** form of this 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>mesh-sdp>igmp-snooping (not supported in access-uplink operating mode)
config>service>vpls>spoke-sdp>igmp-snooping (not supported in access-uplink operating mode)
config>service>vpls>mesh-sdp>snooping>static (not supported in access-uplink operating mode)
config>service>pw-template>igmp-snooping (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies the version of IGMP which is running on this SAP or SDP. This object can be used to configure a router capable of running either value. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN.

When the **send-query** command is configured, all query types generated locally are of the configured **version**. If a report of a version higher than the configured version is received, the report gets dropped and the “wrong version” counter is incremented.

If the **send-query** command is not configured, the **version** command has no effect. The version used on that SAP is the version of the querier.



Note:
IGMPv3 is only supported on 7210 SAS platforms operating in access-uplink mode, or in an R-VPLS on 7210 SAS platforms operating in network mode.

Parameters

version

Specifies the IGMP version.

Values 1 or 2 (in network mode for VPLS services)
 1, 2, or 3 (access-uplink mode)
 1, 2, 3 (for R-VPLS, only in network mode)

to-sap

Syntax

to-sap *sap-id*

no to-sap

Context

config>service>vpls>sap>igmp-snooping>mvr

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the SAP to which the multicast data needs to be copied.

In some scenarios, the multicast traffic should not be copied from the MVR VPLS or MVR R-VPLS to the SAP on which the IGMP message was received (standard MVR behavior) but to another SAP.

Default

no to-sap

Parameters

sap-id

Specifies the SAP to which multicast channels should be copied.

5.8.2.1.12 Routed VPLS commands

allow-ip-int-bind

Syntax

[no] **allow-ip-int-bind**

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

The **allow-ip-int-bind** command that sets a flag on the VPLS or I-VPLS service that enables the ability to attach an IES or VPRN IP interface to the VPLS service to make the VPLS service routable. When the **allow-ip-int-bind** command is not enabled, the VPLS service cannot be attached to an IP interface.

VPLS Configuration Constraints for Enabling allow-ip-int-bind

When attempting to set the allow-ip-int-bind VPLS flag, the system first checks to see if the correct configuration constraints exist for the VPLS service and the network ports. In Release 8.0 the following VPLS features must be disabled or not configured for the allow-ip-int-bind flag to set:

- SAP ingress QoS policies applied to the VPLS SAPs cannot have MAC match criteria defined
- The VPLS service type cannot be B-VPLS or M-VPLS, and it cannot be an I-VPLS service bound to a B-VPLS context

When the VPLS allow-ip-int-bind flag is set on a VPLS service, the preceding features cannot be enabled on the VPLS service.

VPLS Service Name Bound to IP Interface without the allow-ip-int-bind Flag Set

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-bind flag has not been set on the VPLS service context, the system attempt to resolve the service name between the VPLS service and the IP interface fails. After the allow-ip-int-bind flag is successfully set on the VPLS service, either the service name on the VPLS service must be removed and reapplied, or the IP interface must be reinitialized using the **shutdown /no shutdown** commands. This causes the system to reattempt the name resolution process between the IP interface and the VPLS service.

The **no** form of this command resets the allow-ip-int-bind flag on the VPLS service. If the VPLS service currently has an IP interface from an IES or VPRN service attached, the **no allow-ip-int-bind** command fails. When the allow-ip-int-bind flag is reset on the VPLS service, the configuration and hardware restrictions associated with setting the flag are removed. The port network mode hardware restrictions are also removed.

5.8.2.2 VPLS show commands

egress-label

Syntax

egress-label *egress-label1* [*egress-label2*]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays service information using the range of egress labels.

If only the mandatory *egress-label1* parameter is specified, only services using the specified label are displayed.

If both *egress-label1* and *egress-label2* parameters are specified, the services using the range of labels *X* where *egress-label1* ≤ *X* ≤ *egress-label2* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters

- egress-label1***

Specifies the starting egress label value for which to display services using the label range. If only *egress-label1* is specified, services only using *egress-label1* are displayed.

Values 0, 2049 to 131071
- egress-label2***

Specifies the ending egress label value for which to display services using the label range.

Values 2049 to 131071

Default The *egress-label1* value.

fdb-info

Syntax

fdb-info

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Displays global FDB usage information.

Output

The following output is an example of global FDB usage information, and [Table 54: Output fields: FDB information](#) describes the output fields.

Sample output

```
A:7210-SASE# show service fdb-info
=====
Forwarding Database(FDB) Information
=====
Service Id      : 1                Mac Move       : Disabled
Mac Move Rate   : 2                Mac Move Timeout : 10
Table Size      : 8191             Total Count    : 675
```

Learned Count	: 675	Static Count	: 0
Local Age	: 60		
High WaterMark	: 5%	Low Watermark	: 1%
Mac Learning	: Enabl	Discard Unknown	: Dsabl
Mac Aging	: Enabl	Relearn Only	: False
Service Id	: 2	Mac Move	: Disabled
Mac Move Rate	: 2	Mac Move Timeout	: 10
Table Size	: 8191	Total Count	: 0
Learned Count	: 0	Static Count	: 0
Local Age	: 80		
High WaterMark	: 10%	Low Watermark	: 2%
Mac Learning	: Enabl	Discard Unknown	: Dsabl
Mac Aging	: Enabl	Relearn Only	: False
Service Id	: 3	Mac Move	: Disabled
Mac Move Rate	: 2	Mac Move Timeout	: 10
Table Size	: 8191	Total Count	: 675
Learned Count	: 675	Static Count	: 0
Local Age	: 100		
High WaterMark	: 15%	Low Watermark	: 3%
Mac Learning	: Enabl	Discard Unknown	: Dsabl
Mac Aging	: Enabl	Relearn Only	: False
Service Id	: 4	Mac Move	: Disabled
Mac Move Rate	: 2	Mac Move Timeout	: 10
Table Size	: 8191	Total Count	: 0
Learned Count	: 0	Static Count	: 0
Local Age	: 120		
High WaterMark	: 20%	Low Watermark	: 4%
Mac Learning	: Enabl	Discard Unknown	: Dsabl
Mac Aging	: Enabl	Relearn Only	: False
Service Id	: 5	Mac Move	: Disabled
Mac Move Rate	: 2	Mac Move Timeout	: 10
Table Size	: 8191	Total Count	: 0
Learned Count	: 0	Static Count	: 0
Local Age	: 600		
High WaterMark	: 25%	Low Watermark	: 5%
Mac Learning	: Enabl	Discard Unknown	: Dsabl
Mac Aging	: Enabl	Relearn Only	: False
Service Id	: 6	Mac Move	: Disabled
Mac Move Rate	: 2	Mac Move Timeout	: 10
Table Size	: 8191	Total Count	: 675
Learned Count	: 675	Static Count	: 0
Local Age	: 86400		
High WaterMark	: 30%	Low Watermark	: 10%
Mac Learning	: Enabl	Discard Unknown	: Dsabl
Mac Aging	: Enabl	Relearn Only	: False

Total Service FDBs : 6
 Total FDB Configured Size : 49146
 Total FDB Entries In Use : 2025

=====

A:7210-SASE#

Table 54: Output fields: FDB information

Label	Description
Service ID	The value that identifies a service.
Mac Move	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. The rate is computed as the maximum number of relearns allowed in a 5 second interval. The default rate of 10 relearns per second corresponds to 50 relearns in a 5 second period.
Mac Move Timeout	The time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled. A value of zero indicates that the SAP is not 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.
Remote Age	The number of seconds used to age out FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs.
Local Age	The seconds used to age out FDB entries learned on local SAPs.
High WaterMark	The utilization of the FDB table of this service at which a 'table full' alarm is raised by the agent.
Low WaterMark	The utilization of the FDB table of this service at which a 'table full' alarm is cleared by the agent.
Mac Learning	Whether the MAC learning process is enabled in this service.
Discard Unknown	Whether frames received with an unknown destination MAC are discarded in this service.
MAC Aging	Whether the MAC aging process is enabled in this service.

Label	Description
MAC Pinning	Whether MAC pinning is enabled in this service.
Relearn Only	When enabled, indicates that either the FDB table of this service is full or that the maximum system-wide number of MACs supported by the agent has been reached, and therefore MAC learning is temporary 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.

fdb-mac

Syntax

fdb-mac *ieee-address* [**expiry**]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays the FDB entry for a specific MAC address.

Parameters

ieee-address

Specifies 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 entry information for a specific MAC address, and [Table 55: Output fields: FDB MAC](#) describes the output fields.

Sample output

```
*A:ALA-12# show service fdb-mac 00:99:00:00:00:00
=====
Services Using Forwarding Database Mac 00:99:00:00:00:00
=====
ServId  MAC                               Source-Identifier      Type/
Age Last Change
-----
1       00:99:00:00:00:00                 sap:1/2/7:0           Static
=====
*A:ALA-12#
```

Table 55: Output fields: FDB MAC

Label	Description
Service ID	The service ID number.
MAC	The specified MAC address
Source-Identifier	The location where the MAC is defined.
Type/Age	Static - FDB entries created by management
	Learned - dynamic entries created by the learning process
	OAM - entries created by the OAM process
	H - host, the entry added by the system for a static configured subscriber host
	D or DHCP - DHCP-installed MAC. Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease
	P - indicates the MAC is protected by the MAC protection feature

ingress-label

Syntax

ingress-label start-label [end-label]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 the range of labels X where *start-label* <= X <= *end-label* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters

start-label

Specifies 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

Specifies 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 table describes show service ingress-label output fields.

Sample output

Table 56: Output fields: ingress label

Label	Description
Svc ID	The service identifier.
SDP Id	The SDP identifier.
Type	Indicates whether the SDP is spoke.
I.Lbl	The ingress label used by the far-end device to send packets to this device in this service by the SDP.
E.Lbl	The egress label used by this device to send packets to the far-end device in this service by the SDP.
Number of Bindings Found	The number of SDP bindings within the label range specified.

sap-using

Syntax

```
sap-using interface [ip-address | ip-int-name]  
sap-using [ingress | egress] filter filter-id  
sap-using [sap sap-id]  
sap-using [ingress] qos-policy qos-policy-id
```

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays SAP information.
If no optional parameters are specified, the command displays a summary of all defined SAPs.
The optional parameters restrict output to only SAPs matching the specified properties.

Parameters

- ingress**
Specifies matching an ingress policy.
- egress**
Specifies matching an egress policy.
- filter *filter-id***
The ingress or egress filter policy ID for which to display matching SAPs.
Values 1 to 65535

sap-id
Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

Output

The following output is an example of service SAP information, and [Table 57: Output fields: SAP-using](#) describes the output fields.

Sample output

```
*A:ALU_SIM2>config>service>vpls# show service sap-using  
=====  
Service Access Points  
=====  
PortId                               SvcId   Ing.   Ing.   Egr.   Adm   Opr
```


		QoS	Fltr	Fltr		
1/1/1:10	1	1	none	1none	Up	Up
1/1/3:500.*	1	1	none	1none	Up	Up
1/1/1:200	200	1	none	1none	Up	Up
1/1/3:100.200	200	1	none	1none	Up	Up
1/1/1:300	300	1	none	1none	Up	Up

Number of SAPs : 5						

*A:ALU_SIM2>config>service>vpls#						

Table 57: Output fields: SAP-using

Label	Description
Port ID	The ID of the access port where the SAP is defined.
Svc ID	The service identifier.
I.QoS	The SAP ingress QoS policy number specified on the ingress SAP.
I.MAC/IP	The MAC or IP filter policy ID applied to the ingress SAP.
Egr. Fltr	The filter policy ID applied to the egress SAP.
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The administrative state of the SAP.
Opr	The actual state of the SAP.

sdp

Syntax

sdp [*sdp-id* | *far-end ip-addr*] [*detail* | *keep-alive-history*]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays information for the SDPs associated with the service.

If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters

sdp-id

Displays only information for the specified SDP ID. An SDP is a logical mechanism that ties a far-end 7210 SAS to a particular service without having to specifically define far end SAPs. Each SDP represents a method to reach a 7210 SAS router.

Values 1 to 17407

Default All SDPs.

far-end ip-addr

Displays only SDPs matching with the specified system IP address of the far-end destination 7210 SAS M router for the Service Distribution Point (SDP) that is the termination point for a service.

Default SDPs with any far-end IP address.

detail

Displays detailed SDP information.

Output

The following table describes the show service sdp output fields.

Sample output

Table 58: Output fields: SDP

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke.
VC Type	The VC type, ether or vlan.
VC Tag	The explicit dot1Q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.)
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	The IP address of the remote end of the MPLS tunnel defined by this SDP.
Delivery	The type of delivery used by the SDP: MPLS.

Label	Description
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the Keepalive process.
Oper State	The operational state of the Keepalive process.
Hello Time	How often the SDP echo request messages are 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.
Hello Msg Len	The length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	The amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	The number of forwarded ingress packets.
I. Dro. Pkts	The number of dropped ingress packets.
E. Fwd. Pkts.	The number of forwarded egress packets.
E. Fwd. Octets	The number of forwarded egress octets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field.

sdp-using

Syntax

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

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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

Specifies the virtual circuit identifier.

Values 1 to 4294967295

far-end ip-address

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

Default Services with any far-end IP address.

Output

The following output is an example of SDP information, and [Table 59: Output fields: SDP-using](#) describes the output fields.

Sample output

```
*A:ALA-1# show service sdp-using 300
=====
Service Destination Point (Sdp Id : 300)
=====
SvcId      SdpId          Type Far End      Opr State I.Label  E.Label
-----
2          300:2          Spok 10.0.0.13    Up        131070   131070
-----
Number of SDPs : 51
-----
*A:ALA-1#
```

Table 59: Output fields: SDP-using

Label	Description
Svc ID	The service identifier.
Sdp ID	The SDP identifier.

Label	Description
Type	The type of SDP: Spoke.
Far End	The far-end address of the SDP.
Oper State	The operational state of the service.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.

service-using

Syntax

service-using [**epipe**] [**vpls**] [**mirror**] [**customer** *customer-id*]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays the services matching certain usage properties. If no optional parameters are specified, all services defined on the system are displayed.

Parameters

epipe

Displays matching Epipe services.

vpls

Displays matching VPLS instances.

mirror

Displays matching mirror services.

customer *customer-id*

Displays services only associated with the specified customer ID.

Values 1 to 2147483647

Default Services associated with a customer.

Output

The following output is an example of service information, and [Table 60: Output fields: service-using](#) describes the output fields.

Sample output

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

*A:ALA-12#

*A:ALA-12# show service service-using epipe
=====
Services [epipe]
=====
ServiceId   Type      Adm    Opr      CustomerId  Last Mgmt Change
-----
6           Epipe     Up     Up        6           09/22/2006 23:05:58
7           Epipe     Up     Up        6           09/22/2006 23:05:58
8           Epipe     Up     Up        3           09/22/2006 23:05:58
103         Epipe     Up     Up        6           09/22/2006 23:05:58
-----
Matching Services : 4
=====

*A:ALA-12#

*A:ALA-14# show service service-using
=====
Services
=====
ServiceId   Type      Adm    Opr      CustomerId  Last Mgmt Change
-----
10          mVPLS     Down   Down     1           10/26/2006 15:44:57
11          mVPLS     Down   Down     1           10/26/2006 15:44:57
100         mVPLS     Up     Up        1           10/26/2006 15:44:57
101         mVPLS     Up     Up        1           10/26/2006 15:44:57
102         mVPLS     Up     Up        1           10/26/2006 15:44:57
-----
Matching Services : 5
-----

*A:ALA-14#

A:Dut-A>config>service# show service service-using
=====
Services
=====
ServiceId   Type      Adm    Opr      CustomerId  Last Mgmt Change
-----
100         mVPLS     Up     Up        1           07/07/2009 14:39:13
101         uVPLS     Up     Up        1           07/07/2009 14:39:13
```

102	uVPLS	Up	Up	1	07/07/2009 14:39:13
103	uVPLS	Up	Up	1	07/07/2009 14:39:13
104	uVPLS	Up	Up	1	07/07/2009 14:39:13
105	uVPLS	Up	Up	1	07/07/2009 14:39:13
201	VPLS	Up	Up	1	07/07/2009 14:39:13
202	VPLS	Up	Up	1	07/07/2009 14:39:13
203	VPLS	Up	Up	1	07/07/2009 14:39:13
204	VPLS	Up	Up	1	07/07/2009 14:39:13
205	VPLS	Up	Up	1	07/07/2009 14:39:13
300	mVPLS	Up	Up	1	07/07/2009 14:39:13
301	uVPLS	Up	Up	1	07/07/2009 14:39:13
302	uVPLS	Up	Up	1	07/07/2009 14:39:13
303	uVPLS	Up	Up	1	07/07/2009 14:39:13
304	uVPLS	Up	Up	1	07/07/2009 14:39:1
305	uVPLS	Up	Up	1	07/07/2009 14:39:1
401	VPLS	Up	Up	1	07/07/2009 14:39:1
402	VPLS	Up	Up	1	07/07/2009 14:39:1
403	VPLS	Up	Up	1	07/07/2009 14:39:1
404	VPLS	Up	Up	1	07/07/2009 14:39:1
405	VPLS	Up	Up	1	07/07/2009 14:39:1
500	mVPLS	Up	Up	1	07/07/2009 14:39:1
511	uVPLS	Up	Up	1	07/07/2009 14:39:1
513	uVPLS	Up	Up	1	07/07/2009 14:39:1
515	uVPLS	Up	Up	1	07/07/2009 14:39:1
517	uVPLS	Up	Up	1	07/07/2009 14:39:1
519	uVPLS	Up	Up	1	07/07/2009 14:39:1
601	VPLS	Up	Up	1	07/07/2009 14:39:1
602	VPLS	Up	Up	1	07/07/2009 14:39:1
603	VPLS	Up	Up	1	07/07/2009 14:39:1
604	VPLS	Up	Up	1	07/07/2009 14:39:1
605	VPLS	Up	Up	1	07/07/2009 14:39:1
701	VPLS	Up	Up	1	07/07/2009 14:39:1
702	VPLS	Up	Up	1	07/07/2009 14:39:1
703	VPLS	Up	Up	1	07/07/2009 14:39:1
704	VPLS	Up	Up	1	07/07/2009 14:39:1
801	VPLS	Up	Up	1	07/07/2009 14:39:1
802	VPLS	Up	Up	1	07/07/2009 14:39:1
803	VPLS	Up	Up	1	07/07/2009 14:39:1
804	VPLS	Up	Up	1	07/07/2009 14:39:1
805	VPLS	Up	Up	1	07/07/2009 14:39:1
901	VPLS	Up	Up	1	07/07/2009 14:39:1
902	VPLS	Up	Up	1	07/07/2009 14:39:1
903	VPLS	Up	Up	1	07/07/2009 14:39:1
904	VPLS	Up	Up	1	07/07/2009 14:39:1
905	VPLS	Up	Up	1	07/07/2009 14:39:1
906	VPLS	Up	Up	1	07/07/2009 14:39:1
907	VPLS	Up	Up	1	07/07/2009 14:39:1
908	VPLS	Up	Up	1	07/07/2009 14:39:1
909	VPLS	Up	Up	1	07/07/2009 14:39:1
910	VPLS	Up	Up	1	07/07/2009 14:39:1
1101	Epip	Up	Up	1	07/07/2009 14:39:1
1102	Epip	Up	Up	1	07/07/2009 14:39:1
1103	Epip	Up	Up	1	07/07/2009 14:39:1
1104	Epip	Up	Up	1	07/07/2009 14:39:1
1105	Epip	Up	Up	1	07/07/2009 14:39:1
1501	Epip	Up	Up	1	07/07/2009 14:39:1
1502	Epip	Up	Up	1	07/07/2009 14:39:1
1503	Epip	Up	Up	1	07/07/2009 14:39:1
1504	Epip	Up	Up	1	07/07/2009 14:39:1
1505	Epip	Up	Up	1	07/07/2009 14:39:1
2001	Mirror	Up	Up	1	07/07/2009 14:39:1
2002	Mirror	Up	Up	1	07/07/2009 14:39:1
2011	Epip	Up	Up	1	07/07/2009 14:39:1
2012	VPLS	Up	Up	1	07/07/2009 14:39:1

```

3000      mVPLS      Up      Up      1      07/07/2009 14:39:1
4001      VPLS      Up      Up      1      07/07/2009 14:39:1
4002      VPLS      Up      Up      1      07/07/2009 14:39:1
-----
Matching Services : 69
=====
A:Dut-A>config>service#

```

Table 60: Output fields: service-using

Label	Description
Service Id	The service identifier.
Type	The service type configured for the service ID.
Adm	The administrative state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

id

Syntax

id *service-id*

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays information for a particular service-id.

Parameters

service-id

Specifies the unique service identification number that identifies the service in the service domain.

Values service-id: 1 to 214748364
 svc-name: A string up to 64 characters.

- all**
Displays more information about the service.
- base**
Displays basic service information.
- endpoint**
Displays service endpoint information.
- fdb**
Displays FDB entries.
- labels**
Displays labels being used by this service.
- mstp-configuration**
Displays MSTP information.
- sap**
Displays SAPs associated with the service.
- sdp**
Displays SDPs associated with the service.
- stp**
Displays STP information.

all

Syntax
all

Context
show>service>id

Platforms
Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description
This command displays more information for all aspects of the service.

Output
The following outputs are examples of detailed service information, and [Table 61: Output fields: service ID all](#) describes the output fields.

- [Sample output](#), [Sample output for 7210 SAS-T in access-uplink mode](#), [Table 61: Output fields: service ID all](#).

Sample output

```

A:Dut-A>config>service# show service id 305 all
=====
Service Detailed Information
=====
Service Id       : 305                Vpn Id           : 305
Service Type     : uVPLS
Description      : Default tls description for service id 305
Customer Id      : 1
Last Status Change: 07/07/2009 14:39:57
Last Mgmt Change : 07/07/2009 14:39:14
Admin State      : Up                 Oper State        : Up
MTU              : 1514
MTU Check        : Disabled
SAP Count        : 2                  SDP Bind Count    : 4
Send Flush on Fail: Disabled
Uplink Type      : MPLS
Propagate MacFlush: Disabled
-----
Service Destination Points(SDPs)
-----
Sdp Id 1217:305  -(10.20.1.2)
-----
Description      : Default sdp description
SDP Id           : 1217:305            Type              : Spoke
VC Type          : Ether              VC Tag            : n/a
Admin Path MTU   : 0                  Oper Path MTU     : 9186
Far End          : 10.20.1.2          Delivery          : MPLS

Admin State      : Up                 Oper State        : Up
Acct. Pol        : None               Collect Stats     : Disabled
Managed by Service : 300             Prune State      : Not Pruned
Managed by Spoke : 1217:300
Ingress Label    : 130506             Egress Label     : 130516
Admin ControlWord : Not Preferred     Oper ControlWord  : False
Last Status Change: 07/07/2009 18:49:40 Signaling        : TLDP
Last Mgmt Change : 07/07/2009 14:39:14 Force Vlan-Vc    : Disabled
Last Mgmt Change : 07/07/2009 14:39:14
Flags            : None
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
MAC Aging        : Enabled
L2PT Termination : Disabled          BPDU Translation  : Disabled
MAC Pinning      : Disabled
Ignore Standby Sig : False            Block On Mesh Fail: False

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

Statistics        :
I. Fwd. Pkts.     : 13601             I. Fwd. Octs.     : 10676338
E. Fwd. Pkts.     : 65165676         E. Fwd. Octets    : 39462444830

Associated LSP LIST :
Lsp Name          : A_B_17
Admin State       : Up                Oper State        : Up

```

```

Time Since Last Tr*: 05h24m26s
-----
Stp Service Destination Point specifics
-----
Mac Move           : Blockable
Stp Admin State    : Down
Core Connectivity   : Down
Port Role          : N/A
Port Number        : 2049
Port Path Cost     : 10
Admin Edge         : Disabled
Link Type          : Pt-pt
Root Guard         : Disabled
Last BPDUs from    : N/A
Designated Bridge  : N/A

Stp Oper State     : Down
Port State         : Forwarding
Port Priority       : 128
Auto Edge          : Enabled
Oper Edge          : N/A
BPDU Encap         : Dot1d
Active Protocol    : N/A
Designated Port Id : 0

Fwd Transitions    : 0
Cfg BPDUs rcvd     : 0
TCN BPDUs rcvd     : 0
RST BPDUs rcvd     : 0
Bad BPDUs rcvd     : 0
Cfg BPDUs tx       : 0
TCN BPDUs tx       : 0
RST BPDUs tx       : 0
-----
Sdp Id 1317:305  -(10.20.1.3)
-----
Description       : Default sdp description
SDP Id           : 1317:305
VC Type          : Ether
Admin Path MTU    : 0
Far End          : 10.20.1.3
Type             : Spoke
VC Tag           : n/a
Oper Path MTU     : 9186
Delivery         : MPLS

Admin State       : Up
Acct. Pol         : None
Managed by Service : 300
Managed by Spoke : 1317:300
Ingress Label     : 130454
Admin ControlWord : Not Preferred
Last Status Change : 07/07/2009 18:49:43
Last Mgmt Change  : 07/07/2009 14:39:14
Last Mgmt Change  : 07/07/2009 14:39:14
Flags            : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Max Nbr of MAC Addr : No Limit
Learned MAC Addr  : 0
Oper State        : Up
Collect Stats     : Disabled
Prune State       : Not Pruned
Egress Label      : 130591
Oper ControlWord  : False
Signaling         : TLDP
Force Vlan-Vc     : Disabled

Total MAC Addr    : 0
Static MAC Addr   : 0

MAC Learning      : Enabled
MAC Aging         : Enabled
L2PT Termination  : Disabled
MAC Pinning       : Disabled
Discard Unkwn Srce : Disabled
BPDU Translation  : Disabled

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

Statistics        :
I. Fwd. Pkts.     : 10100
E. Fwd. Pkts.     : 65466629
I. Fwd. Octs.     : 7178960
E. Fwd. Octets    : 39665246044

Associated LSP LIST :
Lsp Name          : A_C_17
Admin State       : Up
Time Since Last Tr*: 05h24m23s
Oper State        : Up
-----
Stp Service Destination Point specifics
-----

```

```

-----
Mac Move           : Blockable
Stp Admin State    : Down
Core Connectivity   : Down
Port Role          : N/A
Port Number        : 2050
Port Path Cost     : 10
Admin Edge         : Disabled
Link Type          : Pt-pt
Root Guard         : Disabled
Last BPDU from     : N/A
Designated Bridge  : N/A
Stp Oper State     : Down
Port State         : Forwarding
Port Priority       : 128
Auto Edge          : Enabled
Oper Edge          : N/A
BPDU Encap         : Dot1d
Active Protocol    : N/A
Designated Port Id: 0

Fwd Transitions    : 0
Cfg BPDUs rcvd     : 0
TCN BPDUs rcvd     : 0
RST BPDUs rcvd     : 0
Bad BPDUs rcvd     : 0
Cfg BPDUs tx       : 0
TCN BPDUs tx       : 0
RST BPDUs tx       : 0
-----
Sdp Id 1417:305   -(10.20.1.4)
-----
Description       : Default sdp description
SDP Id           : 1417:305
VC Type          : Ether
Admin Path MTU    : 0
Far End          : 10.20.1.4
Type             : Spoke
VC Tag           : n/a
Oper Path MTU     : 9186
Delivery         : MPLS

Admin State       : Up
Acct. Pol        : None
Managed by Service : 300
Managed by Spoke : 1417:300
Ingress Label     : 130428
Admin ControlWord : Not Preferred
Last Status Change : 07/07/2009 18:13:42
Last Mgmt Change  : 07/07/2009 14:39:14
Last Mgmt Change  : 07/07/2009 14:39:14
Flags            : None
Peer Pw Bits     : None
Peer Fault Ip     : None
Max Nbr of MAC Addr: No Limit
Learned MAC Addr  : 250
Oper State        : Up
Collect Stats     : Disabled
Prune State       : Not Pruned
Egress Label      : 131015
Oper ControlWord  : False
Signaling         : TLDP
Force Vlan-Vc     : Disabled

Total MAC Addr    : 250
Static MAC Addr   : 0

MAC Learning      : Enabled
MAC Aging         : Enabled
L2PT Termination  : Disabled
MAC Pinning       : Disabled
Discard Unkwn Srce: Disabled
BPDU Translation  : Disabled

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

Statistics        :
I. Fwd. Pkts.     : 97516328
E. Fwd. Pkts.     : 166191635
I. Fwd. Octs.     : 47531982212
E. Fwd. Octets    : 67215031404

Associated LSP LIST :
Lsp Name          : A_D_17
Admin State        : Up
Time Since Last Tr*: 09h33m18s
Oper State         : Up
-----
Stp Service Destination Point specifics
-----
Mac Move           : Blockable
Stp Admin State    : Down
Stp Oper State     : Down

```

Core Connectivity	: Down	Port State	: Forwarding
Port Role	: N/A	Port Priority	: 128
Port Number	: 2051	Auto Edge	: Enabled
Port Path Cost	: 10	Oper Edge	: N/A
Admin Edge	: Disabled	BPDU Encap	: Dot1d
Link Type	: Pt-pt	Active Protocol	: N/A
Root Guard	: Disabled		
Last BPDU from	: N/A	Designated Port Id:	0
Designated Bridge	: N/A		
Fwd Transitions	: 1	Bad BPDUs rcvd	: 0
Cfg BPDUs rcvd	: 0	Cfg BPDUs tx	: 0
TCN BPDUs rcvd	: 0	TCN BPDUs tx	: 0
RST BPDUs rcvd	: 0	RST BPDUs tx	: 0

Sdp Id 1617:305 -(10.20.1.6)

Description	: Default sdp description		
SDP Id	: 1617:305	Type	: Spoke
VC Type	: Ether	VC Tag	: n/a
Admin Path MTU	: 0	Oper Path MTU	: 9186
Far End	: 10.20.1.6	Delivery	: MPLS
Admin State	: Up	Oper State	: Up
Acct. Pol	: None	Collect Stats	: Disabled
Managed by Service	: 300	Prune State	: Pruned
Managed by Spoke	: 1617:300		
Ingress Label	: 131060	Egress Label	: 130843
Admin ControlWord	: Not Preferred	Oper ControlWord	: False
Last Status Change	: 07/07/2009 14:40:52	Signaling	: TLDP
Last Mgmt Change	: 07/07/2009 14:39:14	Force Vlan-Vc	: Disabled
Last Mgmt Change	: 07/07/2009 14:39:14		
Flags	: None		
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
MAC Aging	: Enabled		
L2PT Termination	: Disabled	BPDU Translation	: Disabled
MAC Pinning	: Disabled		
KeepAlive Information	:		
Admin State	: Enabled	Oper State	: Alive
Hello Time	: 10	Hello Msg Len	: 0
Max Drop Count	: 3	Hold Down Time	: 10
Statistics	:		
I. Fwd. Pkts.	: 12889	I. Fwd. Octs.	: 6000654
E. Fwd. Pkts.	: 11999	E. Fwd. Octets	: 5208494
Associated LSP LIST	:		
Lsp Name	: A_F_17		
Admin State	: Up	Oper State	: Up
Time Since Last Tr*	: 09h33m18s		

Stp Service Destination Point specifics

Mac Move	: Blockable		
Stp Admin State	: Down	Stp Oper State	: Down
Core Connectivity	: Down		
Port Role	: N/A	Port State	: Discarding
Port Number	: 2052	Port Priority	: 128

```

Port Path Cost      : 10
Admin Edge          : Disabled
Link Type           : Pt-pt
Root Guard          : Disabled
Last BPDUs from     : N/A
Designated Bridge   : N/A
Auto Edge           : Enabled
Oper Edge           : N/A
BPDU Encap          : Dot1d
Active Protocol     : N/A
Designated Port Id : 0

Fwd Transitions     : 0
Cfg BPDUs rcvd      : 0
TCN BPDUs rcvd      : 0
RST BPDUs rcvd      : 0
Bad BPDUs rcvd      : 0
Cfg BPDUs tx         : 0
TCN BPDUs tx         : 0
RST BPDUs tx         : 0
-----
Number of SDPs      : 4
-----
Service Access Points
-----
SAP 1/1/16:305
-----
Service Id          : 305
SAP                  : 1/1/16:305
Dot1Q Ethertype     : 0x8100
Description          : Default sap description for service id 305
Encap                : q-tag
QinQ Ethertype      : 0x8100

Admin State         : Up
Flags               : None
Last Status Change  : 07/07/2009 14:39:57
Last Mgmt Change    : 07/07/2009 14:39:14
Max Nbr of MAC Addr : No Limit
Learned MAC Addr    : 0
Admin MTU           : 9212
Ingress qos-policy  : 10
Ingr IP Fltr-Id     : n/a
Ingr Mac Fltr-Id    : 305
tod-suite           : None
Egr Agg Rate Limit  : max
Mac Learning        : Enabled
Mac Aging           : Enabled
L2PT Termination    : Disabled
Oper State          : Up
Total MAC Addr      : 0
Static MAC Addr     : 0
Oper MTU            : 9212
Egr IP Fltr-Id      : n/a
Egr Mac Fltr-Id     : n/a

Discard Unkwn Srce : Disabled
Mac Pinning         : Disabled
BPDU Translation    : Disabled

Acct. Pol           : None
Collect Stats       : Disabled
-----
Stp Service Access Point specifics
-----
Mac Move             : Blockable
Stp Admin State      : Up
Core Connectivity    : Down
Port Role            : Designated
Port Number          : 2048
Port Path Cost       : 10
Admin Edge           : Disabled
Link Type            : Pt-pt
Root Guard           : Disabled
Last BPDUs from      : 80:04.00:0a:1b:2c:3d:4e
CIST Desig Bridge    : This Bridge
Stp Oper State       : Up
Port State           : Forwarding
Port Priority         : 128
Auto Edge            : Enabled
Oper Edge            : False
BPDU Encap           : Dot1d
Active Protocol      : Rstp
Designated Port      : 34816

Forward transitions: 5
Cfg BPDUs rcvd      : 0
TCN BPDUs rcvd      : 0
RST BPDUs rcvd      : 29
MST BPDUs rcvd      : 0
Bad BPDUs rcvd      : 0
Cfg BPDUs tx         : 0
TCN BPDUs tx         : 0
RST BPDUs tx         : 17610
MST BPDUs tx         : 0
-----
Sap Statistics
-----
Packets              Octets

```

```

Ingress Stats:          66655          39685976
Egress Stats:           65864342       38651746348
-----
Sap per Meter stats
-----
                Packets          Octets

Ingress Meter 1 (Unicast)
For. InProf      : 0              0
For. OutProf     : 0              0

Ingress Meter 2 (Unicast)
For. InProf      : 0              0
For. OutProf     : 0              0

Ingress Meter 3 (Unicast)
For. InProf      : 0              0
For. OutProf     : 0              0

Ingress Meter 4 (Unicast)
For. InProf      : 11406          4291328
For. OutProf     : 12575          4325376

Ingress Meter 11 (Multipoint)
For. InProf      : 0              0
For. OutProf     : 0              0

Ingress Meter 12 (Multipoint)
For. InProf      : 3108           3108000
For. OutProf     : 2235           2235000

Ingress Meter 13 (Multipoint)
For. InProf      : 0              0
For. OutProf     : 0              0

Ingress Meter 14 (Multipoint)
For. InProf      : 8772           5166272
For. OutProf     : 4840           3072000
-----
SAP lag-4:305
-----
Service Id      : 305
SAP             : lag-4:305          Encap           : q-tag
Description     : Default sap description for service id 305

Admin State     : Up                 Oper State      : Up
Flags          : None
Last Status Change : 07/07/2009 14:39:57
Last Mgmt Change  : 07/07/2009 14:39:14
Max Nbr of MAC Addr: No Limit        Total MAC Addr  : 125
Learned MAC Addr : 125               Static MAC Addr : 0
Admin MTU       : 9212               Oper MTU        : 9212
Ingress qos-policy : 10
Ingr IP Fltr-Id  : n/a               Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : 305               Egr Mac Fltr-Id : n/a
tod-suite       : None
Egr Agg Rate Limit : max
Mac Learning     : Enabled            Discard Unkwn Srce: Disabled
Mac Aging        : Enabled            Mac Pinning      : Disabled
L2PT Termination : Disabled           BPDU Translation : Disabled

Acct. Pol       : None                Collect Stats    : Disabled
-----
Stp Service Access Point specifics

```

```

-----
Mac Move           : Blockable
Stp Admin State    : Up
Core Connectivity   : Down
Port Role          : Designated
Port Number        : 2000
Port Path Cost     : 10
Admin Edge         : Disabled
Link Type          : Pt-pt
Root Guard         : Disabled
Last BPDU from     : 80:04:00:0a:1b:2c:3d:4e
CIST Desig Bridge  : This Bridge
Stp Oper State     : Up
Port State         : Forwarding
Port Priority      : 128
Auto Edge         : Enabled
Oper Edge         : False
BPDU Encap        : Dot1d
Active Protocol    : Rstp
Designated Port    : 34768

Forward transitions: 4
Cfg BPDUs rcvd    : 0
TCN BPDUs rcvd    : 0
RST BPDUs rcvd    : 23
MST BPDUs rcvd    : 0
Bad BPDUs rcvd    : 0
Cfg BPDUs tx      : 0
TCN BPDUs tx      : 0
RST BPDUs tx      : 17578
MST BPDUs tx      : 0
-----

```

Sap Statistics

```

-----
Ingress Stats:      Packets      Octets
                   190824363    87464904956
Egress Stats:       97572636    45409567760
-----

```

Sap per Meter stats

```

-----
Ingress Meter 1 (Unicast)
For. InProf         : 0
For. OutProf        : 0
Octets              : 0

Ingress Meter 2 (Unicast)
For. InProf         : 0
For. OutProf        : 0
Octets              : 0

Ingress Meter 3 (Unicast)
For. InProf         : 0
For. OutProf        : 0
Octets              : 0

Ingress Meter 4 (Unicast)
For. InProf         : 56963244
For. OutProf        : 59512115
Octets              : 20851041536
Octets              : 19403302144

Ingress Meter 11 (Multipoint)
For. InProf         : 0
For. OutProf        : 0
Octets              : 0

Ingress Meter 12 (Multipoint)
For. InProf         : 12922550
For. OutProf        : 9452800
Octets              : 12922550000
Octets              : 9452800000

Ingress Meter 13 (Multipoint)
For. InProf         : 0
For. OutProf        : 0
Octets              : 0

Ingress Meter 14 (Multipoint)
For. InProf         : 43268112
For. OutProf        : 6788456
Octets              : 21539479708
Octets              : 2546422464
-----

```

VPLS Spanning Tree Information

```

-----
VPLS oper state    : Up
Stp Admin State    : Up
Core Connectivity   : Down
Stp Oper State     : Up
-----

```



```

Mode : Rstp Vcp Active Prot. : N/A

Bridge Id : 00:0d:00:20:ab:cd:00:01 Bridge Instance Id: 13
Bridge Priority : 0 Tx Hold Count : 6
Topology Change : Inactive Bridge Hello Time : 2
Last Top. Change : 0d 05:21:37 Bridge Max Age : 20
Top. Change Count : 5 Bridge Fwd Delay : 15
MST region revision: 0 Bridge max hops : 20
MST region name :

Root Bridge : This Bridge
Primary Bridge : N/A

Root Path Cost : 0 Root Forward Delay: 15
Rcvd Hello Time : 2 Root Max Age : 20
Root Priority : 13 Root Port : N/A
-----
Forwarding Database specifics
-----
Service Id : 305 Mac Move : Disabled
Mac Move Rate : 2 Mac Move Timeout : 10
Table Size : 500 Total Count : 375
Learned Count : 375 Static Count : 0
Remote Age : 60 Local Age : 60
High WaterMark : 95% Low Watermark : 90%
Mac Learning : Enabl Discard Unknown : Dsabl
Mac Aging : Enabl Relearn Only : False
=====
A:Dut-A>config>service#

```

Sample output for 7210 SAS-T in access-uplink mode

```

*A:SAS-T>show>service>id# all

=====
Service Detailed Information
=====

Service Id : 1 Vpn Id : 0
Service Type : VPLS
Description : (Not Specified)
Customer Id : 1
Last Status Change: 04/29/2001 06:59:15
Last Mgmt Change : 04/28/2001 03:03:03
Admin State : Up Oper State : Up
MTU : 1514
MTU Check : Enabled
SAP Count : 2 SDP Bind Count : 0
Snd Flush on Fail : Disabled
Uplink Type: : MPLS

-----
Service Destination Points(SDPs)
-----
No Matching Entries

-----
Service Access Points
-----

-----
SAP 1/1/1:10.*
-----
Service Id : 1
SAP : 1/1/1:10.* Encap : qinq

```

```

QinQ Dot1p      : Default
Description     : (Not Specified)
Admin State    : Up
Flags          : None
Last Status Change : 04/29/2001 06:59:15
Last Mgmt Change  : 04/28/2001 03:09:30
Dot1Q Ethertype : 0x8100
QinQ Ethertype  : 0x8100

Max Nbr of MAC Addr: No Limit
Learned MAC Addr  : 0
Admin MTU        : 1522
Ingr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : 1
tod-suite        : None
Mac Learning     : Enabled
Mac Aging        : Enabled
BPDU Translation : Disabled
L2PT Termination : Disabled

Total MAC Addr   : 0
Static MAC Addr  : 0
Oper MTU         : 1522
Egr IP Fltr-Id  : n/a
Egr Mac Fltr-Id : n/a

Discard Unkwn Srce: Disabled
Mac Pinning       : Disabled

Acct. Pol        : None
Collect Stats    : Disabled

```

Stp Service Access Point specifics

```

Stp Admin State : Up
Core Connectivity : Down
Port Role       : N/A
Port Number     : 2048
Port Path Cost  : 10
Admin Edge      : Disabled
Link Type       : Pt-pt
Root Guard      : Disabled
Last BPDU from  : N/A
CIST Desig Bridge : N/A
Stp Oper State  : Down
Port State      : Forwarding
Port Priority    : 128
Auto Edge       : Enabled
Oper Edge       : N/A
BPDU Encap      : Dot1d
Active Protocol : N/A
Designated Port : N/A

Forward transitions: 0
Cfg BPDUs rcvd    : 0
TCN BPDUs rcvd    : 0
RST BPDUs rcvd    : 0
MST BPDUs rcvd    : 0
Bad BPDUs rcvd    : 0
Cfg BPDUs tx      : 0
TCN BPDUs tx      : 0
RST BPDUs tx      : 0
MST BPDUs tx      : 0

```

ARP host

```

Admin State      : outOfService
Host Limit       : 1
Min Auth Interval : 15 minutes

```

QoS

Ingress qos-policy : 1

Aggregate Policer

```

rate          : n/a
burst         : n/a

```

Ingress QoS Classifier Usage

```

Classifiers Allocated: 4
Classifiers Used      : 2
Meters Allocated     : 2
Meters Used          : 2

```

Sap Statistics

Ingress Stats:	Packets	Octets
Egress Stats:	142761481188	9707780720784
Extra-Tag Drop Stats:	0	0
	n/a	n/a

Sap per Meter stats		

	Packets	Octets
Ingress Meter 1 (Unicast)		
For. InProf	: 17	1162
For. OutProf	: 0	0
Ingress Meter 11 (Multipoint)		
For. InProf	: 61	4148
For. OutProf	: 142761547917	9707785259394

SAP 1/1/2:10.*		

Service Id	: 1	
SAP	: 1/1/2:10.*	Encap : qinq
QinQ Dot1p	: Default	
Description	: (Not Specified)	
Admin State	: Up	Oper State : Up
Flags	: None	
Last Status Change	: 04/29/2001 07:03:49	
Last Mgmt Change	: 04/28/2001 03:02:15	
Dot1Q Ethertype	: 0x8100	QinQ Ethertype : 0x8100
Max Nbr of MAC Addr:	No Limit	Total MAC Addr : 0
Learned MAC Addr	: 0	Static MAC Addr : 0
Admin MTU	: 1522	Oper MTU : 1522
Ingr IP Fltr-Id	: n/a	Egr IP Fltr-Id : n/a
Ingr Mac Fltr-Id	: n/a	Egr Mac Fltr-Id : n/a
tod-suite	: None	
Mac Learning	: Enabled	Discard Unkwn Srce: Disabled
Mac Aging	: Enabled	Mac Pinning : Disabled
BPDU Translation	: Disabled	
L2PT Termination	: Disabled	
Acct. Pol	: None	Collect Stats : Disabled

Stp Service Access Point specifics		

Stp Admin State	: Up	Stp Oper State : Down
Core Connectivity	: Down	
Port Role	: N/A	Port State : Forwarding
Port Number	: 2049	Port Priority : 128
Port Path Cost	: 10	Auto Edge : Enabled
Admin Edge	: Disabled	Oper Edge : N/A
Link Type	: Pt-pt	BPDU Encap : Dot1d
Root Guard	: Disabled	Active Protocol : N/A
Last BPDU from	: N/A	
CIST Desig Bridge	: N/A	Designated Port : N/A
Forward transitions:	0	Bad BPDUs rcvd : 0
Cfg BPDUs rcvd	: 0	Cfg BPDUs tx : 0
TCN BPDUs rcvd	: 0	TCN BPDUs tx : 0
RST BPDUs rcvd	: 0	RST BPDUs tx : 0
MST BPDUs rcvd	: 0	MST BPDUs tx : 0

```

-----
ARP host
-----
Admin State      : outOfService
Host Limit       : 1                      Min Auth Interval : 15 minutes
-----

QOS
-----
Ingress qos-policy : 1
-----
Aggregate Policer
-----
rate             : n/a                     burst             : n/a
-----
Ingress QoS Classifier Usage
-----
Classifiers Allocated: 4                  Meters Allocated  : 2
Classifiers Used     : 2                  Meters Used       : 2
-----
Sap Statistics
-----
Ingress Stats:      Packets      Octets
                   0             0
Egress Stats:      535194841     36393249188
Extra-Tag Drop Stats: n/a         n/a
-----
Sap per Meter stats
-----
                   Packets      Octets

Ingress Meter 1 (Unicast)
For. InProf        : 0          0
For. OutProf       : 0          0

Ingress Meter 11 (Multipoint)
For. InProf        : 0          0
For. OutProf       : 0          0
-----

VPLS Spanning Tree Information
-----
VPLS oper state    : Up           Core Connectivity : Down
Stp Admin State    : Down        Stp Oper State    : Down
Mode               : Rstp        Vcp Active Prot.  : N/A

Bridge Id          : 80:00:00:25:ba:02:ea: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

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0             Root Forward Delay: 15
Rcvd Hello Time    : 2            Root Max Age       : 20
Root Priority       : 32768        Root Port          : N/A
-----
Forwarding Database specifics
-----
Service Id        : 1             Mac Move           : Disabled
Mac Move Rate     : 2             Mac Move Timeout   : 10

```

```

Mac Move Retries : 3
Table Size       : 250
Learned Count    : 0
Remote Age       : 900
High Watermark   : 95%
Mac Learning     : Enabled
Mac Aging        : Enabled
Total Count      : 0
Static Count     : 0
Local Age        : 300
Low Watermark    : 90%
Discard Unknown  : Disabled
Relearn Only     : False

```

```

-----
Service Endpoints
-----

```

```

No Endpoints found.
=====

```

```

*A:SAS-T>show>service>id#

```

Table 61: Output fields: service ID all

Label	Description
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.
Service Type	The type of service.
SDP Id	The SDP identifier.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service.
Split Horizon Group	Name of the split horizon group for this service.
Description	Description of the split horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.
SDP Id	The SDP identifier.
Type	Indicates whether this service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.

Label	Description
Delivery	The type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	The IP address of the remote end of the MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Hello Time	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 amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field.
Number of SDPs	The total number SDPs applied to this service ID.
Service Id	The service identifier.
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Encap Value	The value of the label used to identify this SAP on the access port.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.

Label	Description
Last Changed	The date and time of the last change.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The SAP ingress QoS policy ID.
Ingress Filter-Id	The SAP ingress filter policy ID.
Egress Filter-Id	The SAP egress filter policy ID.
Multi Svc Site	Indicates the multi-service site that the SAP is a member.
Ingress sched-policy	Indicates the ingress QoS scheduler for the SAP.
Egress sched-policy	Indicates the egress QoS scheduler for the SAP.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
Ingress Stats	The number of received packets/octets for this SAP.
Egress Stats	The number of packets/octets forwarded out of this SAP.
Ingress Meter 1	The index of the ingress QoS meter of this SAP.
High priority offered	The packets or octets count of the high priority traffic for the SAP.
For.InProf	The packets or octets count of the in-profile forwarded traffic for the SAP.
For.OutProf	The number of out of profile traffic packets/octets forwarded.
Managed by Service	The service-id of the management VPLS managing this SAP.
Managed by MSTI	The MST instance inside the management VPLS managing this SAP.
Last BPDU from	The bridge ID of the sender of the last BPDU received on this SAP.
Managed by SAP	The sap-id inside the management VPLS managing this SAP.
Prune state	The STP state inherited from the management VPLS.
Managed by Service	The service-id of the management VPLS managing this spoke-SDP.

Label	Description
Last BPDU from	The bridge ID of the sender of the last BPDU received on this SAP.
Managed by Spoke	The sap-id inside the management VPLS managing this spoke-SDP.
Prune state	The STP state inherited from the management VPLS.

arp

Syntax

arp [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays the ARP table for the VPLS instance. The ARP entries for a subscriber interface are displayed uniquely. Each MAC associated with the subscriber interface child group-interfaces is displayed with each subscriber interface ARP entry for easy lookup.

Parameters

ip-address

All IP addresses.

mac ieee-address

Displays only ARP entries in the ARP table with the specified 48-bit MAC address. The MAC address is in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff, where aa, bb, cc, dd, ee and ff are hexadecimal numbers.

Default All MAC addresses.

sap sap-id

Displays SAP information for the specified SAP ID.

interface

Specifies matching service ARP entries associated with the IP interface.

ip-address

Specifies the IP address of the interface for which to display matching ARP entries.

Values a.b.c.d

ip-int-name

Specifies the IP interface name for which to display matching ARPs.

Output

The following table describes show service-id ARP output fields.

Sample output

Table 62: Output Fields: ARP

Label	Description
IP Address	The IP address
MAC Address	The specified MAC address
	Type Static - FDB entries created by management
	Learned - dynamic entries created by the learning process
	Other - local entries for the IP interfaces created
Expiry	The age of the ARP entry
Interface	The interface applied to the service
SAP	The SAP ID

base**Syntax**

base [msap]

Context

show>service>id

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays basic information about the service ID including service type, description, SAPs and SDP.

Output

The following output is an example of basic service information, and [Table 63: Output fields: base](#) describes the output fields.

Sample output

```

*A:7210SAS# show service id 10 base
=====
Service Basic Information
=====
Service Id      : 10          Vpn Id      : 0
Service Type    : VPLS
Name           : (Not Specified)
Description     : (Not Specified)
Customer Id     : 1
Last Status Change: 02/06/2106 06:28:12
Last Mgmt Change : 01/10/1970 01:55:31
Admin State    : Down        Oper State    : Down
MTU            : Not Applicable Def. Mesh VC Id : 10
SAP Count      : 0
Uplink Type    : L2
SAP Type       : Dot1q Range  Customer vlan: : n/a

-----
Service Access & Destination Points
-----
Identifier      Type      AdmMTU  OprMTU  Adm  Opr
-----
No Matching Entries
=====
*A:7210SAS# show service id 10 base

A:Dut-A# show service id 1 base
=====
Service Basic Information
=====
Service Id : 1 Vpn Id : 0
Service Type : Epipe
Customer Id : 1
Last Status Change: 06/24/2001 00:57:55
Last Mgmt Change : 06/24/2001 00:51:36
Admin State : Up Oper State : Up
MTU : 1514
MTU Check : Disabled
Vc Switching : False
SAP count : 1 SDP Bind Count : 1

-----
Service Access and Destination Points
-----
Identifier Type AdmMTU OprMTU Adm Opr
-----
sap:1/1/21:1 q-tag 1518 1518 Up Up
sdp:1:1 S<100.1.12> n/a 1518 1518 Up Up
=====
A:Dut-A#

```

Table 63: Output fields: base

Label	Description
Service Id	the service identifier
Service Type	the type of service

Label	Description
Description	generic information about the service
Customer Id	the customer identifier
Last Mgmt Change	the date and time of the most recent management-initiated change to this customer
Adm	the administrative state of the service
Oper	the operational state of the service
Mtu	the largest frame size (in octets) that the port can handle
Adm	the largest frame size (in octets) that the SAP can handle
SAP Count	the number of SAPs defined on the service
SAP Type	the type of SAPs allowed in the service - It also describes the applied processing by the node to the packets received on these SAPs.
Identifier	The service access (SAP)
OprMTU	the actual largest service frame size (in octets) that can be transmitted through this port, without requiring the packet to be fragmented
Opr	the operating state of the SAP

fdb

Syntax

fdb [**sap** *sap-id* [**expiry**]] | [**mac** *ieee-address* [**expiry**]] | [**detail**] [**expiry**]

Context

show>service>id

show>service>fdb-mac

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays FDB entries for a specific MAC address.

Parameters

sap sap-id

Specifies the physical port identifier portion of the SAP. See [Common CLI command descriptions](#) for command syntax.

detail

Displays more information.

expiry

Displays time until MAC is aged out.

Output

The following output is an example of service FDB information, and [Table 64: Output fields: FDB](#) describes the output fields.

Sample output

```
A:Dut-A# show service id 305 fdb
=====
Forwarding Database, Service 305
=====
Service Id       : 305           Mac Move       : Disabled
Mac Move Rate    : 2            Mac Move Timeout : 10
Table Size       : 500          Total Count     : 375
Learned Count    : 375          Static Count    : 0
Remote Age       : 60           Local Age       : 60
High WaterMark   : 95%          Low Watermark   : 90%
Mac Learning     : Enabl        Discard Unknown : Dsabl
Mac Aging        : Enabl        Relearn Only    : False
=====
A:Dut-A#
```

Table 64: Output fields: FDB

Label	Description
ServID	The service ID.
MAC	The associated MAC address.
Mac Move	The administrative state of the MAC movement feature associated with this service.
Primary Factor	A factor for the primary ports defining how many MAC-relearn periods should be used to measure the MAC-relearn rate.
Secondary Factor	A factor for the secondary ports defining how many MAC-relearn periods should be used to measure the MAC-relearn rate.
Mac Move Rate	The maximum rate at which MAC's 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 MAS.

Label	Description
	The rate is computed as the maximum number of relearns allowed in a 5 second interval: for example, the default rate of 2 relearns per second corresponds to 10 relearns in a 5 second period.
Mac Move Timeout	The time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled. A value of zero indicates that the SAP is not 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.
Mac Move Retries	The number of times retries are performed for reenabling 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.
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 utilization of the FDB table of this service at which a table full alarm is raised by the agent.
Low Watermark	The utilization of the FDB table of this service at which a table full alarm is cleared by the agent.
Mac Learning	Whether the MAC learning process is enabled
Discard Unknown	Whether frames received with an unknown destination MAC are discarded.
Mac Aging	Whether the MAC aging process is enabled.

Label	Description
Relearn Only	Displays, that when enabled, either the FDB table of this service is full, or that the maximum system-wide number of MA's supported by the agent has been reached, and therefore MAC learning is temporary 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
	Static - statically configured
Last Change	The time of the most recent state changes.

labels

Syntax

labels

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays the labels being used by the service.

Output

The following output is an example of service label information, and [Table 65: Output fields: labels](#) describes the output fields.

Sample output

```
A:Dut-A# show service id 305 labels
=====
Martini Service Labels
=====
Svc Id      Sdp Binding      Type  I.Lbl      E.Lbl
-----
305         1217:305         Spok  130506     130516
305         1317:305         Spok  130454     130591
305         1417:305         Spok  130428     131015
305         1617:305         Spok  131060     130843
-----
Number of Bound SDPs : 4
=====
A:Dut-A#
```

Table 65: Output fields: labels

Label	Description
Svc Id	The service identifier.
Sdp Id	The SDP identifier.
Type	Whether the SDP is spoke.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP.

l2pt

Syntax

- l2pt disabled
- l2pt [detail]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays Layer 2 Protocol Tunnel (L2-PT) route information associated with this service.

Parameters

disabled

Displays only entries with termination disabled. This helps identify configuration errors.

detail

Displays more information.

Output

The following output is an example of L2PT information, and [Table 66: Output fields: L2PT](#) describes the output fields.

Sample output

```
*A:7210SAS>show>service# id 1 l2pt detail

=====
L2pt details, Service id 1
=====

Service Access Points
-----
SapId                L2pt-termination      Admin Bpdu-translation  Oper Bpdu-translation
-----
1/1/1                stp cdp vtp dtp pagp udld  disabled      disabled
-----
Number of SAPs : 1
=====

L2pt summary, Service id 1
=====
      L2pt-term  L2pt-term  Bpdu-trans  Bpdu-trans  Bpdu-trans  Bpdu-trans
      enabled    disabled   auto        disabled    pvst        stp
-----
SAP's 1          0          0           1           0           0
SDP's  0          0          0           0           0           0
-----
Total 1          0          0           1           0           0
=====
*A:7210SAS>show>service#
```

Table 66: Output fields: L2PT

Label	Description
Service id	Displays the 24-bit (0 to 16777215) service instance identifier for the service
L2pt-term enabled	Indicates whether L2PT-termination or BPDU-translation is in use on this service by at least one SAP or spoke-SDP binding. If in use, at least one of L2PT-termination or BPDU-translation is enabled. When enabled, it is not possible to enable STP on this service.

Label	Description
L2pt-term disabled	Indicates that L2PT-termination is disabled
Bpdu-trans auto	Displays the number of L2PT PDUs that are translated before being sent out on a port or SAP
Bpdu-trans disabled	Indicates that BPDU-translation is disabled
SAPs	Displays the number of SAPs with L2PT or BPDU translation enabled or disabled
SDPs	Displays the number of SDPs with L2PT or BPDU translation enabled or disabled
Total	Displays the column totals of L2PT entities
SapId	Displays the ID of the access point where this SAP is defined
L2pt-termination	Displays whether L2pt termination is enabled or disabled
Admin Bpdu-translation	Displays whether Bpdu translation is administratively enabled or disabled
Oper Bpdu-translation	Displays whether Bpdu translation is operationally enabled or disabled
SAP Id	Displays the SAP ID

mac-move

Syntax

mac-move

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays MAC move related information about the service.

mstp-configuration

Syntax

mstp-configuration

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays the MSTP specific configuration data. This command is only valid on a management VPLS.

Output

The following output is an example of MSTP configuration information, and [Table 67: Output fields: MSTP configuration](#) describes the output fields.

Sample output

```
*A:SASMX>show>service>id# mstp-configuration

=====
Mstp configuration info, Service 5
=====
Region Name      : abc
Region Revision  : 0
MST Max Hops     : 20

=====
vlan to MST instance mapping
=====
Instance  Priority  Vlans mapped
-----
2          0
=====
*A:SASMX>show>service>id#
```

Table 67: Output fields: MSTP configuration

Label	Description
Region Name	The MSTP region name.
Region Revision	The MSTP region revision.
MST Max Hops	The MSTP maximum hops specified.

Label	Description
Instance	The MSTP instance number.
Priority	The MSTP priority.
Vlans mapped	The VLAN range of the MSTP instance.

sap

Syntax

sap *sap-id* **detail**

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays information for the SAPs associated with the service.

If no optional parameters are specified, a summary of all associated SAPs is displayed.

Parameters

sap *sap-id*

Specifies the ID that displays SAPs for the service in the *slot/mdalport[.channel]* form. See [Common CLI command descriptions](#) for command syntax.

detail

Displays more information for the SAP.

Output

The following outputs are examples of service SAP information, and [Table 68: Output fields: service ID SAP](#) describes the output fields.

- [Sample output](#), [Sample output for 7210 SAS-Mxp](#), [Table 68: Output fields: service ID SAP](#).

Sample output

```
A:7210>show>service>id# sap 1/1/1:1 detail
```

```
=====
Service Access Points(SAP)
=====
```

```
Service Id      : 1
SAP             : 1/1/1:1          Encap           : q-tag
Description     : (Not Specified)
Admin State     : Up              Oper State      : Down
```

```

Flags : ServiceAdminDown
Last Status Change : 10/05/2010 07:22:04
Last Mgmt Change : 10/05/2010 07:22:05
Dot1Q Ethertype : 0x8100
QinQ Ethertype : 0x8100

Max Nbr of MAC Addr: No Limit
Learned MAC Addr : 0
Admin MTU : 1518
Ingr IP Fltr-Id : n/a
Ingr Mac Fltr-Id : n/a
tod-suite : None
Mac Learning : Enabled
Mac Aging : Enabled
BPDU Translation : Disabled
L2PT Termination : Disabled

Total MAC Addr : 0
Static MAC Addr : 0
Oper MTU : 1518
Egr IP Fltr-Id : n/a
Egr Mac Fltr-Id : n/a

Discard Unkwn Srce: Disabled
Mac Pinning : Disabled

Acct. Pol : None
Collect Stats : Disabled

```

Stp Service Access Point specifics

```

Stp Admin State : Up
Core Connectivity : Down
Port Role : N/A
Port Number : 2048
Port Path Cost : 10
Admin Edge : Disabled
Link Type : Pt-pt
Root Guard : Disabled
Last BPDU from : N/A
CIST Desig Bridge : N/A

Stp Oper State : Down
Port State : Discarding
Port Priority : 128
Auto Edge : Enabled
Oper Edge : N/A
BPDU Encap : Dot1d
Active Protocol : N/A
Designated Port : N/A

Forward transitions: 0
Cfg BPDUs rcvd : 0
TCN BPDUs rcvd : 0
RST BPDUs rcvd : 0
MST BPDUs rcvd : 0

Bad BPDUs rcvd : 0
Cfg BPDUs tx : 0
TCN BPDUs tx : 0
RST BPDUs tx : 0
MST BPDUs tx : 0

```

ARP host

```

Admin State : outOfService
Host Limit : 1
Min Auth Interval : 15 minutes

```

QoS

```

Ingress qos-policy : 5
Egress qos-policy : 1

```

Aggregate Policer (Not Available)

```

rate : n/a
burst : n/a

```

Ingress QoS Classifier Usage

```

Classifiers Allocated: 256
Classifiers Used : 2

Meters Allocated : 32
Meters Used : 2

```

Sap Statistics

```

Ingress Stats:
Egress Stats:

Packets
0
0

Octets
0
0

```

```

-----
Sap per Meter stats
-----
                Packets                Octets

Ingress Meter 1 (Unicast)
For. InProf      : 0                  0
For. OutProf     : 0                  0

Ingress Meter 11 (Multipoint)
For. InProf      : 0                  0
For. OutProf     : 0                  0
=====
*A:SAS>show>service>id# sap 1/1/1:10.* detail
=====
Service Access Points(SAP)
=====
Service Id       : 1
SAP              : 1/1/1:10.*          Encap           : qinq
QinQ Dot1p      : Default
Description      : (Not Specified)
Admin State      : Up                  Oper State       : Up
Flags            : None
Last Status Change : 04/29/2001 06:59:15
Last Mgmt Change  : 04/28/2001 03:09:30
Dot1Q Ethertype  : 0x8100             QinQ Ethertype   : 0x8100

Max Nbr of MAC Addr: No Limit          Total MAC Addr   : 0
Learned MAC Addr  : 0                  Static MAC Addr  : 0
Admin MTU         : 1522                Oper MTU         : 1522
Ingr IP Fltr-Id   : n/a                 Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id  : 1                   Egr Mac Fltr-Id : n/a
tod-suite         : None
Mac Learning      : Enabled              Discard Unkwn Srce: Disabled
Mac Aging         : Enabled              Mac Pinning      : Disabled
BPDU Translation  : Disabled
L2PT Termination  : Disabled

Acct. Pol        : None                  Collect Stats    : Disabled

-----
Stp Service Access Point specifics
-----
Stp Admin State   : Up                  Stp Oper State   : Down
Core Connectivity : Down
Port Role         : N/A                 Port State       : Forwarding
Port Number       : 2048                Port Priority     : 128
Port Path Cost    : 10                  Auto Edge        : Enabled
Admin Edge        : Disabled             Oper Edge        : N/A
Link Type         : Pt-pt                BPDU Encap       : Dot1d
Root Guard        : Disabled             Active Protocol   : N/A
Last BPDU from    : N/A                 Designated Port   : N/A
CIST Desig Bridge : N/A

Forward transitions: 0                   Bad BPDUs rcvd   : 0
Cfg BPDUs rcvd    : 0                   Cfg BPDUs tx     : 0
TCN BPDUs rcvd    : 0                   TCN BPDUs tx     : 0
RST BPDUs rcvd    : 0                   RST BPDUs tx     : 0
MST BPDUs rcvd    : 0                   MST BPDUs tx     : 0
-----
ARP host

```

```

-----
Admin State      : outOfService
Host Limit      : 1                      Min Auth Interval : 15 minutes
-----
QoS
-----
Ingress qos-policy : 1
-----
Aggregate Policer
-----
rate            : n/a                    burst            : n/a
-----
Ingress QoS Classifier Usage
-----
Classifiers Allocated: 4                Meters Allocated  : 2
Classifiers Used     : 2                Meters Used       : 2
-----
Sap Statistics
-----
Ingress Stats:      Packets              Octets
                    142761481188          9707780720784
Egress Stats:       0                    0
Extra-Tag Drop Stats: n/a                n/a
-----
Sap per Meter stats
-----
                    Packets              Octets

Ingress Meter 1 (Unicast)
For. InProf         : 17                  1162
For. OutProf        : 0                    0

Ingress Meter 11 (Multipoint)
For. InProf         : 61                  4148
For. OutProf        : 142761547917        9707785259394
=====

```

Sample output for 7210 SAS-Mxp

```

*A:Dut-A# show service id 10 sap 5/1/1:800 detail
=====
Service Access Points(SAP)
=====
Service Id       : 10
SAP              : 5/1/1:800              Encap            : q-tag
Description      : (Not Specified)
Admin State      : Up                    Oper State       : Down
Flags           : PortOperDown
Last Status Change : 11/07/2017 04:48:25
Last Mgmt Change  : 11/07/2017 05:02:47
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
BGP IPv4 FlowSpec : Disabled
BGP IPv6 FlowSpec : Disabled
tod-suite        : None
Egr Agg Rate CIR  : 0                    Egr Agg Rate PIR  : max
Limit Unused BW   : Disabled

```

Acct. Pol	: None	Collect Stats	: Disabled
Anti Spoofing	: None	Dynamic Hosts	: Enabled
Oper Group	: (none)	Monitor Oper Grp	: (none)
Host Lockout Plcy	: n/a		
Lag Link Map Prof	: (none)		

QoS

Ingress qos-policy	: 17	Egress qos-policy	: 1
Table-based	: enabled		

Aggregate Policer

Rate	: n/a	Burst	: n/a

Egress Aggregate Meter

Rate	: n/a	Burst	: n/a

Ingress QoS Classifier Usage

Classifiers Allocated:	60	Meters Allocated	: 30
Classifiers Used	: 9	Meters Used	: 8

Sap Statistics

	Packets	Octets
Ingress Stats:	0	0
Egress Stats:	0	0
Ingress Drop Stats:	0	0
Extra-Tag Drop Stats:	n/a	n/a

Sap per Meter stats (in/out counter mode)

	Packets	Octets
Ingress Meter 1		
For. InProf	: 0	0
For. OutProf	: 0	0
Ingress Meter 2		
For. InProf	: 0	0
For. OutProf	: 0	0
Ingress Meter 3		
For. InProf	: 0	0
For. OutProf	: 0	0
Ingress Meter 4		
For. InProf	: 0	0
For. OutProf	: 0	0
Ingress Meter 5		
For. InProf	: 0	0
For. OutProf	: 0	0
Ingress Meter 6		
For. InProf	: 0	0
For. OutProf	: 0	0
Ingress Meter 7		
For. InProf	: 0	0
For. OutProf	: 0	0

```

Ingress Meter 8
For. InProf      : 0
For. OutProf     : 0
=====

```

Table 68: Output fields: service ID SAP

Label	Description
Service Id	The service identifier.
SAP	The SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	An Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operational state of the SAP.
Flags	The conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, L2OperDown, Relearn LimitExceeded, ParentIfAdminDown, TodResourceUnavail, Tod MssResourceUnavail, SapParamMismatch, SapIngressNamed PoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipe RingNode
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.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Whether collect stats is enabled.
SAP per Meter stats	
Ingress Meter	The meter ID.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded.
For. OutProf	The number of out-of-profile packets and octets. (rate above CIR and below PIR) forwarded by the ingress meter.

Label	Description
Ingress TD Profile	The profile ID applied to the ingress SAP.
Egress TD Profile	The profile ID applied to the egress SAP.
Alarm Cell Handling	The indication that OAM cells are being processed.
AAL-5 Encap	The AAL-5 encapsulation type.
Aggregate Policer	rate-indicates the rate of the aggregate policer. burst-indicates the burst-size of the aggregate policer.
Loopback Mode	The Ethernet port loopback mode
Loopback Src Addr	The configured loopback source address
Loopback Dst Addr	The configured loopback destination address
No-svc-port used	The port ID of the port on which no service is configured. This port is used for the port loop back with MAC swap functionality.
Table-based	The use of table-based resource classification: Enabled (table-based) or Disabled (CAM-based)

sdp

Syntax

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

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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

Default All SDPs

far-end ip-addr

Displays only SDPs matching with the specified far-end IP address.

Default SDPs with any far-end IP address.

detail

Displays detailed SDP information.

Output

The following output is an example of service SDP information, and [Table 69: Output fields: service ID SDP](#) describes the output fields.

Sample output

```
A:Dut-A>show>service>id# sdp 1217:305
=====
Service Destination Point (Sdp Id : 1217:305)
=====
SdpId          Type IP address    Adm   Opr      I.Lbl    E.Lbl
-----
1217:305       Spok 10.20.1.2      Up    Up        130506   130516
-----
Number of SDPs : 1
=====
A:Dut-A>show>service>id# sdp 1217:305 detail

A:Dut-A>show>service>id#
=====
Service Destination Point (Sdp Id : 1217:305) Details
-----
Sdp Id 1217:305  -(10.20.1.2)
-----
Description      : Default sdp description
SDP Id           : 1217:305                Type           : Spoke
VC Type          : Ether                  VC Tag         : n/a
Admin Path MTU   : 0                     Oper Path MTU  : 9186
Far End          : 10.20.1.2              Delivery       : MPLS

Admin State      : Up                     Oper State     : Up
Acct. Pol       : None                    Collect Stats  : Disabled
Managed by Service : 300                  Prune State   : Not Pruned
Managed by Spoke : 1217:300
Ingress Label    : 130506                  Egress Label   : 130516
Admin ControlWord : Not Preferred              Oper ControlWord : False
Last Status Change : 07/07/2009 18:49:40          Signaling      : TLDP
Last Mgmt Change  : 07/07/2009 14:39:14        Force Vlan-Vc : Disabled
Last Mgmt Change  : 07/07/2009 14:39:14
Flags           : None
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
MAC Aging        : Enabled
L2PT Termination : Disabled                  BPDU Translation : Disabled
MAC Pinning      : Disabled
```

```

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

Statistics            :
I. Fwd. Pkts.         : 13601                      I. Fwd. Octs.        : 10676338
E. Fwd. Pkts.         : 83776987                  E. Fwd. Octets       : 51589499116

Associated LSP LIST :
Lsp Name              : A_B_17
Admin State           : Up                          Oper State           : Up
Time Since Last Tr*: 08h31m06s

-----
Stp Service Destination Point specifics
-----
Mac Move              : Blockable
Stp Admin State       : Down                        Stp Oper State       : Down
Core Connectivity     : Down
Port Role             : N/A                        Port State           : Forwarding
Port Number           : 2049                      Port Priority        : 128
Port Path Cost        : 10                        Auto Edge            : Enabled
Admin Edge            : Disabled                   Oper Edge            : N/A
Link Type             : Pt-pt                      BPDU Encap           : Dot1d
Root Guard            : Disabled                   Active Protocol      : N/A
Last BPDU from        : N/A                        Designated 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

-----
Number of SDPs : 1
=====
* indicates that the corresponding row element may have been truncated.
A:Dut-A>show>service>id#

```

Table 69: Output fields: service ID SDP

Label	Description
Sdp Id	The SDP identifier.
Type	Whether the SDP is spoke.
VC Type	The VC type: ether, vlan, or vpls.
VC Tag	The explicit dot1Q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.)

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 MPLS tunnel defined by this SDP.
Delivery	The type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The current status of the SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the Keepalive process.
Oper State	The operational state of the Keepalive process.

split-horizon-group

Syntax

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

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays service split horizon groups.

stp

Syntax

stp [detail]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays information for the spanning tree protocol instance for the service.

Parameters

detail

Displays more information.

Output

The following outputs are examples of STP information, and [Table 70: Output fields: STP](#) describes the output fields.

- [Sample output](#), [Sample output for 7210 SAS](#), [Table 70: Output fields: STP](#).

Sample output

```
A:Dut-A>show>service>id# stp
=====
Stp info, Service 305
=====
Bridge Id       : 00:0d.00:20:ab:cd:00:01  Top. Change Count : 5
Root Bridge     : This Bridge              Stp Oper State    : Up
Primary Bridge  : N/A                     Topology Change   : Inactive
Mode            : Rstp                     Last Top. Change  : 0d 08:35:16
Vcp Active Prot.: N/A
Root Port       : N/A                     External RPC      : 0
=====
Stp port info
=====
Sap/Sdp Id      Oper-  Port-  Port-  Port-  Oper-  Link-  Active
                State  Role   State  Num   Edge   Type   Prot.
-----
1/1/16:305      Up     Designated Forward 2048   False Pt-pt Rstp
lag-4:305       Up     Designated Forward 2000   False Pt-pt Rstp
1217:305        Up     N/A     Forward 2049   N/A    Pt-pt N/A
1317:305        Up     N/A     Forward 2050   N/A    Pt-pt N/A
1417:305        Up     N/A     Forward 2051   N/A    Pt-pt N/A
1617:305        Pruned N/A     Discard 2052   N/A    Pt-pt N/A
=====
A:Dut-A>show>service>id#
```

```

A:Dut-A>show>service>id# stp detail
=====
Spanning Tree Information
=====
VPLS Spanning Tree Information
-----
VPLS oper state      : Up                Core Connectivity : Down
Stp Admin State      : Up                Stp Oper State      : Up
Mode                 : Rstp              Vcp Active Prot.    : N/A

Bridge Id            : 00:0d.00:20:ab:cd:00:01  Bridge Instance Id: 13
Bridge Priority       : 0                  Tx Hold Count       : 6
Topology Change      : Inactive            Bridge Hello Time    : 2
Last Top. Change     : 0d 08:35:29         Bridge Max Age       : 20
Top. Change Count    : 5                  Bridge Fwd Delay     : 15
MST region revision  : 0                  Bridge max hops      : 20
MST region name      :

Root Bridge          : This Bridge
Primary Bridge       : N/A

Root Path Cost       : 0                  Root Forward Delay: 15
Rcvd Hello Time      : 2                  Root Max Age        : 20
Root Priority         : 13                 Root Port           : N/A
-----
Spanning Tree Sap/Spoke SDP Specifics
-----
SAP Identifier       : 1/1/16:305          Stp Admin State      : Up
Port Role            : Designated          Port State           : Forwarding
Port Number          : 2048                Port Priority         : 128
Port Path Cost       : 10                  Auto Edge            : Enabled
Admin Edge           : Disabled             Oper Edge            : False
Link Type            : Pt-pt                BPDU Encap          : PVST
Root Guard           : Disabled             Active Protocol      : Rstp
Last BPDU from       : 80:04.00:0a:1b:2c:3d:4e
CIST Desig Bridge    : This Bridge          Designated Port      : 34816
Forward transitions: 5                     Bad BPDUs rcvd       : 0
Cfg BPDUs rcvd       : 0                   Cfg BPDUs tx         : 0
TCN BPDUs rcvd       : 0                   TCN BPDUs tx         : 0
RST BPDUs rcvd       : 29                  RST BPDUs tx         : 23488
MST BPDUs rcvd       : 0                   MST BPDUs tx         : 0

SAP Identifier       : lag-4:305           Stp Admin State      : Up
Port Role            : Designated          Port State           : Forwarding
Port Number          : 2000                Port Priority         : 128
Port Path Cost       : 10                  Auto Edge            : Enabled
Admin Edge           : Disabled             Oper Edge            : False
Link Type            : Pt-pt                BPDU Encap          : Dot1d
Root Guard           : Disabled             Active Protocol      : Rstp
Last BPDU from       : 80:04.00:0a:1b:2c:3d:4e
CIST Desig Bridge    : This Bridge          Designated Port      : 34768
Forward transitions: 4                     Bad BPDUs rcvd       : 0
Cfg BPDUs rcvd       : 0                   Cfg BPDUs tx         : 0
TCN BPDUs rcvd       : 0                   TCN BPDUs tx         : 0
RST BPDUs rcvd       : 23                  RST BPDUs tx         : 23454
MST BPDUs rcvd       : 0                   MST BPDUs tx         : 0

SDP Identifier       : 1217:305           Stp Admin State      : Down
Port Role            : N/A                 Port State           : Forwarding
Port Number          : 2049                Port Priority         : 128
Port Path Cost       : 10                  Auto Edge            : Enabled
Admin Edge           : Disabled             Oper Edge            : N/A
Link Type            : Pt-pt                BPDU Encap          : Dot1d
Root Guard           : Disabled             Active Protocol      : N/A

```

```

Last BPDUs from      : N/A
Designated Bridge    : N/A
Fwd Transitions      : 0
Cfg BPDUs rcvd       : 0
TCN BPDUs rcvd       : 0
RST BPDUs rcvd       : 0

Designated Port Id: 0
Bad BPDUs rcvd      : 0
Cfg BPDUs tx        : 0
TCN BPDUs tx        : 0
RST BPDUs tx        : 0

SDP Identifier       : 1317:305
Port Role            : N/A
Port Number          : 2050
Port Path Cost       : 10
Admin Edge           : Disabled
Link Type            : Pt-pt
Root Guard           : Disabled
Last BPDUs from      : N/A
Designated Bridge    : N/A
Fwd Transitions      : 0
Cfg BPDUs rcvd       : 0
TCN BPDUs rcvd       : 0
RST BPDUs rcvd       : 0

Stp Admin State      : Down
Port State           : Forwarding
Port Priority         : 128
Auto Edge            : Enabled
Oper Edge            : N/A
BPDU Encap           : Dot1d
Active Protocol       : N/A

Designated Port Id: 0
Bad BPDUs rcvd      : 0
Cfg BPDUs tx        : 0
TCN BPDUs tx        : 0
RST BPDUs tx        : 0

SDP Identifier       : 1417:305
Port Role            : N/A
Port Number          : 2051
Port Path Cost       : 10
Admin Edge           : Disabled
Link Type            : Pt-pt
Root Guard           : Disabled
Last BPDUs from      : N/A
Designated Bridge    : N/A
Fwd Transitions      : 1
Cfg BPDUs rcvd       : 0
TCN BPDUs rcvd       : 0
RST BPDUs rcvd       : 0

Stp Admin State      : Down
Port State           : Forwarding
Port Priority         : 128
Auto Edge            : Enabled
Oper Edge            : N/A
BPDU Encap           : Dot1d
Active Protocol       : N/A

Designated Port Id: 0
Bad BPDUs rcvd      : 0
Cfg BPDUs tx        : 0
TCN BPDUs tx        : 0
RST BPDUs tx        : 0

SDP Identifier       : 1617:305
Port Role            : N/A
Port Number          : 2052
Port Path Cost       : 10
Admin Edge           : Disabled
Link Type            : Pt-pt
Root Guard           : Disabled
Last BPDUs from      : N/A
Designated Bridge    : N/A
Fwd Transitions      : 0
Cfg BPDUs rcvd       : 0
TCN BPDUs rcvd       : 0
RST BPDUs rcvd       : 0

Stp Admin State      : Down
Port State           : Discarding
Port Priority         : 128
Auto Edge            : Enabled
Oper Edge            : N/A
BPDU Encap           : Dot1d
Active Protocol       : N/A

Designated Port Id: 0
Bad BPDUs rcvd      : 0
Cfg BPDUs tx        : 0
TCN BPDUs tx        : 0
RST BPDUs tx        : 0
=====
A:Dut-A>show>service>id#

*7210-SAS>show>service>id# stp detail

=====
Spanning Tree Information
=====

-----
VPLS Spanning Tree Information
-----
VPLS oper state      : Up
Stp Admin State       : Up
Mode                  : Mstp
Core Connectivity     : Down
Stp Oper State        : Up
Vcp Active Prot.     : N/A

```

```

Bridge Id       : 80:00:00:25:ba:04:66:a0  Bridge Instance Id: 0
Bridge Priority  : 32768                    Tx Hold Count   : 6
Topology Change : Inactive                  Bridge Hello Time : 2
Last Top. Change : 0d 02:54:16              Bridge Max Age    : 20
Top. Change Count : 27                     Bridge Fwd Delay   : 15

Root Bridge     : 40:00:7c:20:64:ac:ff:63
Primary Bridge  : N/A

Root Path Cost  : 10                        Root Forward Delay: 15
Rcvd Hello Time : 2                        Root Max Age       : 20
Root Priority    : 16384                     Root Port          : 2048

MSTP info for CIST :
Regional Root    : 80:00:7c:20:64:ad:04:5f  Root Port          : 2048
Internal RPC     : 10                       Remaining Hopcount: 19
MSTP info for MSTI 1 :
Regional Root    : This Bridge              Root Port          : N/A
Internal RPC     : 0                       Remaining Hopcount: 20
MSTP info for MSTI 2 :
Regional Root    : 00:02:7c:20:64:ad:04:5f  Root Port          : 2048
Internal RPC     : 10                       Remaining Hopcount: 19

```

Spanning Tree Sap Specifics

```

SAP Identifier   : 1/1/7:0                  Stp Admin State   : Up
Port Role        : Root                     Port State        : Forwarding
Port Number      : 2048                     Port Priority      : 128
Port Path Cost   : 10                       Auto Edge         : Enabled
Admin Edge       : Disabled                  Oper Edge         : False
Link Type        : Pt-pt                     BPDU Encap        : Dot1d
Root Guard       : Disabled                  Active Protocol    : Mstp
Last BPDU from   : 80:00:7c:20:64:ad:04:5f Inside Mst Region : True
CIST Desig Bridge : 80:00:7c:20:64:ad:04:5f Designated Port   : 34816
MSTI 1 Port Prio  : 128                      Port Path Cost    : 10
MSTI 1 Desig Brid : This Bridge               Designated Port   : 34816
MSTI 2 Port Prio  : 128                      Port Path Cost    : 10
MSTI 2 Desig Brid : 00:02:7c:20:64:ad:04:5f Designated Port   : 34816
Forward transitions: 17                      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   : 7310                     MST BPDUs tx      : 7277

SAP Identifier   : 1/1/8:0                  Stp Admin State   : Up
Port Role        : Alternate                 Port State        : Discarding
Port Number      : 2049                     Port Priority      : 128
Port Path Cost   : 10                       Auto Edge         : Enabled
Admin Edge       : Disabled                  Oper Edge         : False
Link Type        : Pt-pt                     BPDU Encap        : Dot1d
Root Guard       : Disabled                  Active Protocol    : Mstp
Last BPDU from   : 80:00:7c:20:64:ad:04:5f Inside Mst Region : True
CIST Desig Bridge : 80:00:7c:20:64:ad:04:5f Designated Port   : 34817
MSTI 1 Port Prio  : 128                      Port Path Cost    : 10
MSTI 1 Desig Brid : This Bridge               Designated Port   : 34817
MSTI 2 Port Prio  : 128                      Port Path Cost    : 10
MSTI 2 Desig Brid : 00:02:7c:20:64:ad:04:5f Designated Port   : 34817
Forward transitions: 14                      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   : 7326                     MST BPDUs tx      : 7307

```



```

SAP Identifier      : 1/1/9:0          Stp Admin State   : Up
Port Role          : Designated       Port State       : Forwarding
Port Number        : 2050             Port Priority     : 128
Port Path Cost     : 10               Auto Edge        : Enabled
Admin Edge         : Disabled          Oper Edge        : True
Link Type          : Pt-pt            BPDU Encap       : Dot1d
Root Guard         : Disabled          Active Protocol   : Mstp
Last BPDU from     : N/A              Inside Mst Region : True
CIST Desig Bridge  : This Bridge       Designated Port   : 34818
MSTI 1 Port Prio   : 128              Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge       Designated Port   : 34818
MSTI 2 Port Prio   : 128              Port Path Cost    : 10
MSTI 2 Desig Brid  : This Bridge       Designated Port   : 34818
Forward transitions: 2                 Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                 Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                 TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                 RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                 MST BPDUs tx      : 7415

SAP Identifier      : 1/1/25:0         Stp Admin State   : Up
Port Role          : Alternate         Port State       : Discarding
Port Number        : 2051             Port Priority     : 128
Port Path Cost     : 10               Auto Edge        : Enabled
Admin Edge         : Disabled          Oper Edge        : False
Link Type          : Pt-pt            BPDU Encap       : Dot1d
Root Guard         : Disabled          Active Protocol   : Mstp
Last BPDU from     : 80:00.7c:20:64:ad:04:5f Inside Mst Region : True
CIST Desig Bridge  : 80:00.7c:20:64:ad:04:5f Designated Port   : 34820
MSTI 1 Port Prio   : 128              Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge       Designated Port   : 34819
MSTI 2 Port Prio   : 128              Port Path Cost    : 10
MSTI 2 Desig Brid  : 00:02.7c:20:64:ad:04:5f Designated Port   : 34820
Forward transitions: 10                 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     : 7329             MST BPDUs tx      : 7303

SAP Identifier      : lag-1:0          Stp Admin State   : Up
Port Role          : Alternate         Port State       : Discarding
Port Number        : 2052             Port Priority     : 128
Port Path Cost     : 10               Auto Edge        : Enabled
Admin Edge         : Disabled          Oper Edge        : False
Link Type          : Pt-pt            BPDU Encap       : Dot1d
Root Guard         : Disabled          Active Protocol   : Mstp
Last BPDU from     : 80:00.7c:20:64:ad:04:5f Inside Mst Region : True
CIST Desig Bridge  : 80:00.7c:20:64:ad:04:5f Designated Port   : 34822
MSTI 1 Port Prio   : 128              Port Path Cost    : 10
MSTI 1 Desig Brid  : This Bridge       Designated Port   : 34820
MSTI 2 Port Prio   : 128              Port Path Cost    : 10
MSTI 2 Desig Brid  : 00:02.7c:20:64:ad:04:5f Designated Port   : 34822
Forward transitions: 11                 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     : 7322             MST BPDUs tx      : 7299
=====

```

Sample output for 7210 SAS

```
*A:SAS>show>service>id# stp mst-instance 2
```

```

=====
MSTP specific info for service 5 MSTI 2
=====
Regional Root      : N/A                      Root Port        : N/A
Internal RPC       : 0                        Remaining Hopcount: 20
=====

MSTP port info for MSTI 2
=====
Sap/Sdp Id        Oper-   Port-   Port-   Port-   Same
                  State    Role    State    Num    Region
-----
No data found.
=====
*A:SAS>show>service>id#

Sample output with MSTP information:

*A:SAS>show>service>id# stp mst-instance 2

=====
MSTP specific info for service 5 MSTI 2
=====
Regional Root      : N/A                      Root Port        : N/A
Internal RPC       : 0                        Remaining Hopcount: 20
=====

MSTP port info for MSTI 2
=====
Sap/Sdp Id        Oper-   Port-   Port-   Port-   Same
                  State    Role    State    Num    Region
-----
No data found.
=====
*A:SAS>show>service>id#

```

Table 70: Output fields: STP

Label	Description
Bridge-id	The MAC address used to identify this bridge in the network.
Bridge fwd delay	How fast a bridge changes its state when moving toward the forwarding state.
Bridge Hello time	The amount of time between the transmission of Configuration BPDUs.
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.
Bridge priority	The priority of the Spanning Tree Protocol instance associated with this service.
Topology change	Whether a topology change is currently in progress.

Label	Description
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.
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-id	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	How fast the root changes its state when moving toward the forwarding state.
Root 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 which provides the lowest cost path from this bridge to the root bridge.
SAP Identifier	The ID of the access port where this SAP is defined.
BPDU encap	The type of encapsulation used on BPDUs sent out and received on this SAP.
Port Number	The value of the port number field which is contained in the least significant 12 bits of the 16-bit port ID associated with this SAP.
Priority	The value of the port priority field which is contained in the most significant 4 bits of the 16-bit port ID associated with this SAP.
Cost	The contribution of this port to the path cost of paths toward the spanning tree root which include this port.
Designated Port	The port identifier of the port on the designated bridge for this port's segment.
Designated Bridge	The bridge identifier of the bridge which this port considers to be the designated bridge for this port's segment.

dhcp

Syntax

dhcp

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context display DHCP information for the specified service.

statistics

Syntax

statistics [**sap** *sap-id*]

statistics [**sdp** *sdp-id:vc-id*] (not supported in access-uplink operating mode)

statistics [**interface** *interface-name*]

Context

show>service>id>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays DHCP statistics information.

Parameters

sap *sap-id*

Specifies the physical port identifier portion of the SAP definition.

sdp-id

Specifies the SDP identifier.

Values 1 to 17407

interface *interface-name*

Displays information for the specified IP interface.

Output

The following output is an example of DHCP statistics information, and [Table 71: Output fields: DHCP statistics](#) describes the output fields.

Sample output

```
*A:7210SAS>show>service>id>dhcp# statistics

=====
DHCP Global Statistics, service 1
=====
Rx Packets                : 416554
Tx Packets                : 206405
Rx Malformed Packets     : 0
Rx Untrusted Packets     : 0
Client Packets Discarded  : 0
Client Packets Relayed    : 221099
Client Packets Snooped    : 0
Client Packets Proxied (RADIUS) : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded  : 0
Server Packets Relayed    : 195455
Server Packets Snooped    : 0
DHCP RELEASEs Spoofed    : 0
DHCP FORCERENEWs Spoofed : 0
=====
*A:7210SAS>show>service>id>dhcp#
```

Table 71: Output fields: DHCP statistics

Label	Description
Received Packets	The number of packets received from the DHCP clients. Includes DHCP packets received from both DHCP client and DHCP server.
Transmitted Packets	The number of packets transmitted to the DHCP clients. Includes DHCP packets transmitted from both DHCP client and DHCP server.
Received Malformed Packets	The number of corrupted/invalid packets received from the DHCP clients. Includes DHCP packets received from both DHCP client and DHCP server
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients. In this case, a frame is dropped because the client sending a DHCP packet with Option 82 filled in before "trust" is set under the DHCP interface command.
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded.
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded.

Label	Description
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Server Packets Discarded	The number of packets received from the DHCP server that were discarded.
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded.
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.

summary

Syntax

summary [**interface** *interface-name*]

Context

show>service>id>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays DHCP configuration summary information.

Parameters

interface *interface-name*

Displays information for the specified IP interface.

Output

The following output is an example of summary DHCP information, and [Table 72: Output fields: DHCP summary](#) describes the output fields.

Sample output

```
A:7210SAS# show service id 1 dhcp summary
DHCP Summary, service 1
```

```
=====
Interface Name      Arp    Used/  Info  Admin
SapId/Sdp          Populate Provided  Option  State
-----
egr_1              No      0/0    Replace Up
i_1                No      0/0    Replace Up
-----
```

```
Interfaces: 2
=====
*A:7210SAS>show>service>id>dhcp#
```

Table 72: Output fields: DHCP summary

Label	Description
Interface Name	The name of the router interface.
Arp Populate	Whether or not ARP populate is enabled. 7210 SAS does not support ARP populate.
Used/Provided	7210 SAS does not maintain lease state.
Info Option	Whether Option 82 processing is enabled on the interface.
Admin State	The administrative state.

dhcp6

Syntax
dhcp6

Context
show>service>id

Platforms
7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description
Commands in this context display DHCPv6 information for the specified service.

statistics

Syntax
statistics
statistics [interface interface-name]
statistics sap sap-id
statistics sdp sdp-id:vc-id

Context
show>service>id>dhcp6

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command displays DHCPv6 statistics information.

Parameters

interface-name

Displays information for the specified IP interface, up to 32 characters.

sap-id

Specifies the physical port identifier portion of the SAP definition.

Values		
null		[port-id lag-id]
dot1q		[port-id lag-id pw-id]:qtag1
qinq		[port-id lag-id pw-id]:qtag1.qtag2
dot1q		[port-id lag-id]:cp-id
port-id		slot/mda/port
lag-id		lag-id
		lag keyword
		id 1 to 25
	pw-id	pw-id
		pw keyword
		id 1 to 10239
	cp-id	cp-id
		cp keyword
		id 1 to 1000
	qtag1	*, 0 to 4094
	qtag2	*, 0 to 4094

sdp-id:vc-id

Specifies the SDP identifier.

Values	<i>sdp-id</i> — 1 to 17407
	<i>vc-id</i> — 1 to 4294967295

Output

The following output is an example of DHCPv6 statistics information, and [Table 73: Output fields: DHCPv6 statistics](#) describes the output fields.

Sample output

```
*A:7210# show service id 100 dhcp6 statistics sap 1/1/6

=====
DHCP6 Statistics, service 100  Sap 1/1/6
=====
Client Packets Snooped           : 0
Client Packets Forwarded        : 0
Client Packets Dropped           : 0
Server Packets Snooped           : 0
Server Packets Forwarded        : 0
Server Packets Dropped           : 0
Invalid Packets Dropped          : 0
=====
*A:7210#
```

Table 73: Output fields: DHCPv6 statistics

Label	Description
Client Packets Snooped	The number of packets received from the DHCPv6 clients that were snooped by the node
Client Packets Forwarded	The number of packets received from the DHCPv6 clients that were forwarded by the node
Client Packets Dropped	The number of packets received from the DHCPv6 clients that were discarded by the node
Server Packets Snooped	The number of packets received from the DHCPv6 server that were snooped by the node
Server Packets Forwarded	The number of packets received from the DHCPv6 server that were forwarded by the node
Server Packets Dropped	The number of packets received from the DHCPv6 server that were discarded by the node
Invalid Packets Dropped	The number of corrupted/invalid packets received from the DHCPv6 clients that were discarded by the node

5.8.2.3 IGMP snooping show commands

igmp-snooping

Syntax
igmp-snooping

Context
show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context display IGMP snooping information.

all

Syntax

all

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays detailed information for all aspects of IGMP snooping on the VPLS service.

Output

The following outputs are examples of detailed IGMP snooping information, and [Table 74: Output fields: IGMP snooping all](#) describes the output fields.

- [Sample output](#), [Sample output for 7210 SAS in access-uplink mode](#), [Table 74: Output fields: IGMP snooping all](#)

Sample output

```
*Sample output (7210 SAS-T in network mode)
*A:7210-SAS>show>service>id>igmp-snooping# all

=====
IGMP Snooping info for service 2
=====

-----
IGMP Snooping Base info
-----
Admin State : Down
Querier      : No querier found
-----

Sap/Sdp      Oper  MRtr Send  Max  MVR      Num
Id           State Port Queries Grps From-VPLS Grps
-----
sap:1/1/1    Up    No   No    None 1        1
sap:1/1/4    Up    No   No    None Local  0
-----
```



```

Lcl-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 : 68129
MVR To SAP Cfg Drops    : 0

-----
IGMP Snooping Multicast VPLS Registration info
-----
IGMP Snooping Admin State : Down

MVR Admin State        : Down
MVR Policy              : None
-----
Local SAPs/SDPs
-----
Svc Id      Sap/Sdp      Oper      From      Num Local
            Id           State     VPLS      Groups
-----
2           sap:1/1/1     Up        1         0
2           sap:1/1/4     Up        Local     0
-----
MVR SAPs (from-vpls=2)
-----
Svc Id      Sap/Sdp      Oper      From      Num MVR
            Id           State     VPLS      Groups
-----
No MVR SAPs found.
=====
*A:7210-SAS>show>service>id>igmp-snooping#

```

Sample output for 7210 SAS in access-uplink mode

```

A:7210-SAS>show>service>id# igmp-snooping all

=====
IGMP Snooping info for service 1
=====

-----
IGMP Snooping Base info
-----
Admin State : Up
Querier      : 10.1.1.1 on SAP 1/1/1

-----
Sap/Sdp      Oper      MRtr Send   Max  Max  Num
Id           State     Port Queries Grps Srcs Grps
-----
sap:1/1/1    Up        Yes  No      None None 0
sap:1/1/2    Up        No   No      None None 1
-----

IGMP Snooping Querier info
-----
Sap Id       : 1/1/1
IP Address    : 10.1.1.1
Expires      : 255s

```

Up Time : 0d 16:51:04
Version : 2

General Query Interval : 125s
Query Response Interval : 10.0s
Robust Count : 2

----- IGMP Snooping Multicast Routers

MRouter	Sap/Sdp Id	Up Time	Expires	Version
10.1.1.1	1/1/1	0d 16:51:14	255s	2

Number of mrouter: 1

----- IGMP Snooping Proxy-reporting DB

Group Address	Mode	Up Time	Num Sources
224.0.0.2	exclude	0d 16:51:14	0

Number of groups: 1

----- IGMP Snooping SAP 1/1/1 Port-DB

Group Address	Mode	Type	Up Time	Expires	Num Src

Number of groups: 0

----- IGMP Snooping SAP 1/1/2 Port-DB

Group Address	Mode	Type	Up Time	Expires	Num Src

224.0.0.2	exclude	dynamic	0d 16:51:17	259s	0

Number of groups: 1

----- IGMP Snooping Static Source Groups

----- IGMP Snooping Statistics

Message Type	Received	Transmitted	Forwarded
General Queries	811311	0	811311
Group Queries	0	0	0
Group-Source Queries	0	0	0
V1 Reports	0	0	0
V2 Reports	18030	11928	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

Send Query Cfg Drops : 0
Import Policy Drops  : 0
Exceeded Max Num Groups : 0
Exceeded Max Num Sources : 0
=====

```

Table 74: Output fields: IGMP snooping all

Label	Description
Admin State	The administrative state of the IGMP instance.
Querier	The address of the IGMP querier on the IP subnet to which the interface is attached.
Sap or SDP Id	The SAP or SDP IDs of the service ID.
Oper State	The operational state of the SAP or SDP IDs of the service ID.
Mtrr Port	Displays if the port is a multicast router port.
Send Queries	Whether the send-queries command is enabled or disabled.
Max Num Groups	The maximum number of multicast groups that can be joined on this SAP or SDP.
MVR From VPLS	Specifies MVR from VPLS.
Num MVR Groups	The actual number of multicast groups that can be joined on this SAP or SDP.
MVR From VPLS Cfg Drops	The from VPLS drop count.
MVR To SAP Cfg Drops	The to SAP drop count.
MVR Admin State	The administrative state of MVR.
MVR Policy	The MVR policy name.

mfib

Syntax

mfib [brief] [ip | mac] brief

mfib [**group** *grp-address*]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays the multicast FIB on the VPLS service.

Parameters

brief

Displays a brief output.

group **grp** *grp-address*

Displays the multicast FIB for a specific multicast group address.

Output

The following output is an example of multicast FIB information, and [Table 75: Output fields: MFIB](#) describes the output fields.

Sample output

```
*A:SAS# show service id 1 mfib

=====
Multicast FIB, Service 1
=====
Group Address          Sap/Sdp Id              Svc Id   Fwd/Blk
-----
224.0.0.4              sap:1/1/1              Local    Fwd
-----
Number of entries: 1
=====
A:7210-SAS>show>service>id#
```

Table 75: Output fields: MFIB

Label	Description
Group Address	The IPv4 multicast group address.
SAP ID	The SAP/SDP to which the corresponding multicast stream is forwarded/blocked.
Forwarding/Blocking	Whether the corresponding multicast stream is blocked/forwarded.
Number of Entries	The number of entries in the MFIB.

Label	Description
Forwarded Packets	The number of multicast packets forwarded for the corresponding source/group.
Forwarded Octets	The number of octets forwarded for the corresponding source/group.
Svc ID	The service to which the corresponding multicast stream is forwarded/blocked. Local means that the multicast stream is forwarded/blocked to a SAP or SDP local to the service.

mrouters

Syntax

mrouters [detail]

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays all multicast routers.

Parameters

detail

Displays detailed information.

Output

The following output is an example of multicast router information, and [Table 76: Output fields: IGMP-snooping Mroutes](#) describes the output fields.

Sample output

```
A:7210-SAS>show>service>id# igmp-snooping mroutes
```

```
=====
IGMP Snooping Multicast Routers for service 1
=====
```

MRouter	Sap/Sdp Id	Up Time	Expires	Version
10.1.1.1	1/1/1	0d 16:53:44	254s	2

```
-----
Number of mroutes: 1
=====
```

```
A:7210-SAS>show>service>id#
```


Table 76: Output fields: IGMP-snooping Mrouters

Label	Description
MRouter	The multicast router port.
Sap/Sdp Id	The SAP and SDP ID multicast router ports.
Up Time	The length of time the mrouter has been up.
Expires	The amount of time left before the query interval expires.
Version	The configured version of IGMP running on this interface.
Number of Mrouters	The number of multicast routers.

mvr

Syntax

mvr

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays Multicast VPLS Registration (MVR) information.

Output

The following output is an example of MVR information, and [Table 77: Output fields: MVR](#) describes the output fields.

Sample output

```
*A:7210-SAS>show>service>id>igmp-snooping# mvr

=====
IGMP Snooping Multicast VPLS Registration info for service 2
=====
IGMP Snooping Admin State : Down

MVR Admin State           : Down
MVR Policy                 : None
-----
Local SAPs/SDPs
-----
Svc Id    Sap/Sdp                Oper    From    Num Local
```

	Id	State	VPLS	Groups
2	sap:1/1/1	Up	1	0
2	sap:1/1/4	Up	Local	0

MVR SAPs (from-vpls=2)				

Svc Id	Sap/Sdp Id	Oper State	From VPLS	Num MVR Groups

No MVR SAPs found.				
=====				
*A:7210-SAS>show>service>id>igmp-snooping#				

Table 77: Output fields: MVR

Label	Description
MVR Admin State	The Administrative state.
MVR Policy	The Policy name.
Svc ID	The service identifier.
Sap/Sdp Id	The SAP and SDP IDs of the service ID.
Oper State	The operational state of the SAP and SDP IDs of the svcid.
Mrtr Port	If the port is a multicast router port.
From VPLS	Displays from which VPLS the multicast streams corresponding to the groups learned via this SAP are copied. If local, it is from its own VPLS.
Num Groups	The number of groups learned via this local SAP.

port-db

Syntax

port-db sap *sap-id* [**detail**]

port-db sap *sap-id* **group** *grp-address*

port-db sdp *sdp-id:vc-id* [**detail**]

port-db sdp *sdp-id:vc-id* **group** *grp-address*

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays information about the IGMP snooping port database for the VPLS service.

Parameters

- group *grp-ip-address***

Displays the IGMP snooping port database for a specific multicast group address.
- sap *sap-id***

Displays the IGMP snooping port database for a specific SAP. See [Common CLI command descriptions](#) for command syntax.
- sdp *sdp-id***

Displays only IGMP snooping entries associated with the specified mesh SDP or spoke-SDP. For a spoke-SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 to 17407
- vc-id**

Specifies the virtual circuit ID on the SDP ID for which to display information.

Values 1 — 4294967295

Default For mesh SDPs only, all VC IDs.

Output

The following output is an example of IGMP snooping port database information, and [Table 78: Output fields: port database](#) describes the output fields.

Sample output for 7210 SAS in network mode

```
*A:7210-SAS>show>service>id>igmp-snooping# port-db sap 1/1/1

=====
IGMP Snooping SAP 1/1/1 Port-DB for service 2
=====
Group Address      Type      From-VPLS  Up Time      Expires  MC
                                     Stdby
-----
224.0.0.1          dynamic 1          0d 00:15:57  246s
-----
Number of groups: 1
=====
*A:7210-SAS>show>service>id>igmp-snooping#

=====
*A:MTU-7210#
*A:7210-SAS>show>service>id>igmp-snooping# port-db sap 1/1/1 detail
```

```

=====
IGMP Snooping SAP 1/1/1 Port-DB for service 2
=====
-----
IGMP Group 224.0.0.1
-----
Type           : dynamic
Up Time        : 0d 00:14:30      Expires         : 259s
Compat Mode    : IGMP Version 2
V1 Host Expires : 0s              V2 Host Expires  : 259s
MVR From-VPLS  : 1                MVR To-SAP       : 1/1/4
MC Standby     : no
-----
Number of groups: 1
=====
*A:7210-SAS>show>service>id>i

```

Table 78: Output fields: port database

Label	Description
Group Address	The IP multicast group address for which this entry contains information.
Mode	The type of membership reports received on the interface for the group. In the include mode, reception of packets sent to the specified multicast address is requested only from those IP source addresses listed in the source-list parameter of the IGMP membership report. In exclude mode, reception of packets sent to the specific multicast address is requested from all IP source addresses except those listed in the source-list parameter.
Type	How this group entry was learned. If this group entry was learned by IGMP, the value is set to dynamic. For statically configured groups, the value is set to static.
Compatibility mode	The IGMP mode. This is used in order for routers to be compatible with earlier version routers. IGMPv3 hosts must operate in Version 1 and Version 2 compatibility modes. IGMPv3 hosts must keep state per local interface regarding the compatibility mode of each attached network. A host's compatibility mode is determined from the host compatibility mode variable which can be in one of three states: IGMPv1, IGMPv2 or IGMPv3. This variable is kept per interface and is dependent on the version of general queries heard on that interface as well as the earlier version querier present timers for the interface.
V1 host expires	The time remaining until the local router assumes that there are no longer any IGMP Version 1 members on the IP

Label	Description
	subnet attached to this interface. Upon hearing any IGMPv1 membership report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 leave messages for this group that it receives on this interface.
V2 host expires	The time remaining until the local router assumes that there are no longer any IGMP Version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 membership report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv3 leave messages for this group that it receives on this interface.
Source address	The source address for which this entry contains information.
Up Time	The time since the source group entry was created.
Expires	The amount of time remaining before this entry is aged out.
Number of sources	The number of IGMP group and source specific queries received on this SAP.
Forwarding/Blocking	Whether this entry is on the forward list or block list.
Number of groups	The number of groups configured for this SAP.
From VPLS	Displays from which VPLS the multicast streams corresponding to the groups learned via this SAP are copied. If local, it is from its own VPLS.

proxy-db

Syntax

proxy-db [detail]

proxy-db group *grp-address*

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays information about the IGMP snooping proxy reporting database for the VPLS service.

Parameters

group *grp-ip-address*

Displays the IGMP snooping proxy reporting database for a specific multicast group address.

Output

The following output is an example of proxy reporting database information, and [Table 79: Output fields: proxy database](#) describes the output fields.

Sample output

```
*A:7210# show service id 100 igmp-snooping proxy-db
=====
IGMP Snooping Proxy-reporting DB for service 100
=====
Group Address      Up Time
-----
239.7.7.7          0d 00:05:30
239.7.7.8          0d 00:05:30
239.8.8.8          0d 00:03:42
-----
Number of groups: 3
=====
*A:7210#

*A:T2# show service id 100 igmp-snooping proxy-db detail
=====
IGMP Snooping Proxy-reporting DB for service 100
=====
IGMP Group 239.7.7.7
-----
Up Time : 0d 00:05:43
-----
IGMP Group 239.7.7.8
-----
Up Time : 0d 00:05:43
-----
IGMP Group 239.8.8.8
-----
Up Time : 0d 00:03:55
-----
Number of groups: 3
=====
*A:7210#
```

Table 79: Output fields: proxy database

Label	Description
Group Address	The IP multicast group address for which this entry contains information.
Mode	The type of membership reports received on the interface for the group. In the include mode, reception of packets sent to the specified multicast address is requested only from those

Label	Description
	IP source addresses listed in the source-list parameter of the IGMP membership report.
	In the "exclude" mode, reception of packets sent to the specific multicast address is requested from all IP source addresses except those listed in the source-list parameter.
Up Time	The total operational time in seconds.
Number of groups	The number of IGMP groups.

querier

Syntax

querier

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays information about the IGMP snooping queriers for the VPLS service.

Output

The following output is an example of IGMP snooping querier information, and [Table 80: Output fields: querier](#) describes the output fields.

Sample output

```
*A:7210# show service id 100 igmp-snooping querier
=====
IGMP Snooping Querier info for service 100
=====
Sap Id           : 1/1/1
IP Address       : 10.10.9.9
Expires          : 24s
Up Time          : 0d 00:05:20
Version          : 2

General Query Interval : 10s
Query Response Interval : 10.0s
Robust Count           : 2
=====
*A:7210#

*A:T2# show service id 100 igmp-snooping proxy-db
```

```

=====
IGMP Snooping Proxy-reporting DB for service 100
=====
Group Address      Up Time
-----
239.7.7.7          0d 00:05:30
239.7.7.8          0d 00:05:30
239.8.8.8          0d 00:03:42
-----
Number of groups: 3
=====
*A:T2#

```

Table 80: Output fields: querier

Label	Description
SAP Id	The SAP ID of the service.
IP address	The IP address of the querier.
Expires	The time left, in seconds, that the query expires.
Up time	The length of time the query has been enabled.
Version	The configured version of IGMP.
General Query Interval	The frequency at which host-query packets are transmitted.
Query Response Interval	The time to wait to receive a response to the host-query message from the host.
Robust Count	The value used to calculate several IGMP message intervals.

static

Syntax

static [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays information about static IGMP snooping source groups for the VPLS service.

Parameters

- sap *sap-id***

Displays static IGMP snooping source groups for a specific SAP. See [Common CLI command descriptions](#) for command syntax.
- sdp *sdp-id***

Displays the IGMP snooping source groups for a specific spoke or mesh SDP. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

Values 1 to 17407
- vc-id**

Specifies the virtual circuit ID on the SDP ID for which to display information. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

Values 1 to 4294967295

Default For mesh SDPs only, all VC IDs.

Output

The following output is an example of static IGMP snooping source group information, and [Table 81: Output fields: static](#) describes the output fields.

Sample output

```
*A:7210# show service id 100 igmp-snooping static

=====
IGMP Snooping Static Groups for service 100
=====
-----
IGMP Snooping Static Groups for SAP 1/1/2
-----
Group
-----
228.8.8.8
-----
Static (*,G) entries: 1
=====
*A:7210#
```

Table 81: Output fields: static

Label	Description
Source	Displays the IP source address used in IGMP queries.
Group	Displays the static IGMP snooping source groups for a specified SAP.

statistics

Syntax

statistics [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays IGMP snooping statistics for the VPLS service.

Parameters

sap *sap-id*

Displays IGMP snooping statistics for a specific SAP. See [Common CLI command descriptions](#) for command syntax.

sdp *sdp-id*

Displays the IGMP snooping statistics for a specific spoke or mesh SDP. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

Values 1 to 17407

vc-id

Specifies the virtual circuit ID on the SDP ID for which to display information. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

Values 1 to 4294967295

Default For mesh SDPs only, all VC IDs.

Output

The following outputs are examples of IGMP snooping statistics information, and [Table 82: Output Fields: IGMP-snooping statistics](#) describes the output fields.

- [Sample output for 7210 SAS in network mode](#)
- [Sample output for 7210 SAS-T in access-uplink mode](#)

Sample output for 7210 SAS in network mode

```
*A:7210-SAS>show>service>id>igmp-snooping# statistics
=====
IGMP Snooping Statistics for service 2
=====
Message Type          Received      Transmitted   Forwarded
```

```

-----
General Queries          0          0          0
Group Queries           0          0          0
V1 Reports              0          0          0
V2 Reports             142207        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

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
MCS Failures             : 0

MVR From VPLS Cfg Drops  : 142130
MVR To SAP Cfg Drops     : 0
=====
*A:7210-SAS>show>service>id>igmp-snooping#

```

Sample output for 7210 SAS-T in access-uplink mode

```

A:7210-SAS>show>service>id# igmp-snooping statistics
=====
IGMP Snooping Statistics for service 1
=====
Message Type           Received      Transmitted   Forwarded
-----
General Queries        816014        0             816014
Group Queries          0             0             0
Group-Source Queries   0             0             0
V1 Reports              0             0             0
V2 Reports             18134        11991         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

Send Query Cfg Drops     : 0
Import Policy Drops      : 0
Exceeded Max Num Groups  : 0
Exceeded Max Num Sources : 0
=====

```

```
A:7210-SAS>show>service>id#
```

Table 82: Output Fields: IGMP-snooping statistics

Label	Description
Message Type	The column heading for IGMP messages
General Queries	The number of general query messages received, transmitted, and forwarded
Group Queries	The number of group query messages received, transmitted, and forwarded
Group-Source Queries	The number of group-source query messages received, transmitted, and forwarded
V1 Reports	The number of IGMPv1 or MLDv1 report messages received, transmitted, and forwarded
V2 Reports	The number of IGMPv2 or MLDv2 report messages received, transmitted, and forwarded
V3 Reports	(IGMP only) The number of IGMPv3 report messages received, transmitted, and forwarded
V2 Leaves	(IGMP only) The number of IGMP leave 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	The number of packets dropped due to a bad IGMP 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

Label	Description
Wrong Version	The number of packets dropped due to a wrong version of IGMP or MLD
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

endpoint

Syntax

endpoint [*endpoint-name*]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays service endpoint information.

Parameters

endpoint-name

Specifies an endpoint name created in the **config>service>vpls** context.

Output

The following output is an example of service endpoint information, and [Table 83: Output fields: service ID endpoint](#) describes the output fields.

Sample output

```
*A:Dut-B# show service id 1 endpoint
=====
Service 1 endpoints
=====
Endpoint name       : mcep-t1
Description         : (Not Specified)
Revert time        : 0
```

```

Act Hold Delay           : 0
Ignore Standby Signaling : false
Suppress Standby Signaling : false
Block On Mesh Fail       : true

Psv Mode Active          : No
Tx Active                 : 231:1
Tx Active Up Time         : 0d 00:06:57
Revert Time Count Down    : N/A
Tx Active Change Count     : 5
Last Tx Active Change     : 02/13/2009 22:08:33
-----
Members
-----
Spoke-sdp: 221:1 Prec:1           Oper Status: Up
Spoke-sdp: 231:1 Prec:2           Oper Status: Up
=====
*A:Dut-B#

```

Table 83: Output fields: service ID endpoint

Label	Description
Service endpoints	
Endpoint name	The endpoint name.
Revert time	The revert time setting for the active spoke SDP.
Act Hold Delay	Not applicable
Ignore Standby Signaling	Whether standby signaling is ignored. True: standby signaling is ignored False: standby signaling is not ignored
Suppress Standby Signaling	Whether standby signaling is suppressed. True: standby signaling is suppressed False: standby signaling is not suppressed
Tx Active	The actively transmitting spoke SDP.
Tx Active Up Time	The length of time that the active spoke SDP has been up.
Revert Time Count Down	Not applicable
Tx Active Change Count	The number of times that there has been a change of active spoke SDPs.
Last Tx Active Change	The date and time when a different spoke SDP became the actively transmitting spoke SDP.
Members	

Label	Description
Spoke-sdp	The primary and secondary spoke SDPs that are associated with this endpoint and shows their precedence value (0 precedence indicates the primary spoke SDP).

5.8.2.4 VPLS clear commands

id

Syntax

id *service-id*

Context

clear>service

clear>service>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears commands for a specific service.

Parameters

service-id

Specifies the ID that uniquely identifies a service.

Values service-id: 1 to 214748364
 svc-name: A string up to 64 characters.

statistics

Syntax

statistics

Context

clear>service>stats

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears session statistics for this service.

fdb

Syntax

fdb {**all** | **mac** *ieee-address* | **sap** *sap-id*] | **mesh-sdp** *sdp-id[:vc-id]* | **spoke-sdp** *sdp-id:vc-id*}

Context

clear>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears FDB entries for the service.

Parameters

all

Clears all FDB entries.

mac *ieee-address*

Clears only FDB entries in the FDB table with the specified 48-bit MAC address. The MAC address can be expressed in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

mesh-sdp

Clears only service FDB entries associated with the specified mesh SDP ID. For a mesh SDP, the VC ID is optional. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

spoke-sdp

Clears only service FDB entries associated with the specified spoke-SDP ID. For a spoke-SDP, the VC ID must be specified. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

sdp-id

Specifies the SDP ID for which to clear associated FDB entries. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

vc-id

Specifies the virtual circuit ID on the SDP ID for which to clear associated FDB entries. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

Values			
	sdp-id[:vc-id]	<i>sdp-id</i>	1 to 17407
		<i>vc-id</i>	1 to 4294967295
	sdp-id:vc-id	<i>sdp-id</i>	1 to 17407
		<i>vc-id</i>	1 to 4294967295

mesh-sdp

Syntax

mesh-sdp *sdp-id[:vc-id]* **ingress-vc-label**

Context

clear>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command clears and resets the mesh SDP bindings for the service.

Parameters

sdp-id

Specifies the mesh SDP ID to be reset.

Values 1 to 17407

vc-id

Specifies the virtual circuit ID on the SDP ID to be reset.

Values 1 to 4294967295

Default All VC IDs on the SDP ID.

spoke-sdp

Syntax

spoke-sdp *sdp-id[:vc-id]* {**all** | **counters** | **stp** | **l2pt**}

Context

clear>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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.

stp

Clears all STP statistics associated with the SDP.

l2pt

Clears all L2PT statistics associated with the SDP.

sap

Syntax

sap *sap-id*

Context

clear>service>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears statistics for the SAP bound to the service.

Parameters

sap-id

See [Common CLI command descriptions](#) for command syntax.

all

Clears all queue statistics and STP statistics associated with the SAP.

counters

Clears all queue statistics associated with the SAP.

counters

Syntax

counters

Context

clear>service>statistics>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears all traffic counters associated with the service ID.

l2pt

Syntax

l2pt

Context

clear>service>statistics>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears the l2pt statistics for this service.

mesh-sdp

Syntax

mesh-sdp *sdp-id[:vc-id]* {**all** | **counters** | **stp** | **mrp**}

Context

clear>service>statistics>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command clears the statistics for a particular mesh SDP bind.

Parameters

sdp-id[:vc-id]

sdp-id - [1..17407]

vc-id - [1..4294967295]

all

Clears all queue statistics and STP statistics associated with the SDP.

counters

Clears all queue statistics associated with the SDP.

stp

Clears all STP statistics associated with the SDP.

mrp

Clears all MRP statistics associated with the SDP.

spoke-sdp

Syntax

spoke-sdp *sdp-id[:vc-id]* {**all** | **counters** | **stp** | **l2pt**}

Context

clear>service>statistics>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 and STP statistics associated with the SDP.

counters

Clears all queue statistics associated with the SDP.

stp

Clears all STP statistics associated with the SDP.

l2pt

Clears all L2PT statistics associated with the SDP.

stp

Syntax

stp

Context

clear>service>statistics>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Clears all spanning tree statistics for the service ID.

detected-protocols

Syntax

detected-protocols {**all** | **sap** *sap-id*}

Context

clear>service>id>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

RSTP automatically falls back to STP mode when it receives an STP BPDU. The **clear detected-protocols** command forces the system to revert to the default RSTP mode on the SAP.

Parameters

all

Clears all detected protocol statistics.

sap-id

Clears the specified lease state SAP information. See [Common CLI command descriptions](#) for command syntax.

igmp-snooping

Syntax

igmp-snooping

Context

clear>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context clear IGMP snooping data.

port-db

Syntax

port-db [**sap** *sap-id*] [**group** *grp-address*]

port-db sdp *sdp-id:vc-id* [**group** *grp-address*]

Context

clear>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears the information about the IGMP snooping port database for the VPLS service.

Parameters

sap *sap-id*

Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value. See [Common CLI command descriptions](#) for command syntax.

sdp-id

Clears only IGMP snooping entries associated with the specified mesh SDP or spoke-SDP. For a spoke-SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id

Specifies 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.

group *grp-address*

Clears IGMP snooping statistics matching the specified group address.

querier

Syntax

querier

Context

clear>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears the information about the IGMP snooping queriers for the VPLS service.

5.8.2.5 VPLS debug commands

id

Syntax

id *service-id*

Context

debug>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command debugs commands for a specific service.

Parameters

service-id
Specifies the ID that uniquely identifies a service.

Values service-id: 1 to 214748364
 svc-name: A string up to 64 characters.

event-type

Syntax

[no] event-type {config-change | svc-oper-status-change | sap-oper-status-change | sdpbind-oper-status-change}

Context

debug>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables a particular debugging event type.

The **no** form of this command disables the event type debugging.

Parameters

config-change

Debugs configuration change events.

svc-oper-status-change

Debugs service operational status changes.

sap-oper-status-change

Debugs SAP operational status changes.

sdpbind-oper-status-change

Debugs SDP operational status changes.

sap

Syntax

[no] sap *sap-id*

Context

debug>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables debugging for a particular SAP.

Parameters

sap-id

Specifies the SAP ID.

stp

Syntax

stp

Context

```
debug>service>id
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables the context for debugging STP.

```
all-events
```

Syntax

```
all-events
```

Context

```
debug>service>id>stp
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for all events.

```
bpdu
```

Syntax

```
[no] bpdu
```

Context

```
debug>service>id>stp
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for received and transmitted BPDUs.

core-connectivity

Syntax

[no] **core-connectivity**

Context

debug>service>id>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for core connectivity.

exception

Syntax

[no] **exception**

Context

debug>service>id>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for exceptions.

fsm-state-changes

Syntax

[no] **fsm-state-changes**

Context

debug>service>id>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for FSM state changes.

fsm-timers**Syntax**

[no] fsm-timers

Context

debug>service>id>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for FSM timer changes.

port-role**Syntax**

[no] port-role

Context

debug>service>id>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for changes in port roles.

port-state**Syntax**

[no] port-state

Context

debug>service>id>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for port states.

```
sap
```

Syntax

[no] **sap** *sap-id*

Context

debug>service>id>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for a specific SAP.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

```
sdp
```

Syntax

[no] **sdp** *sdp-id:vc-id*

Context

debug>service>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for a specific SDP.

6 IEEE 802.1ah Provider Backbone Bridging (PBB)

**Note:**

Provider Backbone Bridging (PBB) is supported only on the 7210 SAS-T operating in the network mode.

This chapter provides information about PBB, process overview, and implementation notes.

6.1 IEEE 802.1ah PBB overview

IEEE 802.1ah draft standard (IEEE802.1ah), also known as Provider Backbone Bridges (PBB), defines an architecture and bridge protocols for interconnection of multiple Provider Bridge Networks (PBNs - IEEE802.1ad QinQ networks). PBB is defined in IEEE as a connectionless technology based on multipoint VLAN tunnels. IEEE 802.1ah employs Provider MSTP as the core control plane for loop avoidance and load balancing. As a result, the coverage of the solution is limited by STP scale in the core of large service provider networks. The 7210 SAS-T in network mode supports a native PBB Ethernet backbone deployment.

The IEEE model for PBB is organized around a B-component handling the provider backbone layer and an I-component concerned with the mapping of Customer or Provider Bridge (QinQ) domain (for example, MACs, VLANs) to the provider backbone (for example, B-MACs, B-VLANs), that is, the I-component contains the boundary between the Customer and Backbone MAC domains. PBB encapsulates customer payload in a provider backbone Ethernet header, providing for Customer MAC hiding capabilities. With PBB, 7210 SAS platforms can be used for tier-1/2 aggregation, encapsulating customer service frames in PBB, allowing the PE-rs devices deployed in the metro core to be aware of only provider MAC addresses and for metro service scaling.


7210 SAS platforms fully support only native PBB deployment. They do not support the integrated PBB VPLS model. In particular, 7210 SAS platforms do not support use of SDPs in PBB services.

6.2 PBB features

6.2.1 Integrated PBB-VPLS solution

H-VPLS introduced a service-aware device in a central core location to provide efficient replication and controlled interaction at domain boundaries. The core network facing provider edge (N-PE) devices have knowledge of all VPLS services and customer MAC addresses for local and related remote regions resulting in potential scalability issues as shown in the following figure.

Figure 83: Large H-VPLS deployment




OSSG190

In a large VPLS deployment, it is important to improve the stability of the overall solution and to speed up service delivery. These goals are achieved by reducing the load on the N-PEs and respectively minimizing the number of provisioning touches on the N-PEs.

The integrated PBB-VPLS model provides an additional PBB hierarchy in the VPLS network to address these goals, as shown in the following figure.

Figure 84: Large PBB-VPLS Deployment




PBB encapsulation is added at the user facing PE (U-PE) to hide the customer MAC addressing and topology from the N-PE devices. The core N-PEs need to only handle backbone MAC addressing and do not need to have visibility of each customer VPN. As a result, the integrated PBB-VPLS solution decreases the load in the N-PEs and improves the overall stability of the backbone.

In the preceding figure, 7210 devices can only be used as U-PEs supporting only native Ethernet PBB services.

6.2.2 PBB technology

IEEE 802.1ah specification encapsulates the customer or QinQ payload in a provider header as shown in the following figure.

Figure 85: QinQ payload in provider header example



PBB adds a regular Ethernet header where the B-DA and B-SA are the backbone destination and respectively, source MACs of the edge U-PEs. The backbone MACs (B-MACs) are used by the core N-PE devices to switch the frame through the backbone.

A special group MAC is used for the backbone destination MAC (B-DA) when handling an unknown unicast, multicast or broadcast frame. This backbone group MAC is derived from the I-service instance identifier (ISID) using the rule: a standard group OUI (01-1E-83) followed by the 24 bit ISID coded in the last three bytes of the MAC address.

The BVID (backbone VLAN ID) field is a regular DOT1Q tag and controls the size of the backbone broadcast domain.

The following ITAG (standard Ether-type value of 0x88E7) has the role of identifying the customer VPN to which the frame is addressed through the 24 bit ISID.


6.2.3 PBB mapping to existing VPLS configurations

The IEEE model for PBB is organized around a B-component handling the provider backbone layer and an I-component concerned with the mapping of the customer/provider bridge (QinQ) domain (MACs, VLANs) to the provider backbone (B-MACs, B-VLANs). For example, the I-component contains the boundary between the customer and backbone MAC domains.

The Nokia implementation is extending the IEEE model for PBB to allow support for MPLS pseudowires using a chain of two VPLS context linked together as shown in the following figure.

7210 SAS does not support MPLS pseudowires in a PBB B-component and PBB I-component.

Figure 86: PBB mapping to VPLS constructs



Note:
I-PW and B-PW are not supported on 7210 SAS platforms.

A VPLS context is used to provide the backbone switching component. The white circle marked B, referred to as backbone-VPLS (B-VPLS) operates on backbone MAC addresses providing a core multipoint infrastructure that may be used for one or multiple customer VPNs. The Nokia B-VPLS implementation allows the use of native PBB infrastructures.

Note:
7210 SAS implementation allows the use of only native PBB over Ethernet infrastructures.

Another VPLS context (I-VPLS) can be used to provide the multipoint I-component functionality emulating the ELAN service (see the triangle marked "I" in the preceding figure). Similar to B-VPLS, I-VPLS inherits from the regular VPLS and native Ethernet (SAPs) hand-offs accommodating this way different types of access: for example, direct customer link, QinQ or H-VPLS.

To support PBB ELINE (point-to-point service), the use of an Epipe as I-component is allowed. All Ethernet SAPs supported by a regular Epipe are also supported in the PBB Epipe.

6.2.4 SAP support

This section contains information about the following topics:

- [PBB B-VPLS](#)
- [PBB I-VPLS](#)

6.2.4.1 PBB B-VPLS

- SAPs
 - Ethernet dot1q is supported - this is applicable to most PBB use cases, for example, one backbone VLAN ID used for native Ethernet tunneling

- Ethernet null is supported - this is supported for a direct connection between PBB PEs, for example, no BVID is required
- Default SAP types are blocked in the CLI for the B-VPLS SAP.
- The following rules apply to the SAP processing of PBB frames:
 - For “transit frames” (not destined for a local BMAC), there is no need to process the ITAG component of the PBB Frames. Regular Ethernet SAP processing is applied to the backbone header (BMACs and BVID).
 - If a local I-VPLS instance is associated with the B-VPLS, “local frames” originated/terminated on local I-VPLSes are PBB encapsulated/de-encapsulated using the **pbb-etype** = 0x88e7.

6.2.4.2 PBB I-VPLS

- Port Level
 - All existing Ethernet encapsulation types are supported (for example, null, dot1q, qinq).
- SAPs
 - The I-VPLS SAPs can coexist on the same port with SAPs for other business services, for example, VLL, VPLS SAPs.
 - All existing Ethernet encapsulation are supported: null, dot1q, qinq.


Existing SAP processing rules still apply for the I-VPLS case; the SAP encapsulation definition on Ethernet ingress ports defines which VLAN tags are used to determine the service that the packet belongs to:

- **null encap defined on ingress**
Any VLAN tags are ignored and the packet goes to a default service for the SAP.
- **dot1q encap defined on ingress**
Only first VLAN tag is considered.
- **QinQ encap defined on ingress**
Both VLAN tags are considered; wildcard support for the inner VLAN tag.
- For dot1q/qinq encapsulations, traffic encapsulated with VLAN tags for which there is no definition is discarded.
- Note that any VLAN tag used for service selection on the I-SAP is stripped before the PBB encapsulation is added. Appropriate VLAN tags are added at the remote PBB PE when sending the packet out on the egress SAP.

6.2.5 PBB packet walk-through

This section describes the walk-through for a packet that traverses the B-VPLS and I-VPLS instances using the example of a unicast frame between two customer stations as shown in the following network diagram in the following figure.

Figure 87: PBB packet walk-through



The station with CMAC (customer MAC) X needs to send a unicast frame to CMAC Y through the PBB-VPLS network. A customer frame arriving at PBB-VPLS U-PE1 is encapsulated with the PBB header. The local I-VPLS FIB on U-PE1 is consulted to determine the destination BMAC of the egress U-PE for CMAC Y. In our example, B2 is assumed to be known as the B-DA for Y. If CMAC Y is not present in the U-PE1 forwarding database, the PBB packet is sent in the B-VPLS using the standard group MAC address for the ISID associated with the customer VPN.

Next, only the Backbone Header in green is used to switch the frame through the green B-VPLS/VPLS instances in the N-PEs. At the receiving U-PE2, the CMAC X is learned as being behind BMAC B1; then the PBB encapsulation is removed and the lookup for CMAC Y is performed.

6.2.6 PBB ELINE service

ELINE service is defined in PBB (IEEE 802.1ah) as a point-to-point service over the B-component infrastructure. The Nokia implementation provides support for PBB ELINE through the mapping of multiple Epipe services to a Backbone VPLS infrastructure.

The use of Epipe scales the ELINE services as no MAC switching, learning or replication is required to deliver the point-to-point service.

All packets ingressing the customer SAP are PBB encapsulated and unicast through the B-VPLS "tunnel" using the backbone destination MAC of the remote PBB PE.

All the packets ingressing the B-VPLS destined for the Epipe are PBB de-encapsulated and forwarded to the customer SAP.

6.2.6.1 PBB resiliency for PBB Epipe service

The PBB Epipe service can be protected using G.8032 (the G8032 instance is created to protect the PBB B-VPLS service). For more information and for an example see [Overview of G.8032 operation](#).

6.2.6.2 PBB resiliency for B-VPLS

The following VPLS resiliency mechanisms are also supported in PBB VPLS:


- Native Ethernet resiliency supported in both I-VPLS and B-VPLS contexts.
- Distributed LAG, MC-LAG, RSTP.
- MSTP in a management VPLS monitoring (B- or I-) SAPs.
- The G.8032 is supported for B-VPLS service. The G.8032 support is used only with PBB Epipe service from the current releases and cannot be used with PBB I-VPLS service.

6.2.7 Access multi-homing for native PBB (B-VPLS over SAP infrastructure)

The Nokia PBB implementation allows the operator to use a native Ethernet infrastructure as the PBB core. Native Ethernet tunneling can be emulated using Ethernet SAPs to interconnect the related B-VPLS instances. This kind of solution may fit specific operational environments where Ethernet services was provided in the past using QinQ solution. The drawback is that no LDP signaling is available to provide support for Access Multi-homing for Epipe (pseudowire Active/Standby status) or I-VPLS services (LDP MAC Withdraw). An alternate solution is required.

A PBB network using Native Ethernet core is shown in the following figure. MC-LAG is used to multi-home a number of edge switches running QinQ to PBB BEBs.

Figure 88: Access dual-homing into PBB BEBs - topology view



The interrupted line from the MC-LAG represents the standby, inactive link; the solid line is the active link. The BEBs are dual-homed to two core switches BCB1 and BCB2 using native Ethernet SAPs on the B-VPLS side. Multi-point B-VPLS with MSTP for loop avoidance can be used as the PBB core tunneling.

6.2.8 PBB QoS

The following QoS processing rules apply for PBB B-VPLS SAPs:

- **B-VPLS SAP ingress**
 - If dot1p classification is enabled, the BTAG fields will be used by default to evaluate the internal forwarding class (fc) and discard profile if there is a BTAG field.
 - If dot1p classification is not explicitly enabled or the packets are untagged then the default fc and profile is assigned.
- **B-VPLS SAP egress**
 - If the access port based policy contains FC and profile to dot1p mapping, this entry is used to mark the dot1p bits in the B-TAG of the frame going out of the SAP. The I-Tag of the frame is not modified in any case.
 - If no explicit mapping exists, the related dot1p DE bits are set to zero on both ITAG and BTAG if the frame is originated locally from an I-VPLS. If the frame is transiting the B-VPLS the ITAG stays unchanged, the BTAG is set according to the type of ingress SAP.
 - If the ingress SAP is tagged, the values of the dot1p, DE bits are preserved in the BTAG going out on the egress SAP.
 - If the ingress SAP is untagged, the dot1p, DE bits are set to zero in the BTAG going out on the egress SAP.
- **I-SAP Ingress**

- SAP ingress classification using mac-criteria or IP DSCP is supported.
- **I-SAP Egress**
 - Access port based marking is supported for I-SAPs (dot1q and QinQ SAPs).

6.2.9 PBB ACL support

Filter policies are supported for ingress and egress of PBB I-SAP in both PBB Epipe and PBB VPLS service.

Only MAC criteria Filter policies is available for use with PBB B-SAPs on ingress with the following functionality:

- For PBB B-VPLS B-SAPs, the MAC filter matches the outer MAC header fields (that is, B-DA, B-SA, B-Tag) for traffic received on a B-SAP and forwarded to another B-SAP in the system.
- For PBB B-VPLS B-SAPs, the MAC filter matches the inner MAC header fields (that is, the customer MAC DA, SA and VLAN tags) for traffic received on a B-SAP and forwarded out of an I-SAP in the system.

Only MAC criteria filter policies is available for use with PBB B-SAPs on egress. This filter policy only matches the BCB traffic. BEB traffic (that is, PBB originated traffic) cannot be matched using the egress filter policy attached to PBB B-SAP.

6.2.10 Configuration guidelines

Listed as follows are the configuration guidelines for a PBB service:

- PBB services are supported only on 7210 SAS devices configured in network mode.
- A PBB service instance (identified by the ISID) cannot be used to encapsulate customer payloads with additional VLAN tags, if that service instance is being used to transport frames received on a QinQ access SAP. If a particular service instance is in use by a QinQ access SAP, then the system drops the packets that are received with additional tags on all the SAPs (NULL or Dot1q) using the same instance. Packets received with one or more tags on a NULL SAP, more than one tag on a Dot1q SAP, and more than two tags on a QinQ SAP are classified as packets with additional VLAN tags.
- Service MTU is not available for use.
- Port-based SHG is available for use with I-VPLS and B-VPLS service. Service based SHG is not available for use in an I-VPLS and a B-VPLS service.
- The system uses the internal loopback to flood/replicate BUM traffic received on the B-SAP, to create an additional copy for processing in the I-VPLS context. The system also uses the internal loopback to for egress port mirroring.

The user needs to ensure that aggregate amount of mirrored traffic in the system and the BUM traffic received on a B-SAP does not exceed the available internal loopback bandwidth. Ingress meters can be used to limit the amount of BUM received and processed from a B-SAP and user can limit the number of ports setup for port egress mirroring to control the maximum amount of traffic that needs to be circulated for two pass processing using the internal loopback. NOTE: If only PBB Epipe is used (no I-VPLS service is configured for use), then egress port mirroring can be enabled without affecting PBB traffic, because PBB Epipe traffic does not use the two-pass approach.

- Multiple B-SAPs on the same port cannot be part of the same B-VPLS service. Two B-SAPs on the same port need to be configured in two different services.
- Processing rules for packets received with multiple B-tags on a SAP:
 - If the B-Tag header has two tags, the packet is processed and forwarded appropriately and sent out of an I-SID service or another B-VPLS B-SAP.
 - If the node is acting as a pure BCB (with no ISID/service termination), then the packets are flooded and switched appropriately and if the node is acting as a BCB + BEB, then the packets are flooded and switched appropriately on the B-SAPs, but they will not be switched or flooded to I-SAPs (both VPLS and Epipe I-SAPs).
- PBB I-tag etype is not configurable, it is set to 0x88e7.
- PBB B-tag etype is not configurable; it is set to 0x8100.
- PBB packets received from a destination MAC address other than the one configured in the Epipe service are not accepted by 7210 devices.
- In the current release, PBB packets with UCA bit set are dropped.
- Aging of MAC addresses learned in the B-domain - As long as a Customer MAC (C-MAC) or an Epipe service is associated with a B-SA/B-MAC, do not age out the B-SA. When the last customer MAC ages out or the last Epipe service using the particular B-SA MAC is removed, remove the corresponding B-SA entry. This means that as long as an Epipe service is associated with a particular PBB destination MAC address, the corresponding B-MAC will not age out and will occupy an entry in the Layer 2 learning table. Note, that if only I-VPLS is in use, then aging out of C-MAC will automatically trigger aging out B-MAC, when the last C-MAC associated with the B-MAC is aged out.

6.2.11 Configuration guidelines for the 7210 SAS-T

The following configuration guidelines are specific to the 7210 SAS-T operating in the network mode.

When "discard-unknown" is enabled on a B-VPLS, the following behavior can be observed:

- Unknown unicast (B-DA) packets arriving on a B-SAP are dropped.
- Unknown unicast (C-DA) packets arriving on a B-SAP are processed in the I-VPLS, if the B-DA is not unknown unicast.
- Unknown unicast (C-DA) packets arriving on an I-SAP are not dropped and are flooded in the B-VPLS, because B-DA is equal to the "Group Mcast MAC" and is a known value
- Port based SHG is available for use with both I-VPLS and B-VPLS service. Service based SHG is not available in both.

6.3 Configuration examples

Use the CLI syntax displayed to configure PBB.

6.3.1 PBB ELAN and ELINE

Use the following syntax to bring up PBB B-VPLS - common to both ELAN and ELINE services.

```
config>service# vpls 200 customer 1 b-vpls create
description "This is a B-VPLS."
sap 3/1/3:33 create
description "B-VPLS SAP"
```

Use the following syntax to bring up PBB ELAN.

```
config>service# vpls 2000 customer 6 i-vpls create
description "This is an I-VPLS."
sap 4/1/3:20 create
description "I-VPLS SAP"
backbone-vpls 200
```

Use the following syntax to bring up PBB ELINE.

```
config>service# epipe 1000 customer 10 create pbb-epipe
description "This is an Epipe."
sap 4/1/3:20 create
description "Epipe SAP"
pbb-tunnel 200 backbone-dest-mac 00-01-10-1E-C6-67 isid 752
```

6.3.2 MC-LAG multi-homing for native PBB

This section describes a configuration example for BEB C configuration with the following assumptions:

- BEB C and BEB D are MC-LAG peers
- B-VPLS 100 on BEB C and BEB D
- VPLS 1000 on BEB C and BEB D
- MC-LAG 1 on BEB C and BEB D

```
service pbb
  source-bmac ab-ac-ad-ef-00-00
  port 1/1/1
    ethernet
      encap-type qinq
  lag 1
    port 1/1/1 priority 20
    lacp active administrative-key 32768
  redundancy
    multi-chassis
      peer 1.1.1.3 create
        source-address 10.1.1.1
        mc-lag
          lag 1 lacp-key 1 system-id 00:00:00:01:01:01 system-priority 100
          source-bmac-lsb use-lacp-key
  service vpls 100 bvpls
    sap 2/2/2:100 // bvid 100
    mac-notification
    no shutdown
  service vpls 101 bvpls
    sap 2/2/2:101 // bvid 101
```

```

mac-notification
no shutdown
// no per BVPLS source-bmac configuration, the chassis one (ab-ac-ad-ef-00-00) is used
service vpls 1000 ivpls
backbone-vpls 100
sap lag-1:1000 //automatically associates the SAP with ab-ac-ad-ef-00-01 (first 36 bits
from BVPLS 100 sbmac+16bit source-bmac-lsb)
service vpls 1001 ivpls
backbone-vpls 101
sap lag-1:1001 //automatically associates the SAP with ab-ac-ad-ef-00-01(first 36 bits
from BVPLS 101 sbmac+16bit source-bmac-lsb)

```

6.4 PBB command reference



Note:

PBB CLI commands are supported only on 7210 SAS-T operating in network mode.

6.4.1 Command hierarchies

- [PBB service commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

6.4.1.1 PBB service commands

```

config
- service
- pbb
- mac-name name ieee-address
- no mac-name
- source-bmac ieee-address
- no source-bmac

```

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]
- no vpls service-id
- pbb
- backbone-vpls service-id [isid isid]
- no backbone-vpls
- source-bmac ieee-address
- no source-bmac

```

```

config
- service
- [no] epipe service-id [customer customer-id] [create] [vpn vpn-id] svc-sap-type
{null-star | dot1q-preserve | any}] [customer-vid vlan-id] [pbb-epipe]
- pbb
- tunnel service-id backbone-dest-mac mac-name isid ISID

```

```
- tunnel service-id backbone-dest-mac ieee-address isid ISID
- no tunnel
```

6.4.1.2 Show commands

```
show
- eth-cfm
  - association [ma-index] [detail]
  - cfm-stack-table [port port-id [vlan qtag[.qtag]] | sdp sdp-id[:vc-id] [level 0..7]
[direction up | down]
  - domain [md-index] [association ma-index | all-associations [detail]]
  - mep mep-id domain md-index association ma-index [loopback] [linktrace]
  - mip
- service
  - id service-id
    - i-vpls
    - epipe
    - all
    - base
    - fdb {info | mac ieee-address | sap sap-id | detail | endpoint endpoint} [expiry]
[pbb]
  - stp [detail]
  - isid-using [ISID]
  - pbb
    - mac-name [detail]
  - service-using [b-vpls] [i-vpls]
```

6.4.1.3 Clear commands

```
clear
- service
  - id service-id
    - fdb {all | mac ieee-address | sap sap-id}
    - stp
      - detected-protocols [all | sap sap-id]
  - statistics
    - id service-id
      - counters
      - stp
    - sap sap-id {all | counters | stp}
```

6.4.1.4 Debug commands

```
debug
- service
  - id service-id
    - [no] event-type {config-change | svc-oper-status-change | sap-oper-statuschange}
    - [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
```

6.4.2 Command descriptions

- [PBB service configuration commands](#)
- [PBB show commands](#)
- [PBB clear commands](#)
- [PBB debug commands](#)

6.4.2.1 PBB service configuration commands

```
pbb
```

Syntax

```
pbb
```

Context

```
config>service
```

```
config>service>epipe
```

```
config>service>vpls
```

Platforms

7210 SAS-T (network operating mode)

Description

Commands in this context configure the parameters related to Provider Backbone Bridging (PBB).

```
mac-name
```

Syntax

```
mac-name name ieee-address
```

```
no mac-name name
```

Context

```
config>service>pbb
```

Platforms

7210 SAS-T (network operating mode)

Description

This command configures the MAC name for the MAC address. It associates an ASCII name with an IEEE MAC to improve the PBB Epipe configuration. It can also change the dest-BMAC in one place instead of thousands of Epipes.

Parameters

name

Specifies the MAC name up to 32 characters.

ieee-address

Specifies the MAC address assigned to the MAC name. The value should be input in either a xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format.

source-bmac

Syntax

source-bmac *ieee-address*

no source-bmac

Context

config>service>pbb

config>service>vpls>pbb

Platforms

7210 SAS-T (network operating mode)

Description

This command configures the base source BMAC for the B-VPLS. The first 32 bits must be the same as those configured in the MC-LAG peer. If the base source BMAC under VPLS is not configured, it inherits the chassis-level BMAC.

Parameters

ieee-address

Specifies the MAC address assigned to the BMAC. The value should be input in either a xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format.

tunnel

Syntax

tunnel *service-id* **backbone-dest-mac** *mac-name* **isid** *ISID*

tunnel *service-id* **backbone-dest-mac** *ieee-mac* **isid** *ISID*

no tunnel

Context

```
config>service>epipe>pbb
```

Platforms

7210 SAS-T (network operating mode)

Description

This command configures a PBB tunnel with Backbone VPLS (B-VPLS) service information.

Parameters***service-id***

Specifies the B-VPLS service for the PBB tunnel associated with this service.

Values 1 — 2147483648

backbone-dest-mac {mac-name | ieee-mac}

Specifies the backbone destination MAC-address for PBB packets.

isid ISID

Specifies a 24 bit service instance identifier for the PBB tunnel associated with this service. As part of the PBB frames, it is used at the destination PE as a demultiplexor field.

Values 0 — 16777215

vpls**Syntax**

```
vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type {null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls] [b-vid vid]  
no vpls service-id
```

Context

```
config>service
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates or edits a virtual private LAN services (VPLS) instance. The **vpls** command is used to create or maintain a VPLS service. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites, acting like a zero-hop Layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context. When a service is created, the **customer** keyword and *customer-id* parameter must be specified to associate the service with a customer. The *customer-id* value must already exist, having been created using the **customer** command in the service context. When 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.

When a service is created, the use of the **customer** *customer-id* command is optional for navigating into the service configuration context. Editing a service with the incorrect *customer-id* value specified results 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

Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values *service-id*: 1 to 2147483648

customer *customer-id*

Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

m-vpls

Keyword to create a management VPLS.

b-vpls

Keyword to create a PBB backbone-VPLS service, which can only be configured with SAPs. This keyword is supported only on 7210 SAS-T operating in network mode.

i-vpls

Keyword to create a PBB I-VPLS service, which can only be configured with SAPs. This keyword is only supported when the **svc-sap-type** value **any** is configured. This keyword is supported only on 7210 SAS-T operating in network mode.

create

Keyword that is mandatory when creating a VPLS service instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

svc-sap-type

Specifies the type of service and allowed SAPs in the service.

Values **dot1q-preserve** — Specifies that the allowed SAPs in the service are dot1q. The dot1q ID is not stripped after packets match the SAP.

This option can be configured in conjunction with the **b-vpls** or **r-vpls** keywords.

null-star — Specifies the allowed SAP in the service, which can be null SAPs, dot1q default, Q.* SAP, 0.* SAP, or default QinQ SAP. This option can be configured in conjunction with the **b-vpls** or **r-vpls** keywords.

any — Specifies that all supported SAPs are allowed in the service. This option can be configured in conjunction with the **b-vpls**, **r-vpls**, or **i-vpls** keywords. When these keywords are not configured, **any** can be used with a plain VPLS service, which can be configured with SAPs, spoke-SDPs, and mesh SDPs. See the section [QinQ SAP Configuration Restrictions for 7210 SAS platforms in network operating mode](#) for more information about restrictions related to QinQ SAPs.

Default any

b-vid vid

Specifies the VLAN ID to use when the **svc-sap-type** value is set to **dot1q-preserve**. This parameter is supported only when the **b-vpls** keyword and **svc-sap-type** value **dot1q-preserve** are configured.

Values 1 to 4094

r-vpls

Keyword that allows the VPLS instance to be associated with an IP interface to provide routed VPLS (RVPLS) functionality. When configured with the **svc-sap-type** values **null-star**, **dot1q-preserve**, and **any**, this keyword instantiates an RVPLS service that can be configured only with SAPs.



Note:

The **r-vpls** keyword is not supported in access-uplink mode (that is, in access-uplink mode, routed VPLS service can be configured without using this parameter).

backbone-vpls

Syntax

backbone-vpls *service-id* [**isid** *isid*]

no backbone-vpls

Context

config>service>vpls>pbb

Platforms

7210 SAS-T (network operating mode)

Description

This command configures B-VPLS service associated with the I-VPLS.

Parameters

service-id

Specifies the service ID.

Values 1..2147483648

isis

Specifies the ISID.

Values 0..16777215

6.4.2.2 PBB show commands

eth-cfm

Syntax

eth-cfm

Context

show

Platforms

7210 SAS-T (network operating mode)

Description

This command displays 802.1ag CFM information.

association

Syntax

association [*ma-index*] [**detail**]

Context

show>eth-cfm

Platforms

7210 SAS-T (network operating mode)

Description

This command displays association information.

Parameters

ma-index

Specifies the MA index value.

Values 1 — 4294967295

detail

Displays all association detail.

Output

The following output is an example of Ethernet CFM associations information, and [Table 84: Output fields: ETH-CFM associations](#) describes the output fields.

Sample output

```
*A:alcag1-R6# show eth-cfm association
=====
CFM Association Table
=====
Md-index   Ma-index   Name                CCM-interval Bridge-id
-----
1           1          ivpls                1             5000
=====
*A:alcag1-R6#
```

Table 84: Output fields: ETH-CFM associations

Label	Description
Md-index	Displays the MD index
Ma-index	Displays the MA index
Name	Displays the name of the MA
CCM-interval	Displays the CCM interval (in seconds)
Bridge-id	Displays the bridge ID for the MA. The bridge ID is the same value as the service ID of the service to which the MEP belongs.

cfm-stack-table

Syntax

cfm-stack-table

cfm-stack-table port [*port-id*> [**vlan** *qtag*[. *qtag*]] [**level** 0..7] [**direction** up | down]

cfm-stack-table sdp [*sdp-id*[:*vc-id*]>] [**level** 0..7]] [**direction** up | down]

cfm-stack-table virtual [*service-id*] [*level 0..7*]

Context
show>eth-cfm

Platforms
7210 SAS-T (network operating mode)

Description
This command summarizes all MEPs or MIPs.

Parameters

port-id
Displays information about the specified port.

Values	port-id	slot/mda/port[.channel]
	lag-id	lag-id
	lag	keyword
	id	1 — 200

sdp-id[:vc-id]
Specifies an existing SDP and VC ID.

Values	1 — 17407

qtag
Specifies the qtag value.

Values	0 — 4094

level
Specifies the level.

Values	0 — 7

direction up | down
Indicates the direction in which the maintenance association (MEP or MIP) faces on the bridge port.
down — Displays continuity check information configured away from the MAC relay entity.
up — Displays continuity check information configured toward the MAC relay entity.

service-id
Specifies information about the specified service ID.

Values	1 — 2147483648

Output

The following output is an example of Ethernet CFM stack-table information, and [Table 85: Output fields: ETH-CFM stack table](#) describes the output fields.

Sample output

```
*A:alcag1-R6# show eth-cfm cfm-stack-table
=====
CFM SAP Stack Table
=====
Sap          Level Dir  Md-index  Ma-index  Mep-id Mac-address
-----
1/2/9:5      4      Up    1          1          51      00:ae:ae:ae:ae:ae
=====
CFM SDP Stack Table
=====
Sdp          Level Dir  Md-index  Ma-index  Mep-id Mac-address
-----
No Matching Entries
=====
*A:alcag1-R6#
```

Table 85: Output fields: ETH-CFM stack table

Label	Description
Sap	Displays the SAP identifier.
Sdp	Displays the spoke SDP identifier.
Level	Displays the MD level of the domain.
Dir (direction)	Displays the direction of OAMPDU transmission.
Md-index	Displays the MD index of the domain.
Ma-index	Displays the MA index of the domain.
Mep-id	Displays the MEP identifier.
Mac-address	Displays the MAC address of the MEP.

domain

Syntax

domain [*md-index*] [**association** *ma-index* | **all-associations** [**detail**]]

Context

show>eth-cfm>domain

Platforms

7210 SAS-T (network operating mode)

Description

This command displays domain information.

Parameters

md-index

Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index

Specifies the MA index value.

Values 1 — 4294967295

all-associations

Displays information all maintenance associations.

detail

Displays detailed information.

Output

The following output is an example of Ethernet CFM domain information, and [Table 86: ETH-CFM domain field descriptions](#) describes the output fields.

Sample output

```
*A:alcag1-R6# show eth-cfm domain
=====
CFM Domain Table
=====
Md-index   Level Name                               Format
-----
1          4      ivpls                           charString
=====
*A:alcag1-R6#

*A:alcag1-R6# show eth-cfm mep 51 domain 1 association 1
-----
Mep Information
-----
Md-index      : 1                Direction      : Up
Ma-index      : 1                Admin           : Enabled
MepId         : 51               CCM-Enable     : Enabled
IfIndex       : 38043648         PrimaryVid     : 5
FngState      : fngReset
LowestDefectPri : allDef          HighestDefect   : none
Defect Flags   : None
Mac Address    : 00:ae:ae:ae:ae:ae CcmLtmPriority  : 7
CcmTx         : 775              CcmSequenceErr  : 0
CcmLastFailure Frame:
None
XconCcmFailure Frame:
None
*A:alcag1-R6#
```

Table 86: ETH-CFM domain field descriptions

Label	Description
Domain	
Md-index	Displays the MD index of the domain
Level	Displays the MD level of the domain
Name	Displays the name of the MD
Format	Displays the format for the MD name
Domain Associations	
Md-index	Displays the MD index of the domain
Direction	Displays the direction of OAMPDU transmission
Ma-index	Displays the MA index of the association
Admin	Displays the administrative status of the MEP
Mepld	Displays the MEP identifier
CCM-Enable	Displays the status of the CCM (enabled or disabled)
IfIndex	Displays the index of the interface
PrimaryVid	Displays the identifier of the primary VLAN
FngState	Indicates the different states of the Fault Notification Generator
LowestDefectPri	Displays the lowest priority defect (a configured value) that is allowed to generate a fault alarm
HighestDefect	Identifies the highest defect that is present (for example, if defRDICCM and defXconCCM are present, the highest defect is defXconCCM)
Defect Flags	Displays the number of defect flags
Mac Address	Displays the MAC address of the MEP
CcmLtmPriority	Displays the priority value transmitted in the linktrace messages (LTM)s and CCMs for this MEP. The MEP must be configured on a VLAN.
CcmTx	Displays the number of Continuity Check Messages (CCM) sent. The count is taken from the last polling interval (every 10 s).

Label	Description
CcmSequenceErr	Displays the number of CCM errors
CcmLastFailure Frame	Displays the frame that caused the last CCM failure
XconCcmFailure Frame	Displays the frame that caused the XconCCMFailure

mep

Syntax

mep *mep-id* **domain** *md-index* **association** *ma-index* [**loopback**] [**linktrace**]

Context

show>eth-cfm>domain

Platforms

7210 SAS-T (network operating mode)

Description

This command displays Maintenance Endpoint (MEP) information.

Parameters

mep-id

Specifies the maintenance association end point identifier.

Values 1 — 8191

md-index

Specifies the maintenance domain (MD) index value.

Values 1 — 4294967295

ma-index

Specifies the MA index value.

Values 1 — 4294967295

loopback

Displays loopback information for the specified MEP.

linktrace

Displays linktrace information for specified MEP.

Output

The following output is an example of MEP information, and [Table 87: Output fields: MEP](#) describes the output fields.

Sample output

```
*A:alcag1-R6# oam eth-cfm loopback 00:af:af:af:af:af mep 51 domain 1 association 1
eth-cfm Loopback Test Initiated: Mac-Address: 00:af:af:af:af:af, out sap: 1/2/9:5
Sent 1 packets, received 1 packets [0 out-of-order, 0 Bad Msdu] -- OK
*A:alcag1-R6#

*A:alcag1-R6# oam eth-cfm linktrace 00:af:af:af:af:af mep 51 domain 1 association 1
Index Ingress Mac      Egress Mac      Relay      Action
-----
1      00:00:00:00:00:00    00:AF:AF:AF:AF:AF  rlyHit     terminate
-----
No more responses received in the last 5 seconds.
*A:alcag1-R6#
```

Table 87: Output fields: MEP

Label	Description
Index	Displays the index of the domain
Relay	Displays the value returned in the Relay Action field
IngressMac	Displays the MAC address returned in the ingress MAC address field
Action	Displays the value returned in the Action field of the linktrace message
EgressMac	Displays the MAC address returned in the egress MAC address field

mip

Syntax
mip

Context
show>eth-cfm>mip

Platforms
7210 SAS-T (network operating mode)

Description
This command displays Maintenance Intermediate Point (MIP) information.

Output
The following output is an example of Ethernet CFM MIP information, and [Table 88: Output fields: MIP](#) describes the output fields.

Sample output

```
*A:7210SAS# show eth-cfm mip
=====
CFM SAP MIP Table
=====
Sap                               Mip-Enabled    Mip Mac Address
-----
1/1/16                           yes            00:a1:b1:c1:d1:e1
=====
CFM SDP MIP Table
=====
Sdp                               Mip-Enabled    Mip Mac Address
-----
456:123                          yes            00:a2:b2:c2:d2:e2
=====
*A:7210SAS#
```

Table 88: Output fields: MIP

Label	Description
Mip-Enabled	Displays the state of the MIP service
Mip Mac Address	Indicates the Mac Address of the MIP

id

Syntax

id service-id

Context

show>service

Platforms

7210 SAS-T (network operating mode)

Description

This command displays information about a specific service ID.

Parameters

service-id

The unique service identification number that identifies the service in the service domain.

Values service-id: 1 to 214748364

all

Displays detailed information about the service.

- base**
Displays basic service information.
- fdb**
Displays FDB entries.
- epipe**
Displays the Epipe services associated with the B-VPLS service.
- i-vpls**
Displays the I-VPLS services associated with this B-VPLS service.
- stp**
Display STP information.

all

Syntax
all

Context
show>service>id

Platforms
7210 SAS-T (network operating mode)

Description
This command displays detailed information for all aspects of the service.

Output
The following output is an example of detailed information for all aspects of a service, and [Table 89: Output fields: service ID all](#) describes the output fields.

Sample output

```
Sample output for PBB Epipe:

*A:7210-SAS>show>service# id 1000 all

=====
Service Detailed Information
=====
Service Id       : 1000                Vpn Id          : 0
Service Type     : Epipe
Description      : (Not Specified)
Customer Id      : 1
Last Status Change: 04/04/2001 22:18:48
Last Mgmt Change  : 04/04/2001 21:28:34
Admin State      : Up                  Oper State       : Up
MTU              : n/a
MTU Check        : n/a
Vc Switching     : False
SAP Count        : 1                  SDP Bind Count   : 0
```

```

Uplink Type:      : MPLS
-----
Service Destination Points(SDPs)
-----
No Matching Entries
-----
Service Access Points
-----
-----
SAP 1/1/15:1000
-----
Service Id      : 1000
SAP             : 1/1/15:1000          Encap           : q-tag
Description     : (Not Specified)
Admin State     : Up                  Oper State      : Up
Flags           : None
Last Status Change : 04/04/2001 21:29:23
Last Mgmt Change  : 04/04/2001 21:28:34
Dot1Q Ethertype  : 0x8100            QinQ Ethertype   : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU       : 1518                Oper MTU        : 1518
Ingr IP Fltr-Id : n/a                Egr IP Fltr-Id  : n/a
Ingr Mac Fltr-Id : n/a                Egr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a              Egr IPv6 Fltr-Id : n/a
tod-suite       : None
Endpoint        : N/A

Acct. Pol       : None                Collect Stats    : Disabled
-----
QoS
-----
Ingress qos-policy : 1
-----
Aggregate Policer
-----
rate            : n/a                burst           : n/a
-----
Ingress QoS Classifier Usage
-----
Classifiers Allocated: 4              Meters Allocated : 2
Classifiers Used      : 1              Meters Used       : 1
-----
Sap Statistics
-----
Ingress Stats:      Packets      Octets
Egress Stats:       0            0
Ingress Drop Stats: 0            0

Extra-Tag Drop Stats: n/a          n/a
-----
Sap per Meter stats (in/out counter mode)
-----
Ingress Meter 1 (Unicast)
For. InProf         : 0            0
For. OutProf        : 0            0
-----

```

```

PBB Tunnel Point
-----
B-vpls      Backbone-dest-MAC Isid      AdmMTU OperState Flood Oper-dest-MAC
-----
2           8c:90:d3:79:b2:65 1000    1514   Up        Yes   8c:90:d3:79:b2:65
-----
Last Status Change: 04/04/2001 22:18:48
Last Mgmt Change:   04/04/2001 22:18:48
-----

Service Endpoints
-----
No Endpoints found.
=====
*A:7210-SAS>show>service#

Sample output for I-VPLS:

*A:7210-SAS>show>service# id 200 all

=====
Service Detailed Information
=====
Service Id      : 200                Vpn Id          : 0
Service Type    : i-VPLS
Description     : (Not Specified)
Customer Id     : 1
Last Status Change: 04/04/2001 22:14:30
Last Mgmt Change : 04/04/2001 22:15:06
Admin State     : Up                 Oper State      : Up
MTU             : n/a
MTU Check       : n/a
SAP Count       : 1                 SDP Bind Count  : 0
Snd Flush on Fail : Disabled
Uplink Type     : MPLS
b-Vpls Id       : 2                 Oper ISID       : 200
b-Vpls Status   : Up
-----

Split Horizon Group specifics
-----

Service Destination Points(SDPs)
-----
No Matching Entries
-----

Service Access Points
-----

SAP 1/1/15:200
-----
Service Id      : 200
SAP             : 1/1/15:200         Encap           : q-tag
Description     : (Not Specified)
Admin State     : Up                 Oper State      : Up
Flags           : None
Last Status Change : 04/04/2001 22:14:30
Last Mgmt Change  : 04/04/2001 22:14:22
Dot1Q Ethertype : 0x8100            QinQ Ethertype  : 0x8100
Split Horizon Group: (Not Specified)

```

Max Nbr of MAC Addr:	No Limit	Total MAC Addr	: 0
Learned MAC Addr	: 0	Static MAC Addr	: 0
Admin MTU	: 1518	Oper MTU	: 1518
Ingr IP Fltr-Id	: n/a	Egr IP Fltr-Id	: 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		
Mac Learning	: Enabled	Discard Unkwn Srce:	Disabled
Mac Aging	: Enabled	Mac Pinning	: Disabled
BPDU Translation	: Disabled		
L2PT Termination	: Disabled		
Acct. Pol	: None	Collect Stats	: Disabled

 Stp Service Access Point specifics

Stp Admin State	: Up	Stp Oper State	: Down
Core Connectivity	: Down		
Port Role	: N/A	Port State	: Forwarding
Port Number	: 2049	Port Priority	: 128
Port Path Cost	: 10	Auto Edge	: Enabled
Admin Edge	: Disabled	Oper Edge	: N/A
Link Type	: Pt-pt	BPDU Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: N/A
Last BPDU from	: N/A		
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
TC bit BPDUs rcvd	: 0	TC bit BPDUs tx	: 0
RST BPDUs rcvd	: 0	RST BPDUs tx	: 0
MST BPDUs rcvd	: 0	MST BPDUs tx	: 0

 ARP host

Admin State	: outOfService		
Host Limit	: 1	Min Auth Interval	: 15 minutes

 QoS

Ingress qos-policy : 1

 Aggregate Policer

rate	: n/a	burst	: n/a

 Ingress QoS Classifier Usage

Classifiers Allocated:	4	Meters Allocated	: 2
Classifiers Used	: 2	Meters Used	: 2

 Sap Statistics

	Packets	Octets
Ingress Stats:	0	0
Egress Stats:	0	0
Ingress Drop Stats:	0	0
Extra-Tag Drop Stats:	n/a	n/a

Sap per Meter stats (in/out counter mode)

	Packets	Octets
Ingress Meter 1 (Unicast)		
For. InProf	: 0	0
For. OutProf	: 0	0
Ingress Meter 11 (Multipoint)		
For. InProf	: 0	0
For. OutProf	: 0	0

VPLS Spanning Tree Information

VPLS oper state	: Up	Core Connectivity	: Down
Stp Admin State	: Down	Stp Oper State	: Down
Mode	: Rstp	Vcp Active Prot.	: N/A
Bridge Id	: 80:00:00:25:ba:08:f6:20	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
Root Bridge	: N/A		
Primary Bridge	: N/A		
Root Path Cost	: 0	Root Forward Delay	: 0
Rcvd Hello Time	: 0	Root Max Age	: 0
Root Priority	: 0	Root Port	: N/A

Forwarding Database specifics

Service Id	: 200	Mac Move	: Disabled
Mac Move Rate	: 2	Mac Move Timeout	: 10
Mac Move Retries	: 3		
Table Size	: 250	Total Count	: 0
Learned Count	: 0	Static Count	: 0
Remote Age	: 900	Local Age	: 300
High Watermark	: 95%	Low Watermark	: 90%
Mac Learning	: Enabled	Discard Unknown	: Disabled
Mac Aging	: Enabled	Relearn Only	: False

Sample output for B-VPLS service:

*A:7210-SAS>show>service# id 2 all

Service Detailed Information

Service Id	: 2	Vpn Id	: 0
Service Type	: b-VPLS		
Description	: (Not Specified)		
Customer Id	: 1		
Last Status Change	: 04/04/2001 22:13:57		
Last Mgmt Change	: 04/04/2001 22:13:57		
Admin State	: Up	Oper State	: Up
MTU	: n/a		

```

MTU Check      : n/a
SAP Count      : 1
Snd Flush on Fail : Disabled
Uplink Type    : MPLS
Oper Backbone Src : 00:25:ba:08:f6:20
i-Vpls Count   : 1
Epipe Count    : 1

-----
Split Horizon Group specifics
-----
Service Destination Points(SDPs)
-----
No Matching Entries
-----
Service Access Points
-----

SAP 1/1/2:2
-----
Service Id      : 2
SAP             : 1/1/2:2
Description     : (Not Specified)
Admin State     : Up
Flags          : None
Last Status Change : 04/04/2001 22:13:57
Last Mgmt Change  : 04/04/2001 22:13:54
Dot1Q Ethertype : 0x8100
PBB Ethertype   : 0x88e7
Split Horizon Group: (Not Specified)

Max Nbr of MAC Addr: No Limit
Learned MAC Addr   : 0
Admin MTU          : 1518
Ingr Mac Fltr-Id   : n/a
tod-suite         : None
Mac Learning       : Enabled
Mac Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled

Total MAC Addr    : 0
Static MAC Addr   : 0
Oper MTU          : 1518
Egr Mac Fltr-Id   : n/a

Discard Unkwn Srce: Disabled
Mac Pinning       : Disabled

Acct. Pol        : None
Collect Stats    : Disabled

-----
Stp Service Access Point specifics
-----
Stp Admin State   : Up
Core Connectivity : Down
Port Role        : N/A
Port Number      : 2048
Port Path Cost    : 10
Admin Edge       : Disabled
Link Type        : Pt-pt
Root Guard       : Disabled
Last BPDU from    : N/A
CIST Desig Bridge : N/A

Stp Oper State    : Down
Port State       : Forwarding
Port Priority     : 128
Auto Edge        : Enabled
Oper Edge        : N/A
BPDU Encap       : Dot1d
Active Protocol   : N/A
Designated Port   : N/A

Forward transitions: 0
Cfg BPDUs rcvd    : 0
TCN BPDUs rcvd    : 0
TC bit BPDUs rcvd : 0

Bad BPDUs rcvd    : 0
Cfg BPDUs tx      : 0
TCN BPDUs tx      : 0
TC bit BPDUs tx   : 0

```

```

RST BPDUs rcvd      : 0
MST BPDUs rcvd      : 0
RST BPDUs tx        : 0
MST BPDUs tx         : 0
-----
ARP host
-----
Admin State          : outOfService
Host Limit           : 1
Min Auth Interval    : 15 minutes
-----
QOS
-----
Ingress qos-policy   : 1
-----
Aggregate Policer
-----
rate                 : n/a
burst                 : n/a
-----
Ingress QoS Classifier Usage
-----
Classifiers Allocated: 4
Classifiers Used      : 2
Meters Allocated      : 2
Meters Used           : 2
-----
Sap Statistics
-----
Ingress Stats:      Packets      Octets
                   0             0
Egress Stats:       0             0
Ingress Drop Stats: 0             0
Extra-Tag Drop Stats: n/a         n/a
-----
Sap per Meter stats (in/out counter mode)
-----
                   Packets      Octets

Ingress Meter 1 (Unicast)
For. InProf         : 0          0
For. OutProf        : 0          0

Ingress Meter 11 (Multipoint)
For. InProf         : 0          0
For. OutProf        : 0          0
-----
VPLS Spanning Tree Information
-----
VPLS oper state     : Up
Stp Admin State      : Down
Mode                 : Rstp
Core Connectivity    : Down
Stp Oper State       : Down
Vcp Active Prot.     : N/A

Bridge Id            : 80:00:00:25:ba:08:f6:20
Bridge Priority       : 32768
Topology Change      : Inactive
Last Top. Change     : 0d 00:00:00
Top. Change Count    : 0
Bridge Instance Id    : 0
Tx Hold Count        : 6
Bridge Hello Time     : 2
Bridge Max Age        : 20
Bridge Fwd Delay      : 15

Root Bridge          : N/A
Primary Bridge        : N/A

Root Path Cost        : 0
Rcvd Hello Time       : 0
Root Priority          : 0
Root Forward Delay    : 0
Root Max Age          : 0
Root Port             : N/A
-----

```



```

Forwarding Database specifics
-----
Service Id       : 2                Mac Move       : Disabled
Mac Move Rate    : 2                Mac Move Timeout : 10
Mac Move Retries : 3
Table Size       : 250              Total Count    : 0
Learned Count    : 0                Static Count   : 0
Remote Age       : 900              Local Age      : 300
High Watermark   : 95%              Low Watermark  : 90%
Mac Learning     : Enabled           Discard Unknown : Disabled
Mac Aging        : Enabled           Relearn Only   : False
-----

Related i-Vpls services for b-Vpls service 2
-----
i-Vpls SvcId      Oper ISID        Admin          Oper
-----
200                200                Up              Up
-----

Number of Entries : 1
-----

Related Epipe services for b-Vpls service 2
-----
Epipe SvcId      Oper ISID        Admin          Oper
-----
1000              1000              Up              Up
-----

Number of Entries : 1
-----

Service Endpoints
-----
No Endpoints found.
=====
*A:7210-SAS>show>service#

```

Table 89: Output fields: service ID all

Label	Description
Service Id	The service identifier.
Service Type	Specifies the type of service.
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.

Label	Description
Vc Switching	Displays the status of VC switching.
SAP Count	The number of SAPs specified for this service.
Uplink Type	Displays the mode of the device.
Vpn Id	The number which identifies the VPN.
Oper State	The operational state of the service.
SAP	Displays the SAP ID.
Encap	The value of the label used to identify this SAP on the access port.
QinQ Ethertype	Displays the configured QinQ Ethertype value.
Dot1Q Ethertype	Displays the configured Dot1Q Ethertype value.
Split Horizon Group	Displays the split horizon group information.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the port to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the port to the far-end router, without requiring the packet to be fragmented.
Ingr IP Fltr-Id	Ingress IP filter ID.
Egr IP Fltr-Id	Egress IP filter ID.
Ingr Mac Fltr-Id	Ingress MAC filter ID.
Egr Mac Fltr-Id	Egress MAC filter ID.
Ingr IPv6 Fltr-Id	Ingress IPv6 filter ID.
Egr IPv6 Fltr-Id	Egress IPv6 filter ID.
Endpoint	Displays the endpoint name.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
QoS parameters	
Ingress qos-policy	The SAP ingress QoS policy ID.
Classifiers Allocated	Displays the number of classifiers allocated.
Classifiers Used	Displays the number of classifiers used.

Label	Description
Meters Allocated	Displays the number of meters allocated.
Meters Used	Displays the number of meters used.
Ingress Stats	The number of received packets/octets for this SAP.
Egress Stats	The number of packets/octets forwarded out of this SAP.
PBB Tunnel Point parameters	
B-vpls	Displays the B-VPLS ID.
Backbone-dest-MAC	Displays the back bone destination MAC address.
Isid	Displays the ISID number.
Flood	Specifies whether or not the traffic is flooded in the B-VPLS for the Destination instead of unicast. If the backbone destination MAC is in the B-VPLS FDB, it is unicast.
Oper-dest-MAC	Displays the operational destination MAC address.
i-Vpls Count	Displays the count of I-VPLS bound to B-VPLS.
b-Vpls Status	Displays the operational state of the B-VPLS service.
Epipe Count	Displays the count of Epipe bound to B-VPLS.

base

Syntax

base

Context

show>service>id

Platforms

7210 SAS-T (network operating mode)

Description

This command displays basic information about the service including service type, description and SAPs.

Output

The following output is an example of basic service information, and [Table 90: Output fields: base](#) describes the output fields.

Sample output

Sample output for PBB Epipe service:

*A:7210-SAS>show>service# id 1000 base

```
=====
Service Basic Information
=====
```

```
Service Id       : 1000                Vpn Id          : 0
Service Type     : Epipe
Description      : (Not Specified)
Customer Id      : 1
Last Status Change: 04/04/2001 22:18:48
Last Mgmt Change : 04/04/2001 21:28:34
Admin State      : Up                  Oper State       : Up
MTU              : n/a
MTU Check        : n/a
Vc Switching     : False
SAP Count        : 1                  SDP Bind Count   : 0
Uplink Type      : MPLS
```

```
-----
Service Access & Destination Points
-----
```

Identifier	Type	AdmMTU	OprMTU	Adm	Opr
sap:1/1/15:1000	q-tag	1518	1518	Up	Up

```
-----
PBB Tunnel Point
-----
```

B-vpls	Backbone-dest-MAC	Isid	AdmMTU	OperState	Flood	Oper-dest-MAC
2	8c:90:d3:79:b2:65	1000	1514	Up	Yes	8c:90:d3:79:b2:65

```
Last Status Change: 04/04/2001 22:18:48
Last Mgmt Change:   04/04/2001 22:18:48
=====
```

*A:7210-SAS>show>service#

Sample output for I-VPLS service:

*A:7210-SAS>show>service# id 200 base

```
=====
Service Basic Information
=====
```

```
Service Id       : 200                Vpn Id          : 0
Service Type     : i-VPLS
Description      : (Not Specified)
Customer Id      : 1
Last Status Change: 04/04/2001 22:14:30
Last Mgmt Change : 04/04/2001 22:15:06
Admin State      : Up                  Oper State       : Up
MTU              : n/a
MTU Check        : n/a
SAP Count        : 1                  SDP Bind Count   : 0
Snd Flush on Fail : Disabled
Uplink Type      : MPLS
b-Vpls Id        : 2                  Oper ISID        : 200
b-Vpls Status    : Up
```

```

-----
Service Access & Destination Points
-----
Identifier                Type                AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/15:200            q-tag              1518    1518    Up   Up
=====
*A:7210-SAS>show>service#

Sample output for B-VPLS service:

*A:7210-SAS>show>service# id 2 base

=====
Service Basic Information
=====
Service Id      : 2                Vpn Id          : 0
Service Type    : b-VPLS
Description     : (Not Specified)
Customer Id     : 1
Last Status Change: 04/04/2001 22:13:57
Last Mgmt Change : 04/04/2001 22:13:57
Admin State     : Up                Oper State      : Up
MTU             : n/a
MTU Check       : n/a
SAP Count       : 1                SDP Bind Count  : 0
Snd Flush on Fail : Disabled
Uplink Type     : MPLS
Oper Backbone Src : 00:25:ba:08:f6:20
i-Vpls Count    : 1
Epipe Count     : 1

-----
Service Access & Destination Points
-----
Identifier                Type                AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/2:2              q-tag              1518    1518    Up   Up
=====
*A:7210-SAS>show>service#

```

Table 90: Output fields: base

Label	Description
Service Id	The service identifier.
Service Type	Specifies the type of service.
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.

Label	Description
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.
Vc Switching	Displays the status of VC switching.
SAP Count	The number of SAPs specified for this service.
Uplink Type	Displays the mode of the device.
Vpn Id	The number which identifies the VPN.
Oper State	The operational state of the service.
SAP	Displays the SAP ID.
Encap	The value of the label used to identify this SAP on the access port.
Vpn Id	The number which identifies the VPN.
Oper State	The operational state of the service.
SAP	Displays the SAP ID.
PBB Tunnel Point	
B-vpls	Displays the B-VPLS ID.
Backbone-dest-MAC	Displays the back bone destination MAC address.
Isid	Displays the ISID number.
Flood	Specifies whether or not the traffic is flooded in the B-VPLS for the Destination instead of unicast. If the backbone destination MAC is in the B-VPLS FDB, it is unicast.
b-Vpls Status	Displays the operational state of the B-VPLS service.
b-Vpls Id	Displays the B-VPLS ID.

fdb

Syntax

fdb {**info** | **mac** *ieee-address* | **sap** *sap-id* | **detail** | **endpoint** *endpoint*}
[expiry] **[pbb]**

Context

show>service>id

Platforms

7210 SAS-T (network operating mode)

Description

This command displays FDB entries for a specific MAC address.

Parameters

- sap sap-id**
Specifies the physical port identifier portion of the SAP.
- detail**
Displays detailed information.
- expiry**
Displays time until MAC is aged out.
- endpoint**
Displays endpoint information.
- pbb**
Displays PBB information.

Output

The following output is an example of FDB entry information, and [Table 91: Output fields: FDB](#) describes the output fields.

Sample output

```
*A:7210-SAS>show>service# id 200 fdb

=====
Forwarding Database, Service 200
=====
Service Id       : 200           Mac Move       : Disabled
Mac Move Rate    : 2            Mac Move Timeout : 10
Mac Move Retries : 3
Table Size       : 250          Total Count      : 0
Learned Count    : 0            Static Count     : 0
Remote Age       : 900          Local Age        : 300
High Watermark   : 95%         Low Watermark    : 90%
Mac Learning     : Enabled      Discard Unknown  : Disabled
Mac Aging        : Enabled      Relearn Only     : False
=====

*A:7210-SAS>show>service#

*A:7210-SAS>show>service# id 2 fdb

=====
Forwarding Database, Service 2
=====
Service Id       : 2            Mac Move       : Disabled
Mac Move Rate    : 2            Mac Move Timeout : 10
Mac Move Retries : 3
```

```

Table Size      : 250          Total Count      : 0
Learned Count   : 0           Static Count     : 0
Remote Age      : 900         Local Age       : 300
High Watermark  : 95%        Low Watermark  : 90%
Mac Learning    : Enabled     Discard Unknown : Disabled
Mac Aging       : Enabled     Relearn Only   : False
=====

```

```
*A:7210-SAS>show>service#
```

Table 91: Output fields: FDB

Label	Description
Service Id	Displays the service ID.
Mac Move Rate	Displays the maximum rate at which MAC's 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 MAs. The rate is computed as the maximum number of relearns allowed in a 5 second interval: for example, the default rate of 2 relearns per second corresponds to 10 relearns in a 5 second period.
Mac Move Retries	Displays the number of times retries are performed for reenabling the SAP.
Table Size	Specifies the maximum number of learned and static entries allowed in the FDB of this service.
Learned Count	Displays the current number of learned entries in the FDB of this service.
Remote Age	Displays the number of seconds used to age out FDB entries learned on an SAP. These entries correspond to MAC addresses learned on remote SAPs.
High Watermark	Displays the utilization of the FDB table of this service at which a table full alarm is raised by the agent.
Mac Learning	Specifies whether the MAC learning process is enabled
Mac Aging	Indicates whether the MAC aging process is enabled.
Mac Move	Displays the administrative state of the MAC movement feature associated with this service.
Mac Move Timeout	Displays 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 zero indicates that the SAP is automatically re-enabled after being disabled. If after the SAP is re-enabled

Label	Description
	it is disabled again, the effective retry timeout is doubled to avoid thrashing.
Total Count	Displays the total number of learned entries in the FDB of this service.
Static Count	Displays the current number of static entries in the FDB of this service.
Local Age	Displays the number of seconds used to age out FDB entries learned on local SAPs.
Low Watermark	Displays the utilization of the FDB table of this service at which a table full alarm is cleared by the agent.
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded.
Relearn Only	Displays, that when enabled, either the FDB table of this service is full, or that the maximum system-wide number of MAC's supported by the agent has been reached, and therefore MAC learning is temporary disabled, and only MAC relearns can take place.

stp

Syntax

stp [detail]

Context

show>service>id

Platforms

7210 SAS-T (network operating mode)

Description

This command displays information for the spanning tree protocol (STP) instance for the service.

Parameters

detail

Displays detailed information.

Output

The following output is an example of service STP information, and [Table 92: Output fields: STP](#) describes the output fields.

Sample output

```

*A:7210-SAS>show>service# id 200 stp

=====
Stp info, Service 200
=====
Bridge Id      : 80:00:00:25:ba:08:f6:20  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 Id      Oper-   Port-   Port-   Port-   Oper-   Link-   Active
                  State    Role    State    Num    Edge    Type    Prot.
-----
Backbone VPLS    Up      N/A     Forward  2048    N/A     N/A     N/A
1/1/15:200      Up      N/A     Forward  2049    N/A     Pt-pt  N/A

=====
*A:7210-SAS>show>service#

```

Table 92: Output fields: STP

Label	Description
Bridge Id	Specifies the MAC address used to identify this bridge in the network.
Top. Change Count	Specifies the total number of topology changes detected by the Spanning Tree Protocol instance associated with this service since the management Entity was last reset or initialized.
Root Bridge	Specifies the bridge identifier of the root of the spanning tree as determined by the Spanning Tree Protocol instance associated with this service. This value is used as the Root Identifier parameter in all Configuration BPDUs originated by this node.
Stp Oper State	Displays the operational state of the STP
Primary Bridge	Specifies the bridge identifier of the root of the spanning tree as determined by the Spanning Tree Protocol instance associated with this service. This value is used as the Root Identifier parameter in all Configuration BPDUs originated by this node.
Topology Change	Specifies whether a topology change is currently in progress.
Mode	Displays the mode of the STP

Label	Description
Last Top. Change	Specifies the time (in hundredths of a second) since the last time a topology change was detected by the Spanning Tree Protocol instance associated with this service.
Root Port	Specifies the port number of the port which provides the lowest cost path from this bridge to the root bridge.
Backbone VPLS	Displays the ID of the B-VPLS

isid-using

Syntax

isid-using [*ISID*]

Context

show>service

Platforms

7210 SAS-T (network operating mode)

Description

This command displays information about services using ISID.

Parameters

ISID

Displays the service using the specified I-component Service ID (ISID).

Values 0 — 16777215

Output

The following output is an example of information about services using ISID, and [Table 93: Output fields: ISID-using](#) describes the output fields.

Sample output

```
*A:7210-SAS>show>service# isid-using
```

```
=====
Services
=====
```

SvcId	ISID	Type	b-Vpls	Adm	Opr	SvcMtu	CustId
100	100	i-VPLS	1	Up	Up	1514	1
200	200	i-VPLS	2	Up	Up	1514	1
1000	1000	Epipe	2	Up	Up	1514	1
3000	3000	Epipe	1	Up	Up	1514	1

```
-----
```

```
Matching Services : 4
```

```
-----
=====
*A:7210-SAS>show>service#
```

Table 93: Output fields: ISID-using

Label	Description
SvcId	The service identifier.
ISID	Displays the ISID number.
Type	Indicates the type of service.
b-Vpls	Displays the B-VPLS ID.
Adm	Specifies the operating status of the service.
Opr	The current status of the service.
SvcMtu	Indicates the service MTU value.
CustId	Displays the customer ID.

i-vpls

Syntax

i-vpls

Context

show>service>id

Platforms

7210 SAS-T (network operating mode)

Description

This command displays I-VPLS services associated with the B-VPLS service. This command only applies when the service is a B-VPLS.

Output

The following output is an example of I-VPLS information, and [Table 94: Output fields: I-VPLS](#) describes the output fields.

Sample output

```
*A:7210-SAS>show>service# id 2 i-vpls
```

```
=====
Related i-Vpls services for b-Vpls service 2
```

```

=====
i-Vpls SvcId      Oper ISID      Admin      Oper
-----
200              200              Up          Up
-----
Number of Entries : 1
-----
*A:7210-SAS>show>service#

```

Table 94: Output fields: I-VPLS

Label	Description
i-Vpls SvcId	Displays the service ID of the I-VPLS service.
Oper ISID	Displays the ISID number.
Admin	Specifies the operating status of the service.
Oper	The current status of the service.

epipe

Syntax

epipe

Context

show>service>id

Platforms

7210 SAS-T (network operating mode)

Description

This command displays Epipe information for the PBB service.

Output

The following output is an example of Epipe information for a PBB service, and [Table 95: Output fields: service ID Epipe](#) describes the output fields.

Sample output

```

*A:7210-SAS>show>service# id 2 epipe
=====
Related Epipe services for b-Vpls service 2
=====
Epipe SvcId      Oper ISID      Admin      Oper
-----
1000            1000            Up          Up
-----

```

```
Number of Entries : 1
-----
*
=====
*A:7210-SAS>show>service# id 200 epipe
```

Table 95: Output fields: service ID Epipe

Label	Description
Epipe SvclId	Displays the service ID of the Epipe service bound to the B-VPLS service.
Oper ISID	Displays the ISID number.
Admin	Specifies the operating status of the service.
Oper	The current status of the service.

isid-using

Syntax

isid-using [ISID]

Context

show>service

Platforms

7210 SAS-T (network operating mode)

Description

This command displays the services using ISID.

Parameters

ISID

Displays the service using the specified I-component Service ID (ISID).

Values 0 — 16777215

Output

The following output is an example of information about services using ISID, and [Table 96: Output fields: ISID-using](#) describes the output fields.

Sample output

```
*A:7210-SAS>show>service# isid-using
=====
Services
```

```

=====
SvcId      ISID      Type      b-Vpls      Adm  Opr  SvcMtu  CustId
-----
100        100        i-VPLS    1            Up   Up   1514    1
200        200        i-VPLS    2            Up   Up   1514    1
1000       1000       Epipe     2            Up   Up   1514    1
3000       3000       Epipe     1            Up   Up   1514    1
-----
Matching Services : 4
-----
*A:7210-SAS>show>service#

```

Table 96: Output fields: ISID-using

Label	Description
SvcId	The service identifier.
ISID	Displays the ISID number.
Type	Indicates the type of service.
b-Vpls	Displays the B-VPLS ID.
Admin	Specifies the operating status of the service.
Oper	The current status of the service.
SvcMtu	Indicates the service MTU value.
Customer Id	Displays the customer ID.

service-using

Syntax

service-using [b-vpls] [i-vpls]

Context

show>service

Platforms

7210 SAS-T (network operating mode)

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

b-vpls
Displays matching Epipe services.

i-vpls
Displays matching VPLS instances.

Output

The following output is an example of service information, and [Table 97: Output fields: Service-using](#) describes the output fields.

Sample output

```
*A:7210-SAS>show>service# service-using b-vpls

=====
Services [bvpls]
=====
ServiceId      Type      Adm      Opr      CustomerId      Last Mgmt Change
-----
1              b-VPLS    Up       Up        1               04/04/2001 23:22:12
2              b-VPLS    Up       Up        1               04/04/2001 22:13:57
-----
Matching Services : 2
=====
*A:7210-SAS>show>service#
```

Table 97: Output fields: Service-using

Label	Description
Service Id	The service identifier.
Type	Indicates the type of service.
Admin	Specifies the operating status of the service.
Oper	The current status of the service.
Customer Id	Displays the customer ID.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.

mac-name

Syntax

mac-name [detail]

Context

show>service>pbb

Platforms

7210 SAS-T (network operating mode)

Description

This command displays information about a specific MAC name.

Parameters

detail

Displays detail information.

Output

The following output is an example of information for a specific MAC name, and [Table 98: Output fields: MAC name](#) describes the output fields.

Sample output

```
*A:7210-SAS>show>service# pbb mac-name test detail

=====
Services Using MAC name='test' addr='00:25:ba:08:f6:23'
=====
Svc-Id                ISID
-----
No Matching Entries
=====
*A:7210-SAS>show>service#
```

Table 98: Output fields: MAC name

Label	Description
Svc-Id	The service identifier.
ISID	Displays the ISID number.
Name	Displays the MAC name.
Addr	Displays the MAC address

6.4.2.3 PBB clear commands

id

Syntax

id service-id

Context

clear>service

```
clear>service>statistics
```

Platforms

7210 SAS-T (network operating mode)

Description

This command clears commands for a specific service.

Parameters

service-id

Specifies the ID that uniquely identifies a service.

Values service-id: 1 — 214748364

statistics

Syntax

statistics

Context

```
clear>service>stats
```

Platforms

7210 SAS-T (network operating mode)

Description

This command clears session statistics for this service.

fdb

Syntax

fdb {all | mac *ieee-address* | sap *sap-id* }

Context

```
clear>service>id
```

Platforms

7210 SAS-T (network operating mode)

Description

This command clears FDB entries for the service.

Parameters

all

Clears all FDB entries.

mac *ieee-address*

Clears only FDB entries in the FDB table with the specified 48-bit MAC address. The MAC address can be expressed in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers.

sap-id

Specifies the physical port identifier portion of the SAP definition.

sap

Syntax

sap *service-id*

Context

clear>service>statistics

Platforms

7210 SAS-T (network operating mode)

Description

This command clears statistics for the SAP bound to the service.

Parameters

sap-id

See [Common CLI command descriptions](#) for the command syntax.

counters

Syntax

counters

Context

clear>service>statistics>id

Platforms

7210 SAS-T (network operating mode)

Description

This command clears all traffic queue counters associated with the service ID.

stp

Syntax

stp

Context

clear>service>statistics>id

Platforms

7210 SAS-T (network operating mode)

Description

Clears all spanning tree statistics for the service ID.

detected-protocols

Syntax

detected-protocols {**all** | **sap** *sap-id*}

Context

clear>service>id>stp

Platforms

7210 SAS-T (network operating mode)

Description

RSTP automatically falls back to STP mode when it receives an STP BPDU. This command forces the system to revert to the default RSTP mode on the SAP.

Parameters

all

Clears all detected protocol statistics.

sap-id

Clears the specified lease state SAP information.

6.4.2.4 PBB debug commands

Id

Syntax

id *service-id*

Context

debug>service

Platforms

7210 SAS-T (network operating mode)

Description

This command debugs commands for a specific service.

Parameters

service-id

The ID that uniquely identifies a service.

Values service-id: 1 — 214748364

event-type

Syntax

[no] **event-type** {**config-change** | **svc-oper-status-change** | **sap-oper-status-change**}

Context

debug>service>id

Platforms

7210 SAS-T (network operating mode)

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 related events.

svc-oper-status-change

Debugs service operational status changes.

sap-oper-status-change

Debugs SAP operational status changes.

sap**Syntax**

[no] **sap** *sap-id*

Context

debug>service>id

Platforms

7210 SAS-T (network operating mode)

Description

This command enables debugging for a particular SAP.

Parameters

sap-id

Specifies the SAP ID.

stp**Syntax**

stp

Context

debug>service>id

Platforms

7210 SAS-T (network operating mode)

Description

This command enables the context for debugging STP.

all-events**Syntax**

all-events

Context

debug>service>id>stp

Platforms

7210 SAS-T (network operating mode)

Description

This command enables STP debugging for all events.

bpdu**Syntax**

[no] bpdu

Context

debug>service>id>stp

Platforms

7210 SAS-T (network operating mode)

Description

This command enables STP debugging for received and transmitted BPDUs.

core-connectivity**Syntax**

[no] core-connectivity

Context

debug>service>id>stp

Platforms

7210 SAS-T (network operating mode)

Description

This command enables STP debugging for core connectivity.

exception**Syntax**

[no] exception

Context

debug>service>id>stp

Platforms

7210 SAS-T (network operating mode)

Description

This command enables STP debugging for exceptions.

fsm-state-changes**Syntax**

[no] fsm-state-changes

Context

debug>service>id>stp

Platforms

7210 SAS-T (network operating mode)

Description

This command enables STP debugging for FSM state changes.

fsm-timers**Syntax**

[no] fsm-timers

Context

debug>service>id>stp

Platforms

7210 SAS-T (network operating mode)

Description

This command enables STP debugging for FSM timer changes.

port-role**Syntax**

[no] port-role

Context

debug>service>id>stp

Platforms

7210 SAS-T (network operating mode)

Description

This command enables STP debugging for changes in port roles.

port-state**Syntax**

[no] port-state

Context

debug>service>id>stp

Platforms

7210 SAS-T (network operating mode)

Description

This command enables STP debugging for port states.

sap**Syntax**

[no] sap *sap-id*

Context

debug>service>id>stp

Platforms

7210 SAS-T (network operating mode)

Description

This command enables STP debugging for a specific SAP.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.

7 Internet Enhanced Service

This chapter provides information about Internet Enhanced Services when 7210 SAS-T is operated in network mode and in access-uplink mode, and 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE is operated in Network mode, the process overview, and implementation notes.



Note:

For 7210 SAS platforms operating in network mode, IES can provide services or in-band management of the node.

For 7210 SAS platforms operating in access-uplink mode, IES is designed for in-band management of the node only.

This chapter explicitly notes if a feature is supported on 7210 SAS platforms operating in network or access-uplink mode.

7.1 IES service overview

Internet Enhanced Service (IES) is a routed connectivity service where the subscriber communicates with an IP router interface to send and receive Internet traffic. An IES has one or more logical IP routing interfaces each with a SAP which acts as the access point to the subscriber network.



Note:


For 7210 SAS platforms operating in access-uplink mode, IES is designed only for in-band management of the node.

IES allows IP interfaces to participate in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber be unique between other provider addressing schemes and potentially the entire Internet. While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be reserved for service IP provisioning, and be administered by a separate, but subordinate address authority.

IP interfaces defined within the context of an IES service must have a SAP associated as the uplink access point to the subscriber network. Multiple IES services are created to segregate subscriber owned IP interfaces.

The following figure shows a diagram of Internet enhanced service.

Figure 89: Internet Enhanced Service



OSSG023

The IES service provides in-band management connectivity. Other features include:

- Multiple IES services are created to separate IP interfaces.
- More than one IES service can be created for a single customer ID.
- More than one IP interface can be created within a single IES service ID. All IP interfaces created within an IES service ID belong to the same customer.

In access-uplink mode, the IES services provide IP connectivity to the node for in-band management of the node. Most of the management tasks supported with the out-of-band management port are supported with in-band management.

7.2 IES features

This section describes various general service features and any special capabilities or considerations as they relate to IES services.

7.2.1 IP interfaces

IES customer IP interfaces can be configured with most of the options found on the core IP interfaces. The advanced configuration options supported are:

- VRRP - for IES services with more than one IP interface (available only in network mode)
- Secondary IP addresses (available only on 7210 SAS-T (network mode), 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE platforms)
- ICMP Options

In network mode, configuration options found on core IP interfaces not supported on IES IP interfaces are:

- NTP broadcast receipt

7.2.1.1 IPv6 support for IES IP interfaces (access-uplink operating mode)



Note:

IPv6 addressing is supported for IES IP interfaces in access-uplink mode.

In access-uplink mode, IES IP interfaces associated with access-uplink SAPs support IPv6 addressing. IPv6 can be used for in-band management of the node using the IES IP interface.



Note:

IPv6 IES IP interfaces on access-uplink SAPs are only supported on 7210 SAS-T operating in access-uplink mode.

IPv4 and IPv6 route table lookup entries are shared. Before adding routes for IPv6 destinations, route entries in the routed lookup table needs to be allocated for IPv6 addresses. This can be done using the CLI command **config> system> resource-profile> router> max-ipv6-routes**. This command allocates route entries for /64 IPv6 prefix route lookups. The system does not allocate any IPv6 route entries by default and user needs to allocate some resources before using IPv6. For the command to take effect the node must be rebooted after making the change. For more information, see the following example and the *7210 SAS-Mxp, R6, R12, S, Sx, T Interface Configuration Guide*.

A separate route table (or a block in the route table) is used for IPv6 /128-bit prefix route lookup. A limited amount of IPv6 /128 prefixes route lookup entries is supported. The software enables lookups in this table by default (that is no user configuration is required to enable Ipv6 /128-bit route lookup).



Note:

IPv6 interfaces can be created without allocating IPv6 route entries.

Following features and restrictions is applicable for IPv6 IES IP interfaces:

- IPv6 interfaces supports only static routing.
- Only port-based ingress QoS policies are supported.
- IPv6 filter policies can be used on SAP ingress and egress.
- Routing protocols, such as OSPFv3, and others are not supported.
- A limited amount of IPv6 /128 prefixes route lookup entries is supported.

7.2.1.2 IPv6 support for IES IP interfaces (network operating mode)

IES IPv6 IP interfaces provide IPv6 connectivity in the routing base instance. It can be used to connect IPv6 networks over an IPv4 cloud using 6PE mechanisms. For more information about the 6PE, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Routing Protocols Guide*.

IPv4 and IPv6 route table lookup entries are shared. Before adding routes for IPv6 destinations, route entries in the routed lookup table needs to be allocated for IPv6 addresses. This can be done using the CLI command **config> system> resource-profile> router> max-ipv6-routes**. This command allocates route entries for /64 IPv6 prefix route lookups. The system does not allocate any IPv6 route entries by default and user needs to allocate some resources before using IPv6. For the command to take effect the node must be rebooted after making the change. For more information, see the following example and the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide*.

A separate route table (or a block in the route table) is used for IPv6 /128-bit prefix route lookup. A limited amount of IPv6 /128 prefixes route lookup entries is supported. The software enables lookups in this table by default (that is no user configuration is required to enable Ipv6 /128-bit route lookup).

In addition, the number IP subnets can be configured by the user using the command **configure> system>resource-profile>router>max-ip-subnets**. Suitable default are assigned to this parameter. Users can increase the number of subnets if they plan to more IPv6 addresses per IPv6 interface.

Following features and restrictions is applicable for IPv6 IES IP interfaces:

- IPv6 interfaces supports static routing, OSPv3, and IS-IS.
- A limited amount of IPv6 /128 prefixes route lookup entries is supported on 7210 SAS platforms.

7.2.1.3 Encapsulations

The following SAP encapsulation is supported on IES services in both network mode and access-uplink mode:

- Ethernet null
- Ethernet dot1q
- Ethernet QinQ

In 7210 SAS-T access-uplink mode, the following access-uplink SAP encapsulations are supported:

- Ethernet QinQ (access-uplink QinQ SAP)

7.2.2 Routing protocols

IES IP interfaces are restricted to routing protocols that can be configured on the interface. IES IP interfaces support the following routing protocols:

- RIP (only supported on the 7210 SAS-Mxp)
- OSPF
- IS-IS
- eBGP for the IPv4 and IPv6 address families (MPBGP is not supported)
- IGMP
- PIM
- BFD



Note:

The SAP for the IES IP interface is created at the IES service level, but the routing protocols for the IES IP interface are configured at the routing protocol level for the main router instance.

7.2.2.1 CPE connectivity check

Static routes are used within many IES services. Unlike dynamic routing protocols, there is no way to change the state of routes based on availability information for the associated CPE. CPE connectivity check adds flexibility so that unavailable destinations will be removed from the service provider routing tables dynamically and minimize wasted bandwidth.

The availability of the far-end static route is monitored through periodic polling. The polling period is configured. If the poll fails a specified number of sequential polls, the static route is marked as inactive.

An ICMP ping mechanism is used to test the connectivity. If the connectivity check fails and the static route is de-activated, the router will continue to send polls and reactivate any routes that are restored.

7.2.3 QoS policies

When applied to 7210 SAS IES services, service ingress QoS policies only create the unicast meters defined in the policy. The multipoint meters are not created on the service. With IES services, service egress QoS policies function as with other services where the class-based queues are created as defined in the policy.

On 7210 SAS ingress, only meters are supported on all the platforms.



Note:

QoS policies only create the unicast meters defined in the policy if PIM is not configured on the associated IP interface; if PIM is configured, the multipoint meters are applied as well.

In access-uplink mode, IES IP interface associated with an access SAP supports use of service ingress QoS policies. IES IP interface associated with an access-uplink SAP does not support use of service ingress QoS policies. IES IP interfaces associated with an access-uplink SAP share the port based ingress and egress QoS policies.

Note that both MAC and IPv4 criteria can be used in the QoS policies for traffic classification in an IES.

7.2.3.1 CPU QoS for IES interfaces in access-uplink mode

In access-uplink mode, IES IP interface bound to routed VPLS services, IES IP interface on access SAPs and IES IP interface on access-uplink SAPs are designed for use with inband management of the node. Consequently, they share a common set of queues for CPU bound management traffic. All CPU bound traffic is policed to predefined rates before being queued into CPU queues for application processing. The system uses meters per application or a set of applications. It does not allocate meters per IP interface. The possibility of CPU overloading has been reduced by use of these mechanisms. Users must use appropriate security policies either on the node or in the network to ensure that this does not happen.

7.2.3.2 CPU QoS for IES access interfaces in network mode

Traffic bound to CPU received on IES access interfaces are policed/rate-limited and queued into CPU queues. The software allocates a policer per IP application or a set of IP applications, for rate-limiting CPU bound IP traffic from all IES access SAPs. The policers CIR/PIR values are set to appropriate values based on feature scaling and these values are not user configurable. The software allocates a set of queues for CPU bound IP traffic from all IES access SAPs. The queues are either shared by a set of IP applications or in some cases allocated to an IP application. The queues are shaped to appropriate rate based on feature scaling. The shaper rate is not user configurable.



Note:

- The instance of queues and policers used for traffic received on network port IP interfaces is different for traffic received from access port IP interfaces. Additionally, the network CPU queues are accorded higher priority than the access CPU queues. This is done to provide better security and mitigate the risk of access traffic affecting network traffic.
- The 7210 SAS-Mxp allows the user to configure the IP differentiated services code point (DSCP) value for self-generated traffic. On the 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE, IP DSCP marking of self-generated traffic is not user-configurable and is assigned by software.

7.2.4 Filter policies

In network mode, only IP filter policies can be applied to IES services.

In access-uplink mode, only IP filter policies can be applied to IES service when either access SAP or access-uplink SAP is associated with the service.

7.2.5 VRRP support for IES IP interfaces in network operating mode



Note:

IPv4 is supported for IES IPv4 interfaces in network operating mode only.

- VRRP is not supported in access-uplink operating mode.
- For IPv6 interfaces, VRRP is not supported in both network and access-uplink operating mode.

The Virtual Router Redundancy Protocol (VRRP) for IPv4 is defined in the IETF RFC 3768, Virtual Router Redundancy Protocol. VRRP describes a method of implementing a redundant IP interface shared between two or more routers on a common LAN segment, allowing a group of routers to function as one virtual router. When this IP interface is specified as a default gateway on hosts directly attached to this LAN, the routers sharing the IP interface prevent a single point of failure by limiting access to this gateway address. For more information about use of VRRP, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.

7.3 Configuring an IES service with CLI

This section provides information to configure IES services using the command line interface.

7.3.1 Basic configuration

The most basic IES service configuration has the following entities:

- customer ID (see [Configuring customers accounts](#))
- an interface to create and maintain IP routing interfaces within IES service ID
- a SAP on the interface specifying the access port and encapsulation values

Example

The following is a sample configuration output of an IES service on ALA-48 on an access-uplink SAP (applicable for access-uplink mode only).

```
*A:ALA-48>config>service# info
-----
  ies 1000 customer 50 create
    description "to internet"
    interface "to-web" create
      address 10.1.1.1/24
      sap 1/1/5:0.* create
    exit
  exit
  no shutdown
-----
*A:ALA-48>config>service#
```

The following is a sample basic IES service configuration for IPv6, along with the use of max-ipv6-routes on the 7210 SAS.

Example

The following is a sample allocation of IPv6 routes on the node.

```
*A:7210SAS>config>system>res-prof# info
-----
max-ipv6-routes 1000
-----

NOTE: the node must be rebooted after the above change.

*A:ALA-50>config>service# info
-----
ies 1000 customer 50 vpn 1000 create
description "to inband-mgmt"
interface "to-mgmt" create
  ipv6
  address 2001:db8::1/24
  sap 1/1/10:100.* create
  exit
exit
no shutdown
-----
*A:ALA-50>config>service#
```

Example

The following is a sample configuration output of an IES service on ALA-50.

```
*A:ALA-50>config>service# info
-----
ies 1000 customer 50 vpn 1000 create
description "to internet"
interface "to-web" create
  address 10.1.1.1/24
  sap 1/1/10:100 create
  exit
exit
no shutdown
-----
*A:ALA-50>config>service#
```

7.3.2 Common configuration tasks

About this task

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

Procedure

- Step 1.** Associate an IES service with a customer ID.
- Step 2.** Associate customer ID with the service.
- Step 3.** Assign an IP address.
- Step 4.** Create an interface.

- Step 5.** Define SAP parameters on the interface
- Select nodes and ports.
 - Optionally, select filter policies (configured in the **config>filter** context).
- Step 6.** Enable service.

7.3.3 Configuring IES Components

7.3.3.1 Configuring an IES service

Use the following syntax to create an IES service.

Example: Basic IES service configuration

The following is a sample basic IES service configuration output.

```
A:ALA-48>config>service#
-----
...
    ies 1001 customer 1730 create
        description "to-internet"
        no shutdown
    exit
-----
A:ALA-48>config>service#
```

7.3.3.2 Configuring IES interface parameters

7.3.3.2.1 Network mode

Example

The following is a sample IES configuration output with interface parameters in network mode.

```
*A:7210-SAS>config>service>ies>if# info
-----
    arp-timeout 10000
    allow-directed-broadcasts
    icmp
        ttl-expired 120 38
    exit
    arp-populate
    ip-mtu 1000
    host-connectivity-verify interval 500 timeout 50 retry-count 15
    delayed-enable 150
    bfd 150 receive 300 multiplier 15 echo-receive 3000
    local-proxy-arp
    remote-proxy-arp
    loopback
*A:7210-SAS>config>service>ies>if#
-----
```

7.3.3.2 Access-uplink mode

Example

The following is a sample IES configuration with interface parameters in access-uplink mode.

```
*A:7210-SAS>config>service>ies>if# info
-----
  arp-timeout 10000
  allow-directed-broadcasts
  icmp
      ttl-expired 120 38
  exit
  ip-mtu 1000
-----
*A:7210-SAS>config>service>ies>if#
```

7.3.3.3 Configuring IES SAP parameters

A SAP is a combination of a port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique within a router.

When configuring IES access SAP parameters, a default QoS policy is applied to each SAP ingress. Additional QoS policies must be configured in the **config>qos** context. Filter policies are configured in the **config>filter** context and must be explicitly applied to a SAP. There are no default filter policies.

Example: IES SAP configuration output

```
-----
*A:ALA-A>config>service>ies>if# info
-----
  address 10.10.36.2/24
  sap 1/1/3:100 create
      ingress
          qos 101
      exit
  exit
-----
*A:ALA-A>config>service>ies>if#
```

7.3.3.4 Configuring VRRP

Configuring VRRP parameters on an IES interface is optional and is available only in network mode and is not supported in access-uplink mode. VRRP can be configured in either an owner or non-owner mode. The owner is the VRRP router whose virtual router IP address is the same as the real interface IP address. This is the router that responds to packets addressed to one of the IP addresses for ICMP pings, TCP connections and related addresses. All other virtual router instances participating in this message domain should have the same VRID configured and cannot be configured as an owner.

Example: IES configuration

The following is a sample IES configuration output.

```
*A:ALA-A>config>service>ies>if# info
-----
address 10.10.36.2/24
vrrp 2 owner
    backup 10.10.36.2
    authentication-type password
    authentication-key "3WErEDozxyQ" hash
exit
-----
*A:ALA-A>config>service#
```

7.3.4 Service management tasks

This section discusses the service management tasks.

7.3.4.1 Modifying IES service parameters

Existing IES service parameters in the CLI or NMS can be modified, added, removed, enabled or disabled. The changes are applied immediately to all services when the changes are applied.

To display a list of customer IDs, use the **show service customer** command.

Enter the parameters (such as description SAP information) and then enter the new information.

Example: Modified service configuration

```
*A:ALA-A>config>service>ies# info
-----
ies 1000 customer 50 create
    description "This is a new description"
    interface "to-web" create
        address 10.1.1.1/24
        mac 00:dc:98:1d:00:00
        sap 1/1/5:0.* create
    exit
    exit
    no shutdown
exit
-----
*A:ALA-A>config>service#
```

7.3.4.2 Deleting an IES service

An IES service cannot be deleted until SAPs and interfaces are shut down and deleted and the service is shutdown on the service level.

Use the following syntax to delete an IES service.

```
config>service#
[no] ies service-id
shutdown
[no] interface ip-int-name
shutdown
    [no] sap sap-id
    shutdown
```

7.3.4.3 Disabling an IES service

Use the following syntax to shut down an IES service without deleting the service parameters.

```
config>service> ies service-id
shutdown
```

7.3.4.4 Re-enabling an IES service

Use the following syntax to re-enable an IES service that was shut down.

```
config>service> ies service-id
[no] shutdown
```

Example:

```
config>service# ies 2000
config>service>ies# no shutdown
config>service>ies# exit
```

7.4 IES services command reference

7.4.1 Command hierarchies

- [IES configuration commands](#)
 - [Global commands](#)
 - [Interface commands \(network operating mode\)](#)
 - [RVPLS commands \(network operating mode\)](#)
 - [RVPLS commands \(access-uplink operating mode\)](#)
 - [Interface SAP commands \(network and access-uplink operating mode\)](#)
 - [IES SAP QoS and filter commands \(for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE\)](#)
 - [IES SAP QoS and filter commands \(for 7210 SAS-Mxp\)](#)
 - [Interface commands \(access-uplink operating mode\)](#)
 - [VRRP commands](#)
 - [Interface IPv6 commands](#)
- [Show commands](#)

7.4.1.1 IES configuration commands

7.4.1.1.1 Global commands



Note:

Global IES commands are supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

```
config
- service
- ies service-id [customer customer-id] [create] [vpn vpn-id]
- no ies service-id
- description description-string
- no description
- service-name service-name
- no service-name
- [no] shutdown
```

7.4.1.1.2 Interface commands (network operating mode)

```
config
- service
- ies service-id [customer customer-id] [create] [vpn vpn-id]
- [no] interface ip-int-name [create]
- address {ip-address/mask | ip-address netmask} [broadcast all-ones | host-
ones]
- no address {ip-address/mask | ip-address netmask}
- arp-timeout seconds
- no arp-timeout
- bfd transmit-interval [receive receive-interval] [multiplier multiplier]
[echo-receive echo-interval]
- no bfd
- cflowd-parameters
- sampling {unicast|multicast} type {interface} [direction {ingress-only}]
- no sampling {unicast|multicast}
- dhcp
- description description-string
- no description
- gi-address ip-address [src-ip-addr]
- no gi-address
- [no] option
- action {replace | drop | keep}
- no action
- [no] circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tuple]
- [no] remote-id [mac | string string]
- [no] vendor-specific-option
- [no] client-mac-address
- [no] sap-id
- [no] service-id
- string text
- no string
- [no] system-id
- no relay-plain-bootp
- relay-plain-bootp
- no server
```

```

- server server1 [server2...(upto 8 max)]
- [no] shutdown
- [no] trusted
- description description-string
- no description
- icmp
  - redirects [number seconds]
  - no redirects
  - ttl-expired [number seconds]
  - no ttl-expired
  - unreachable [number seconds]
  - no unreachable
- ip-mtu octets
- no ip-mtu
- [no] loopback
- [no] sap sap-id [create]
- secondary {ip-address/mask | ip-address netmask} [broadcast {all-ones | host-
ones}] [igp-inhibit]
- no secondary {ip-address/mask | ip-address netmask}
- static-arp ip-address ieee-mac-address
- no static-arp ip-address
- [no] static-arp ieee-mac-addr unnumbered
- [no] shutdown
- [no] vrrp virtual-router-id

```

7.4.1.1.3 Interface commands (access-uplink operating mode)

```

config
- service
  - ies service-id [customer customer-id]
  - interface ip-int-name [create]
  - [no] interface ip-int-name
    - address {[ip-address/mask | ip-address netmask} [broadcast all-ones | host-
ones]]
    - no address
    - arp-timeout seconds
    - no arp-timeout
    - allow-directed-broadcasts
    - no allow-directed-broadcasts
    - description long 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
        - [no] circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tuple]
        - [no] remote-id [mac | string string]
        - [no] vendor-specific-option
          - [no] client-mac-address
          - [no] sap-id
          - [no] service-id
          - string text
          - no string
          - [no] system-id
      - icmp
        - [no] mask-reply
        - redirects [number seconds]

```

```

- no redirects
- ttl-expired [number seconds]
- no ttl-expired
- unreachable [number seconds]
- no unreachable
- ip-mtu octets
- no ip-mtu
- [no] ipv6
- [no] loopback
- [no] sap sap-id [create]
- [no] shutdown
- [no] static-arp ip-address [ieee-address]

```

7.4.1.1.4 RVPLS commands (network operating mode)

```

config
- service
- ies service-id [customer customer-id] [vpn vpn-id]
- interface ip-interface-name [create]
- no interface ip-interface-name
- vpls service-name
- no vpls
- ingress
- [no] enable-table-classification
- routed-override-qos-policy policy-id
- no routed-override-qos-policy
- v4-routed-override-filter ip-filter-id
- no v4-routed-override-filter
- v6-routed-override-filter ip-filter-id
- no v6-routed-override-filter

```

7.4.1.1.5 RVPLS commands (access-uplink operating mode)

```

config
- service
- ies service-id [customer customer-id] [vpn vpn-id]
- interface ip-interface-name [create]
- no interface ip-interface-name
- vpls service-name
- no vpls
- ingress
- v4-routed-override-filter ip-filter-id
- no v4-routed-override-filter

```

7.4.1.1.6 Interface SAP commands (network and access-uplink operating mode)

```

config
- service
- ies service-id [customer customer-id] [vpn vpn-id] [create]
- [no] interface ip-int-name
- [no] sap sap-id [create]
- accounting-policy acct-policy-id
- no accounting-policy
- collect-stats
- no collect-stats

```

```

- description description-string
- no description
- ingress
  - meter-override
  - no meter-override
    - meter meter-id [create]
    - no meter meter-id
      - adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
      - no adaptation-rule
      - cbs size [kbits | bytes | kbytes]
      - no cbs
      - mbs size [kbits | bytes | kbytes]
      - no mbs
      - mode mode
      - no mode
      - rate cir cir-rate [pir pir-rate]
      - no rate
- statistics
  - ingress
    - counter-mode {in-out-profile-count | forward-drop-count}
- [no] tod-suite tod-suite-name
- [no] shutdown

```

7.4.1.1.7 IES SAP QoS and filter commands (for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE)

```

config
- service
  - ies service-id [customer customer-id] [vpn vpn-id] [create]
  - [no] interface ip-int-name
    - [no] sap sap-id [create]
    - egress
      - aggregate-meter-rate rate-in-kbps [burst burst-in-kbits] [enable-
stats]
      - no aggregate-meter-rate
      - filter ip ip-filter-id
      - filter ipv6 ipv6-filter-id
      - no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
    - ingress
      - aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
      - no aggregate-meter-rate
      - filter ip ip-filter-id
      - filter [ipv6 ipv6-filter-id]
      - no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
      - qos policy-id
      - no qos

```

7.4.1.1.8 IES SAP QoS and filter commands (for 7210 SAS-Mxp)

```

config
- service
  - ies service-id [customer customer-id] [vpn vpn-id] [create]
  - [no] interface ip-int-name
    - [no] sap sap-id [create]
    - egress
      - agg-rate-limit agg-rate
      - no agg-rate-limit

```



```

- filter [ip ip-filter-id]
- filter [ipv6 ipv6-filter-id]
- filter [mac mac-filter-id]
- no filter [ip ip-filter-id] [ipv6 ipv6 -filter-id] [mac mac-filter-
id]
- qos policy-id
- no qos
- ingress
- aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
- no aggregate-meter-rate
- filter [ip ip-filter-id]
- filter [ipv6 ipv6-filter-id]
- filter [mac mac-filter-id]
- no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
- qos policy-id [enable-table-classification]
- no qos

```

7.4.1.1.9 VRRP commands



Note:

IES VRRP commands are only supported on 7210 SAS platforms operating in network mode.

```

config
- service
- ies service-id [customer customer-id] [vpn vpn-id]
- interface ip-int-name
- ipv6
- vrrp virtual-router-id [owner]
- no vrrp virtual-router-id
- [no] backup ip-address
- init-delay seconds
- no init-delay
- [no] master-int-inherit
- message-interval {[seconds] [milliseconds milliseconds]}
- no message-interval
- [no] ping-reply
- policy vrrp-policy-id
- no policy
- [no] preempt
- priority priority
- no priority
- [no] shutdown
- [no] standby-forwarding
- [no] telnet-reply
- [no] traceroute-reply
- vrrp virtual-router-id [owner]
- no vrrp virtual-router-id
- authentication-key {authentication-key | hash-key} [hash | hash2]
- no authentication-key
- [no] backup ip-address
- [no] bfd-enable [service-id] interface interface-name dst-ip ip-address
- init-delay seconds
- no init-delay
- [no] master-int-inherit
- message-interval {[seconds] [milliseconds milliseconds]}
- no message-interval
- [no] ping-reply
- policy vrrp-policy-id
- no policy
- [no] preempt

```

```

- priority priority
- no priority
- [no] shutdown
- [no] ssh-reply
- [no] standby-forwarding
- [no] telnet-reply
- [no] traceroute-reply

```

7.4.1.1.10 Interface IPv6 commands

```

config
- service
  - ies service-id [customer customer-id] [create]
  - [no] interface ip-int-name [create]
    - ipv6
    - no ipv6
      - [no] address ipv6-address/prefix-length [eui-64] [preferred]
      - icmp6
        - [no] packet-too-big number seconds
        - [no] param-problem number seconds
        - [no] redirects number seconds
        - [no] time-exceeded number seconds
        - [no] unreachable number seconds
      - [no] link-local-address ipv6-address [preferred]
      - [no] local-proxy-nd
      - [no] neighbor ipv6-address mac-address
      - [no] proxy-nd-policy policy-name [policy-name...(up to 5 max)]

```

7.4.1.2 Show commands

```

show
- service
  - customer [customer-id] [site customer-site-name]
  - sap-using [sap sap-id]
  - sap-using interface [ip-address | ip-int-name]
  - sap-using [ingress | egress] filter filter-id
  - service-using [ingress] qos-policy qos-policy-id
  - service-using [ies] [customer customer-id]
  - id service-id
    - all
    - arp [ip-address] | [mac ieee-address] | [sap sap-id] | [interface ip-int-name]
    - base
    - interface [ip-address | ip-int-name] [detail]

```

7.4.2 Command descriptions

- [IES service configuration commands](#)
- [Show commands](#)

7.4.2.1 IES service configuration commands

- [Generic commands](#)

- [IES global commands](#)
- [IES interface IPv6 commands](#)
- [IES interface commands](#)
- [IES interface ICMP commands](#)
- [IES SAP commands](#)
- [IES QoS and filter commands](#)

7.4.2.1.1 Generic commands

description

Syntax

description *long description-string*

no description

Context

config>service>ies

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

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

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

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

Default

No description associated with the configuration context.

Parameters

string

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

shutdown

Syntax

[no] shutdown

Context

config>service>ies

config>service>ies>if

config>service>ies>if>ipv6>vrrp (7210 SAS-Mxp only)

config>service>ies>if>vrrp

config>service>ies>if>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described as follows in Special Cases.

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

Special Cases

IES

The default administrative status of an IES service is down. While the service is down, all its associated virtual router interfaces are operationally down. The administrative state of the service is not reflected in the administrative state of the virtual router interface.

For example if:

- 1) An IES service is operational and an associated interface is shut down.
- 2) The IES service is administratively shutdown and brought back up.
- 3) The interface shutdown remains in administrative shutdown state.

A service is regarded as operational provided that one IP Interface is operational.

IES IP Interfaces

When the IP interface is shutdown, it enters the administratively and operationally down states. For a SAP bound to the IP interface, no packets are transmitted out the SAP and all packets received on the SAP are dropped while incrementing the packet discard counter.

VRRP Protocol Handling

On all 7210 SAS platforms, VRRP is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure>service>ies>if>vrrp** command instantiates the protocol in the **no shutdown** state and resources are allocated to enable the node to process the protocol.

To deallocate resources, you must issue the **configure>service>ies>if>vrrp>shutdown** and **configure>service>ies>if>no vrrp** commands to allow the node to boot up correctly after the reboot. It is not sufficient to only issue a **configure>service>ies>if>vrrp>shutdown** command.

Note that the resources for VRRP are allocated when the VRRP context is enabled either in the base routing instance or the VPRN service instance. Resources are deallocated when the configuration of the last VRRP context under either base routing instances or VPRN service is removed or shutdown.

VRRPv3 Protocol Handling

On all 7210 SAS platforms, VRRPv3 is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure>service>ies>if>ipv6>vrrp** command instantiates the protocol in the **no shutdown** state and resources are allocated to enable the node to process the protocol.

To deallocate resources, you must issue the **configure>service>ies>if>ipv6>vrrp>shutdown** and **configure>service>ies>if>ipv6>no vrrp** commands to allow the node to boot up correctly after the reboot. It is not sufficient to only issue a **configure>service>ies>if>vrrp>ipv6>shutdown** command.

Note that the resources for VRRPv3 are allocated when the VRRPv3 context is enabled either in the base routing instance or the VPRN service instance. Resources are deallocated when the configuration of the last VRRPv3 context under either base routing instances or VPRN service is removed or shutdown.

DHCP Protocol Handling for 7210 SAS-Mxp

When the **no shutdown** command is issued in the **configure>service>ies>if>dhcp** context under the first IPv4 interface configured in the VPRN service instance, resources are allocated to enable the node to process the protocol. The resources are deallocated when you issue the **configure>service>ies>if>dhcp>shutdown** command for the last IPv4 interface enabled to use DHCP relay IPv4 configured in the VPRN service instance.

7.4.2.1.2 IES global commands

ies

Syntax

ies *service-id* **customer** *customer-id* [**create**] [**vpn** *vpn-id*]

no ies *service-id*

Context

config>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates or maintains an Internet Enhanced Service (IES) 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.

IP interfaces defined within the context of an IES service ID must have a SAP created.

When a service is created, the **customer** keyword and *customer-id* must be specified 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 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.

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 results in an error.

More than one IP interface may be created within a single IES service ID.



Note:

For 7210 SAS platforms operating in access-uplink mode, IES is used for in-band management of the node and cannot be used to deliver services. Typically, a single IP interface per IES is sufficient for management purposes.

The **no** form of this command deletes the IES service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces defined within the service ID have been shutdown and deleted.

Default

No IES service instances exist until they are explicitly created.

Parameters

service-id

Specifies the unique service identification number or string identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values *service-id*: 1 to 2147483648

customer *customer-id*

Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn vpn-id

Specifies the VPN ID assigned to the service. This parameter is only supported on 7210 SAS platforms operating in network mode.

Values 1 to 2147483647

service-name**Syntax**

service-name *service-name*

no service-name

Context

config>service>epipe

config>service>ies

config>service>vpls

config>service>vpn

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures an optional service name, up to 64 characters, which adds a name identifier to a specific service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7210 SAS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a specific service when 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).

7.4.2.1.3 IES interface IPv6 commands**ipv6****Syntax**

[no] ipv6

Context

```
config>service>ies>if
config>service>vprn>if (not supported in access-uplink operating mode)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command enables IPv6 for an IES interface.

address

Syntax

```
address ipv6-address/prefix-length [eui-64]
no address ipv6-address/prefix-length
```

Context

```
config>service>ies>if>ipv6
config>service>vprn>if>ipv6 (not supported in access-uplink operating mode)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command assigns an IPv6 address to the IES interface.

Parameters

ipv6-address/prefix-length	
Specifies the IPv6 address on the interface.	
Values	ipv6-address/prefix: ipv6-address x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x [0 — FFFF]H d [0 — 255]D prefix-length 1 to 128

eui-64
When this keyword is specified, a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from MAC address on Ethernet interfaces. For interfaces without a MAC address, for example ATM interfaces, the Base MAC address of the chassis is used.

icmp6

Syntax

icmp6

Context

config>service>ies>if>ipv6

config>service>vprn>if>ipv6 (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command configures ICMPv6 parameters for the IES interface.

packet-too-big

Syntax

packet-too-big [*number seconds*]

no packet-too-big

Context

config>service>ies>if>ipv6>icmp6

config>service>vprn>if>ipv6>icmp6 (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command specifies whether packet-too-big ICMPv6 messages should be sent. When enabled, ICMPv6 packet-too-big messages are generated by this interface.

The **no** form of this command disables the sending of ICMPv6 packet-too-big messages.

Default

100 10

Parameters***number***

Specifies the number of packet-too-big ICMPv6 messages to send in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds

Specifies the time frame in seconds that is used to limit the number of packet-too-big ICMPv6 messages issued.

Values 1 to 60

Default 10

param-problem**Syntax**

param-problem [*number seconds*]

no packet-too-big

Context

config>service>ies>if>ipv6>icmp6

config>service>vprn>if>ipv6>icmp6 (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command specifies whether parameter-problem ICMPv6 messages should be sent. When enabled, parameter-problem ICMPv6 messages are generated by this interface.

The **no** form of this command disables the sending of parameter-problem ICMPv6 messages.

Default

100 10

Parameters***number***

Specifies 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

Specifies the time frame in seconds that is used to limit the number of parameter-problem ICMPv6 messages issued.

Values 1 to 60

Default 10

redirects

Syntax

redirects [*number seconds*]
no redirects

Context

config>service>ies>if>ipv6>icmp6
config>service>vprn>if>ipv6>icmp6 (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command configures ICMPv6 redirect messages. When enabled, ICMPv6 redirects are generated when routes are not optimal on this router and another router on the same subnetwork has a better route to alert that node that a better route is available.

The **no** form of this command disable ICMPv6 redirect messages.

Default

100 10

Parameters

number

Specifies the number of version 6 redirects to be issued in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds

Specifies the time frame in seconds that is used to limit the number of version 6 redirects issued.

Values	1 to 60
Default	10

time-exceeded

Syntax

time-exceeded [*number seconds*]
no time-exceeded

Context

config>service>ies>if>ipv6>icmp6
config>service>vprn>if>ipv6>icmp6 (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command specifies whether time-exceeded ICMPv6 messages should be sent. When enabled, ICMPv6 time-exceeded messages are generated by this interface.
The **no** form of this command disables ICMPv6 time-exceeded messages.

Default

100 10

Parameters

number

Specifies 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

Specifies the time frame in seconds that is used to limit the number of "time-exceeded" ICMPv6 message to be issued.

Values	1 to 60
Default	10

unreachables

Syntax

unreachables [*number seconds*]
no unreachables

Context

config>service>ies>if>ipv6>icmp6
config>service>vprn>if>ipv6>icmp6 (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command specifies that ICMPv6 host and network unreachable messages are generated by this interface.
The **no** form of this command disables ICMPv6 host and network unreachable messages.

Default

100 10

Parameters

number

Specifies the number of destination unreachable ICMPv6 messages issued in the time frame specified by the *seconds* parameter.

Values	10 to 1000
Default	100

seconds

Specifies the time frame in seconds that is used to limit the number of destination unreachable ICMPv6 messages to be issued.

Values	1 to 60
Default	10

link-local-address

Syntax

link-local-address *ipv6-address* [**preferred**]

no link-local-address

Context

config>service>ies>if>ipv6

config>service>vprn>if>ipv6 (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command configures the IPv6 link local address.

local-proxy-nd

Syntax

[no] local-proxy-nd

Context

config>service>ies>if>ipv6

config>service>vprn>if>ipv6 (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command enables local proxy neighbor discovery on the interface.

The **no** form of this command disables local proxy neighbor discovery.

proxy-nd-policy

Syntax

proxy-nd-policy *policy-name* [*policy-name...*(up to 5 max)]

no proxy-nd-policy

Context

config>service>ies>if>ipv6

config>service>vprn>if>ipv6 (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command applies a proxy neighbor discovery policy for the interface.

Parameters

policy-name

Specifies an existing neighbor discovery policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. The specified policy names must already be defined.

neighbor

Syntax

neighbor *ipv6-address mac-address*

no neighbor *ipv6-address*

Context

config>service>ies>if>ipv6

config>service>vprn>if>ipv6 (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command configures IPv6-to-MAC address mapping on the IES interface.



Note:

This command is not supported for RVPLS interfaces.

Parameters

ipv6-address

Specifies the IPv6 address of the interface for which to display information.

Values x:x:x:x:x:x:x (eight 16-bit pieces)
 x: [0 to FFFF]H
 d: [0 to 255]D
 prefix-length [1 to 128]

mac-address

Specifies the 48-bit MAC address for the IPv6-to-MAC address mapping in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

7.4.2.1.4 IES interface commands

interface

Syntax

interface *ip-int-name* [create]

no interface *ip-int-name*

Context

config>service>ies

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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, is used to create and maintain IP routing interfaces within IES service IDs. The **interface** command can be executed in the context of an IES service ID. The IP interface created is associated with the service core network routing instance and default routing.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config service ies interface** (that is, the network core router instance). Interface names must not be in the dotted-decimal notation of an IP address. For example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All IES IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

**Note:**

- See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for information about how to allocate addresses toward IP subnets using the following commands:
 - **configure> system> resource-profile> router> max-ip-subnets** (applies to the 7210 SAS-T (network mode), 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE)
 - **configure> system> global-res-prof> router> max-ip-subnets** (applies to the 7210 SAS-Sx/S 1/10GE (standalone-VC))
- Before using IPv6, resources for IPv6 routes must be allocated. See the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide* for information about how to use the following commands:
 - **configure> system> resource-profile> router> max-ipv6-routes** (applies to the 7210 SAS-T (network mode), 7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone), and 7210 SAS-Sx 10/100GE)
 - **configure> system> global-res-prof> router> max-ipv6-routes** (applies to the 7210 SAS-Sx/S 1/10GE (standalone-VC))

The **no** form of this command removes IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For IES services, the IP interface must be shutdown before the SAP on that interface may be removed.

Parameters***ip-int-name***

Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

If *ip-int-name* already exists within the service ID, the context is changed to maintain that IP interface. If *ip-int-name* already exists within another service ID, an error occurs and context is not changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

address**Syntax**

address {*ip-address/mask* | *ip-address netmask*}

address *ip-address mask*

no address

Context

config>service>ies>if

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command assigns an IP address or IP subnet to an IES IP router interface. Only one IP address can be associated with an IP interface. An IP address must be assigned to each IES IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7210 SAS.

The IP address for the interface can be entered in either the Classless Inter-Domain Routing (CIDR) or traditional dotted-decimal notation. The **show** commands display CIDR notation and is stored in configuration files.

The **no** form of this command removes the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Address	Admin state	Oper state
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable. The address and admin states are the only controlling variable and can be set independently. If an address is assigned to an interface that is in an administratively up state, it becomes operationally up.

Parameters

ip-address

Specifies the IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted-decimal notation. Allowed values are IP addresses in the range 1.0.0.0 to 223.255.255.255 (with support of /31 subnets).

Values a.b.d.c (no multicast/broadcast address)

/

The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the *"/"* and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted-decimal mask must follow the prefix.

mask

Specifies the subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to

determine the host portion of the IP address. Allowed values are integers in the range 0 – 30. Note that a mask length of 32 is reserved for system IP addresses.

Values 0 to 32

netmask

Specifies the subnet mask in dotted-decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted-decimal mask. The *mask* parameter indicates the complete mask that is used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted-decimal addresses in the range 128.0.0.0 to 255.255.255.254. A mask of 255.255.255.255 is reserved for system IP addresses.

Values a.b.c.d (network bits all 1 and host bits all 0)

arp-timeout

Syntax

arp-timeout *seconds*

no arp-timeout

Context

config>service>ies>if

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the minimum time in seconds 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 **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

The **no** form of this command restores **arp-timeout** to the default value.

Default

14400 seconds

Parameters

seconds

Specifies the minimum number of seconds a learned ARP entry is stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries is not aged.

Values 0 to 65535

cflowd-parameters

Syntax

cflowd-parameters

Context

config>service>ies>interface

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure traffic sampling for the interface.

sampling

Syntax

sampling {unicast|multicast} type {interface} [direction {ingress-only}]

no sampling {unicast|multicast}

Context

config>service>ies>interface>cflowd-parameters

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables traffic sampling for the interface. See "Configuration Notes" in the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide* for more information.

The **no** form of this command disables traffic sampling for the interface.

Default

no sampling

Parameters

unicast

Keyword to enable unicast sampling.

multicast

Keyword to enable multicast sampling.

type

Keyword to configure the cflowd sampling type.

interface

Keyword to configure interface cflowd sampling type.

direction

keyword to configure the direction of the cflowd analysis.

ingress-only

Keyword to configure the ingress direction only for cflowd analysis.

dhcp

Syntax

dhcp

Context

config>service>ies >if

config>service>vpn >if

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure DHCP relay parameters.

gi-address

Syntax

gi-address *ip-address* [*src-ip-addr*]

no gi-address

Context

config>service>ies>if>dhcp (Network and Access-uplink Operating Mode)

config>service>vpn >if>dhcp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the gateway interface address for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. The GI address is needed, when the router functions as a DHCP relay, to distinguish between different interfaces.

By default, the GI address used in the relayed DHCP packet is the primary IP address of a normal IES interface. Specifying the GI address allows the user to choose a secondary address. For group interfaces a GI address must be specified under the group interface DHCP context or subscriber-interface DHCP context in order for DHCP to function.

Default

no gi-address

Parameters

ip-address

Specifies the host IP address to be used for DHCP relay packets.

src-ip-address

Specifies that this GI address is to be the source IP address for DHCP relay packets.

action

Syntax

action {replace | drop | keep}

no action

Context

config>service>ies >if>dhcp>option (Network and Access-uplink Operating Mode)

config>service>vprn >if>dhcp>option (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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

no action

Parameters

replace

Specifies that 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

Specifies that the DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.

keep

Specifies that 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 case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router inserts its own VSO into the Option 82 field. This is only done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router does not create the Option 82 field. In this case, no VSO is 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 (Network and Access-uplink Operating Mode)

config>service>vprn >if>dhcp>option (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the router to send either an ASCII tuple, or the interface index (If Index), on the specified SAP ID in the **circuit-id** sub-option of the DHCP packet.

When disabled, the **circuit-id** sub-option of the DHCP packet is left empty.

The **no** form of this command returns the system to the default.

Default

circuit-id ascii-tuple

Parameters

ascii-tuple

Specifies that the ASCII-encoded concatenated tuple is used which consists of the access-node-identifier, service-id, and interface-name, separated by "|".

ifindex

Specifies that the interface index is 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 is used.

vlan-ascii-tuple

Specifies that the format includes VLAN ID, dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Therefore, when the Option 82 bits are stripped, dot1p bits is copied to the Ethernet header of an outgoing packet.

option

Syntax

[no] option

Context

config>service>ies >if>dhcp (Network and Access-uplink Operating Mode)

config>service>vprn >if>dhcp (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.

The **no** form of this command returns the system to the default.

Default

no option

remote-id

Syntax

remote-id [mac | string *string*]

no remote-id

Context

config>service>ies >if>dhcp>option (Network and Access-uplink Operating Mode)
config>service>vprn >if>dhcp>option (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command sends the MAC address of the remote end (typically the DHCP client) in the **remote-id** sub-option of the DHCP packet. This command identifies the host at the other end of the circuit.

When disabled, the **remote-id** sub-option 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 the MAC address of the remote end is encoded in the sub-option.

string string

Specifies the remote-id.

vendor-specific-option

Syntax

[no] vendor-specific-option

Context

config>service>ies >if>dhcp>option (Network and Access-uplink Operating Mode)
config>service>vprn >if>dhcp>option (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the vendor specific sub-option of the DHCP relay packet.

client-mac-address

Syntax

[no] client-mac-address

Context

config>service>ies >if>dhcp>option>vendor (Network and Access-uplink Operating Mode)

config>service>vprn >if>dhcp>option>vendor (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command sends the MAC address in the vendor specific sub-option of the DHCP relay packet.

The **no** form of this command disables sending the MAC address in the vendor specific sub-option of the DHCP relay packet.

sap-id

Syntax

[no] sap-id

Context

config>service>ies >if>dhcp>option>vendor (Network and Access-uplink Operating Mode)

config>service>vprn >if>dhcp>option>vendor (not supported in access-uplink operating mode)

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command sends the SAP ID in the vendor specific suboption of the DHCP relay packet.

The **no** form of this command disables sending the SAP ID in the vendor specific suboption of the DHCP relay packet.

allow-directed-broadcasts

Syntax

[no] allow-directed-broadcasts

Context

```
config>service>ies>if
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command forwards directed broadcasts out of the IP interface. A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The `allow-directed-broadcasts` command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface.

When enabled, a frame destined to the local subnet on this IP interface is 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 is counted in the normal discard counters for the egress SAP.

By default, directed broadcasts are not allowed and are discarded at this egress IP interface.

The **no** form of this command disables forwarding of directed broadcasts out of the IP interface.

Default

`no allow-directed-broadcasts` — Directed broadcasts are dropped.

delayed-enable

Syntax

delayed-enable seconds [init-only]

no delayed-enable

Context

```
config>service>ies>if
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

Description

This command delays making interface operational by the specified number of seconds. In environments with many subscribers, it can take time to synchronize the subscriber state between peers when the subscriber-interface is enabled (perhaps, after a reboot). To ensure that the state has time to be synchronized, the `delayed-enable` timer can be specified. The optional parameter **init-only** can be added to use this timer only after a reboot.

Default

no delayed-enable

Parameters***seconds***

Specifies the number of seconds to delay before the interface is operational.

Values 1 to 1200

ip-mtu**Syntax**

ip-mtu *octets*

no ip-mtu

Context

config>service>ies>if

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the maximum IP transmit unit (packet) for the interface.

The MTU that is advertised from the IES size is:

MINIMUM((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))

By default (for Ethernet network interface) if no ip-mtu is configured, the packet size is (1568 - 14) = 1554.

The **no** form of this command returns the default value.

Default

no ip-mtu

Parameters***octets***

Specifies the number of octets in the IP-MTU.

Values 512 to 9000

loopback

Syntax

[no] loopback

Context

config>service>ies>if

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated IES interface cannot be bound to a SAP.

Note that you can configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap** command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

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>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command assigns up to 64 secondary IP addresses to the interface, including the primary IP address. Each address can be configured in an IP address, IP subnet, or broadcast address format.

Default

n/a

Parameters

ip-address

Specifies the IP address of the IP interface. The *ip-address* portion of the address command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

/

The forward slash is a parameter delimiter that separates the *ip-address* portion of the IP address from the *mask* that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "/" and the *mask* parameter. If a forward slash does not immediately follow the *ip-address*, a dotted-decimal *netmask* must follow the prefix.

mask

Specifies the subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask* parameter. The *mask* parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 1 to 32. A mask length of 32 is reserved for system IP addresses.

Values 1 to 32

netmask

Specifies the subnet mask in dotted-decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted-decimal mask. The *netmask* parameter indicates the complete mask that is used in a logical 'AND' function to derive the local subnet of the IP address. A netmask of 255.255.255.255 is reserved for system IP addresses.

Values a.b.c.d (network bits all 1 and host bits all 0)

broadcast {all-ones | host-ones}

The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones**, which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the broadcast type to **host-ones** after being configured as **all-ones**, the **address** command must be executed with the **broadcast** parameter defined. 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**) is received by the IP interface

Values **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* or *netmask* with all of the host bits set to binary 1. This is the default broadcast address used by an IP interface.

Default	host-ones

igp-inhibit

Specifies that the secondary IP address should not be recognized as a local interface by the running IGP.

static-arp

Syntax

static-arp *ip-address ieee-mac-address*

no static-arp *ip-address*

Context

config>service>ies>if

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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 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 with the new MAC address.

The **no** form of this command removes a static ARP entry.

Parameters

ip-address

Specifies the IP address for the static ARP in IP address dotted-decimal notation.

ieee-mac-address

Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

static-arp

Syntax

[no] static-arp *ieee-mac-addr unnumbered*

Context

config>service>ies>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures a static address resolution protocol (ARP) entry associating an unnumbered interface 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 unnumbered interface.

If an entry for a particular unnumbered interface already exists, and a new MAC address is configured for the interface, the existing MAC address is replaced with the new MAC address.

The **no** form of this command removes a static ARP entry.

Parameters

ieee-mac-addr

Specifies the 48-bit MAC address for the static ARP in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

unnumbered

Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. When this command is configured, it overrides any dynamic ARP.

vpls

Syntax

vpls *service-name*

Context

config>service
config>service>ies>if

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command, within the IP interface context, is used to bind the IP interface to the specified service name.

The system does not attempt to resolve the service name provided until the IP interface is placed into the administratively up state (no shutdown). When the IP interface is administratively up, the system scans the available VPLS services that have the allow-ip-int-binding flag set for a VPLS service associated with the name. If the service name is bound to the service name when the IP interface is already in the administratively up state, the system immediately attempts to resolve the specific name.

If a VPLS service is found associated with the name and with the allow-ip-int-binding flag set, the IP interface is attached to the VPLS service allowing routing to and from the service virtual ports when the IP interface is operational.

A VPLS service associated with the specified name that does not have the allow-ip-int-binding flag set or a non-VPLS service associated with the name is ignored and is not attached to the IP interface.

If the service name is applied to a VPLS service after the service name is bound to an IP interface and the VPLS service allow-ip-int-binding flag is set at the time the name is applied, the VPLS service is automatically resolved to the IP interface if the interface is administratively up or when the interface is placed in the administratively up state.

If the service name is applied to a VPLS service without the allow-ip-int-binding flag set, the system does not attempt to resolve the applied service name to an existing IP interface bound to the name. To rectify this condition, the flag must first be set and then the IP interface must enter or reenter the administratively up state.

While the specified service name may be assigned to only one service context in the system, it is possible to bind the same service name to more than one IP interface. If two or more IP interfaces are bound to the same service name, the first IP interface to enter the administratively up state (if currently administratively down) or to reenter the administratively up state (if currently administratively up) when a VPLS service is configured with the name and has the allow-ip-int-binding flag set is attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

When an IP interface is attached to a VPLS service, the name associated with the service cannot be removed or changed until the IP interface name binding is removed. Also, the allow-ip-int-binding flag cannot be removed until the attached IP interface is unbound from the service name. Unbinding the service name from the IP interface causes the IP interface to detach from the VPLS service context. The IP interface may then be bound to another service name or a SAP or SDP binding may be created for the interface using the sap or spoke-sdp commands on the interface.

Parameters

service-name

Specifies the service-name parameter is required when using the IP interface vpls command and specifies the service name that the system attempts to resolve to an allow-ip-int-binding enabled VPLS service associated with the name. The specified name is expressed as an ASCII string comprised of up to 32 characters. It does not need to already be associated with a service and the system does not check to ensure that multiple IP interfaces are not bound to the same name.

ingress

Syntax

ingress

Context

config>service>ies>if>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

The command under the vpls binding context defines the routed ip-filter-id optional filter overrides.

enable-table-classification

Syntax

[no] enable-table-classification

Context

config>service>ies>if>vpls>ingress

Platforms

7210 SAS-Mxp

Description

This command enables and disables the use of IP DSCP table-based classification to assign FC and profile on a per-interface ingress basis.

The match-criteria configured in the service ingress policy, which require CAM resources, are ignored. Only meters from the service ingress policy are used (and the meters still require CAM resources). If an IP DSCP classification policy is configured in the VPLS SAP ingress policy, it is not used to assign FC and profile.

The **no** form of this command disables table-based classification. When disabled, the IP ingress packets within a VPLS service attached to the IP interface use the SAP ingress QoS policy applied to the virtual port used by the packets, when defined.

Default

no enable-table-classification

routed-override-qos-policy

Syntax

routed-override-qos-policy *policy-id*
no routed-override-qos-policy

Context

config>service>ies>if>vpls>ingress

Platforms

7210 SAS-Mxp

Description

This command specifies an IP DSCP classification policy that is applied to all ingress packets entering the VPLS service. The DSCP classification policy overrides any existing SAP ingress QoS policy applied to SAPs for packets associated with the routing IP interface. The routed override QoS policy is optional and when it is not defined or it is removed, the IP routed packets use the existing SAP ingress QoS policy configured on the VPLS virtual port.

The **no** form of this command is used to remove the IP DSCP classification policy from the ingress IP interface. When removed, the IP ingress routed packets within a VPLS service attached to the IP interface use the SAP ingress QoS policy applied to the virtual port used by the packets, when defined.

Default

no routed-override-qos-policy

Parameters

policy-id

Specifies the ID for the routed override QoS policy. Allowed values are an integer that corresponds to an existing IP DSCP classification policy in the **configure>qos>dscp-classification** context.

Values 1 to 65535

v4-routed-override-filter

Syntax

v4-routed-override-filter *ip-filter-id*
no v4-routed-override-filter

Context

config>service>ies>if>vpls>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command specifies an IPv4 filter ID applied to all ingress packets entering the VPLS service. The filter overrides existing ingress IPv4 filters applied to SAPs for packets associated with the routing IP interface. The override filter is optional, and if it is not defined or is removed, the IPv4 routed packets use the existing ingress IP filter on the VPLS virtual ports.

The **no** form of this command removes the IP routed override filter from the ingress IP interface.

Parameters

ip-filter-id

Specifies the integer filter ID value for the IPv4 filter policy. The filter ID must already exist within the IPv4 filters created in the **configure>filter>ip-filter** context.

Values 1 to 65535

v6-routed-override-filter

Syntax

v6-routed-override-filter *ip-filter-id*

no v6-routed-override-filter

Context

config>service>ies>if>vpls>ingress

Platforms

7210 SAS-Mxp

Description

This command specifies an IPv6 filter ID applied to all ingress packets entering the VPLS service. The filter overrides existing ingress IPv6 filters applied to SAPs for packets associated with the routing IP interface. The override filter is optional, and if it is not defined or is removed, the IPv6 routed packets use the existing ingress IP filter on the VPLS virtual ports.

The **no** form of this command removes the IP routed override filter from the ingress IP interface.

Parameters

ip-filter-id

Specifies the integer filter ID value for the IPv6 filter policy. The filter ID must already exist within the IPv6 filters created in the **configure>filter>ip-filter** context.

Values 1 to 65535

7.4.2.1.5 Interface VRRP commands

vrrp

Syntax

vrrp *virtual-router-id* [**owner**]

no vrrp *virtual-router-id*

Context

config>service>ies>if

config>service>ies>if>ipv6 (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures a Virtual Router ID (VRID) on the service IP interface. A VRID is internally represented in conjunction with the IP interface name. This allows the VRID to be used on multiple IP interfaces while representing different virtual router instances.

Two VRRP nodes can be defined on an IP interface. One, both, or none may be defined as owner. The nodal context of **vrrp** *virtual-router-id* is used to define the configuration parameters for the VRID.

The **no** form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the VRID. The VRID does not need to be shut down to remove the virtual router instance.

Parameters

virtual-router-id

Specifies a new virtual router ID or one that can be modified on the IP interface.

Values 1 to 255

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

config>service>ies>if>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command, within the **vrrp** *virtual-router-id* context, is used to assign a simple text password authentication key to generate master VRRP advertisement messages and validate received VRRP advertisement messages.

The **authentication-key** command is one of the few commands not affected by the presence of the **owner** keyword. If simple text password authentication is not required, this command is not required. 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, altering the simple text password used when **authentication-type** password authentication method is used by the virtual router instance. The **authentication-type password** command does not need to be executed before defining the **authentication-key** command.

To change the current in-use password key on multiple virtual router instances:

- Identify the current master.
- Shutdown 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 key.

The **no** form of this command restores the default null string to the value of key.

Default

No default. The authentication data field contains the value 0 in all 16 octets.

Parameters

authentication-key

Specifies the *key* parameter 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 *key* parameter is expressed as a string consisting of up to eight alpha-numeric characters. Spaces must be contained in quotation marks (" "). The quotation marks are not considered part of the string.

The string is case sensitive and is left-justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet (starting with IETF RFC bit position 0) 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

The hash key. The key can be any combination of ASCII characters up to 22 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks ("").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

backup**Syntax**

[no] **backup** *ip-address*

Context

config>service>ies>if>vrrp

config>service>ies>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures virtual router IP addresses for the interface.

bfd-enable**Syntax**

[no] **bfd-enable** [*service-id*] **interface** *interface-name* **dst-ip** *ip-address*

Context

```
config>service>ies>if>vrrp
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command assigns a bi-directional forwarding (BFD) session, providing a heart-beat mechanism for the VRRP instance. There can only be one BFD session assigned to a specified VRRP instance, but multiple VRRP instances can use the same BFD session. If the specified interface is configured with centralized BFD, the BFD transmit and receive intervals must be 300 ms or longer.

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 configured using the BFD command under the IP interface. The virtual router initiates the BFD session after the specified interface is configured with BFD.

Parameters***service-id***

Specifies the service ID of the interface that is running BFD.

Values *service-id* — 1 to 2147483648

svc-name — Specifies an existing service name up to 64 characters in length.

interface-name

Specifies the name of the interface that is 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 (7210 SAS-Mxp only)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures a VRRP initialization delay timer.

Default

no init-delay

Parameters

seconds

Specifies the initialization delay timer for VRRP, in seconds.

Values 1 to 65535

master-int-inherit

Syntax

[no] master-int-inherit

Context

config>service>ies>if>vrrp

config>service>ies>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command allows the master instance to dictate the master down timer (non-owner context only).

Default

no master-int-inherit

message-interval

Syntax

message-interval {[seconds] [milliseconds milliseconds]}

no message-interval

Context

config>service>ies>if

config>service>ies>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

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 participating as a virtual router. Any VRRP advertisement message received with an Advertisement Interval field different from the virtual router instance configured message-interval value is silently discarded.

The **message-interval** command is available in both non-owner and owner **vrrp** *virtual-router-id* nodal contexts. If the message-interval command is not executed, the default message interval of 1 second is used.

The **no** form of this command restores the default message interval value of 1 second to the virtual router instance.

Parameters

seconds

Specifies the number of seconds that transpire before the advertisement timer expires.

Values 1 to 255

Default 1

milliseconds *milliseconds*

Specifies the milliseconds time interval between sending advertisement messages. This parameter is not supported on single-slot chassis.

Values 100 to 900

ping-reply

Syntax

[no] ping-reply

Context

config>service>ies>if>vrrp

config>service>ies>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the non-owner master to reply to ICMP Echo Requests directed at the virtual router instances IP addresses. The ping request can be received on any routed interface.

Ping must not be disabled at the management security level (either on the parental IP interface or based on the Ping source host address). When ping-reply is not enabled, ICMP Echo Requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP Echo Requests regardless of the setting of ping-reply configuration.

The ping-reply command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the ping-reply command is not executed, ICMP Echo Requests to the virtual router instance IP addresses are silently discarded.

The **no** form of this command restores the default operation of discarding all ICMP Echo Request messages destined to 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 (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command associates a VRRP priority control policy with the virtual router instance (non-owner context only).

Parameters

vrrp-policy-id

Specifies a VRRP priority control policy.

Values 1 to 9999

preempt

Syntax

preempt

no preempt

Context

```
config>service>ies>if>vrrp
```

```
config>service>ies>if>ipv6>vrrp (7210 SAS-Mxp only)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command overrides an existing non-owner master to the virtual router instance. Enabling preempt mode is recommended for correct operation of the base-priority and vrrp-policy-id definitions on the virtual router instance. If the virtual router cannot preempt an existing non-owner master, the affect of the dynamic changing of the in-use priority is greatly diminished.

The **preempt** command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. The owner may not be preempted due to the fact that 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 virtual router instances only preempts when preempt is set and the current master has an in-use message priority value less than the virtual router instances 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 the virtual router in-use priority value
- Equal to the in-use priority value and the source IP address (primary IP address) is greater than the virtual router instance primary IP address

The **no** form of this command prevents a non-owner virtual router instance from preempting another, less desirable virtual router. Use the preempt command to restore the default mode.

Default

preempt

priority

Syntax

priority *priority*

no priority

Context

```
config>service>ies>if>vrrp
```

```
config>service>ies>if>ipv6>vrrp (7210 SAS-Mxp only)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures a specific priority value to the virtual router instance. In conjunction with an optional policy command, the base-priority is used to derive the in-use priority of the virtual router instance.

The **priority** command is only available in the non-owner **vrrp virtual-router-id** nodal context. 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 to base-priority.

Parameters

base-priority

Specifies the base-priority parameter configures the base priority used by the virtual router instance. If a VRRP priority control policy is not also defined, the base-priority is the 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

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the non-owner master to reply to SSH Requests directed at the virtual router instance IP addresses. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parental IP interface or based on the SSH source host address). Correct login and CLI command authentication is still 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 regardless of the ssh-reply configuration.

The **ssh-reply** command is only available in non-owner **vrrp virtual-router-id** nodal context. If the ssh-reply command is not executed, SSH packets to the virtual router instance IP addresses are silently discarded.

The **no** form of this command restores the default operation of discarding all SSH packets destined to the non-owner virtual router instance IP addresses.

Default

no ssh-reply

standby-forwarding**Syntax**

[no] standby-forwarding

Context

config>service>ies>if>vrrp

config>service>ies>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the forwarding of packets by a standby router.

The **no** form of this command specifies that a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router real MAC address.

Default

no standby-forwarding

telnet-reply**Syntax**

[no] telnet-reply

Context

config>service>ies>if>vrrp

config>service>ies>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the non-owner master to reply to TCP port 23 Telnet Requests directed at the virtual router instance IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Correct login and CLI command authentication is still enforced.

When 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 in non-owner VRRP nodal context. If the telnet-reply command is not executed, Telnet packets to the virtual router instance IP addresses are silently discarded.

The **no** form of this command restores the default operation of discarding all Telnet packets destined to 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 (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.

A non-owner backup virtual router never responds to such traceroute requests regardless of the **traceroute-reply** command status.

Default

no traceroute-reply

7.4.2.1.6 IES interface ICMP commands

icmp

Syntax

icmp

Context

config>service>ies>if

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context configure Internet Control Message Protocol (ICMP) parameters on an IES service

mask-reply

Syntax

[no] mask-reply

Context

config>service>ies>if>icmp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance replies to mask requests.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

Default

mask-reply

redirects

Syntax

redirects [*number seconds*]

no redirects

Context

config>service>ies>if>icmp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.

When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

The **redirects** command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of redirect messages that can be issued on the interface for a specific time interval.

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default

redirects 100 10

Parameters

number

Specifies the maximum number of ICMP redirect messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 to 1000

seconds

Specifies the time frame in seconds used to limit the *number* of ICMP redirect messages that can be issued.

Values 1 to 60

ttl-expired

Syntax

ttl-expired *number seconds*

no ttl-expired

Context

config>service>ies>if>icmp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default

ttl-expired 100 10

Parameters

number

Specifies the maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 to 1000

seconds

Specifies 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
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The **unreachables** command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *time* parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a specific time interval.

By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 10 per 60 second time interval.

The **no** form of this command disables the generation of icmp destination unreachable messages on the router interface.

Default

```
unreachables 100 10
```

Parameters***number***

Specifies the maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 to 1000

seconds

Specifies the time frame in seconds used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 to 60

7.4.2.1.7 IES SAP commands

```
sap
```

Syntax

```
sap sap-id [create]
```

```
no sap sap-id
```

Context

config>service>ies>if

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP does not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access uplink port using the **configure port port number ethernet mode access uplink** command.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service is 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 is also deleted.

Default

No SAPs are defined.

Special Cases

IES

A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Attempts to create a second SAP on an IP interface fail and generate an error; the original SAP is not affected.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

port-id

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

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the *slot_number/MDA_number/port_number* format. For example 1/1/1 specifies port 1 on MDA 1 in slot 1.

The *port-id* must reference a valid port type. The port must be configured as an uplink access port.

create

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

7.4.2.1.8 IES QoS and filter commands

egress

Syntax

egress

Context

config>service>ies>if>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context apply egress policies and filter policies.

ingress

Syntax

ingress

Context

config>service>ies>if>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Commands in this context apply ingress policies and filter policies.

filter

Syntax

filter ip *ip-filter-id* **ipv6** *ipv6-filter-id*
no filter

Context

config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command associates a filter policy with an ingress or egress Service Access Point (SAP). Filter policies control the forwarding and dropping of packets based on the matching criteria.

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

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to the match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP. The filter ID is not removed from the system.

Special Cases

IES

Only IP filters are supported on an IES IP interface, and the filters only apply to routed traffic.

Parameters

ip *ip-filter-id*

Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the **configure>filter>ip-filter** context.

Values 1 to 65535

ipv6 *ipv6-filter-id*

Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

qos

Syntax

qos *policy-id*

qos *policy-id* [**enable-table-classification**]

no qos *policy-id*

Context

config>service>ies>if>sap>egress (7210 SAS-Mxp only)

config>service>ies>if>sap>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP) or IP interface. QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined before associating the policy with a SAP or IP interface. If the *policy-id* does not exist, an error is returned.

The **qos** command is used to associate both ingress and egress QoS policies. The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress, and only allows egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second policy of same or different type replaces the earlier one with the new policy.



Note:

SAP egress QoS policies are only supported on the 7210 SAS-Mxp.

On the 7210 SAS-Mxp (ingress), using the **enable-table-classification** keyword enables the use of IP DSCP tables to assign FC and profile on a per-SAP ingress basis. The match-criteria configured in the service ingress policy, which require CAM resources, are ignored. Only meters from the service ingress policy are used (and the meters still require CAM resources). The IP DSCP classification policy configured in the SAP ingress policy is used to assign FC and profile. The default FC is assigned from the SAP ingress policy.

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



Note:

On the 7210 SAS-Mxp, when the interface is associated with RVPLS, the behavior of the **qos** command is affected. See the **config>service>ies>if>vpls> ingress>enable-table-classification** and **routed-override-qos-policy** commands for more information about classification behavior for RVPLS.

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

Parameters

policy-id

Specifies the ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.

Values 1 to 65535

enable-table-classification

Enables the use of table-based classification instead of CAM-based classification at SAP ingress. The FC and profile are taken from the IP DSCP classification policy configured in the ingress policy, along with the meters from the SAP ingress policy. Match-criteria entries in the SAP ingress policy are ignored.

tod-suite

Syntax

tod-suite *tod-suite-name*

no tod-suite

Context

config>service>ies>if>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

This command applies a time-based policy (filter or QoS policy) to the service SAP. The suite name must already exist in the **config>cron** context.

Default

no tod-suite

Parameters

tod-suite-name

Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP. The suite can be applied to more than one SAP.

7.4.2.1.9 IES interface SAP statistics commands

statistics

Syntax

statistics

Context

config>service>ies>if>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

Commands in this context configure the counters associated with SAP ingress and egress.

ingress

Syntax

ingress

Context

config>service>ies>if>sap>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

Description

Commands in this context configure the ingress SAP statistics counters.

counter-mode

Syntax

counter-mode {in-out-profile-count | forward-drop-count}

Context

config>service>ies>if>sap>statistics>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command sets the counter mode for the counters associated with sap ingress meters or policers. A pair of counters is available with each meter. These counters count different events based on the counter mode value.



Note:

The counter mode can be changed if an accounting policy is associated with a SAP. If the counter mode is changed, the counters associated with the meter are reset and the counts are cleared. If an accounting policy is in use when the counter-mode is changed, a new record is written into the current accounting file.

Execute the following sequence of commands on the specified SAP to ensure that the correct statistics are collected when the counter-mode is changed:

1. Execute the **config service ies interface sap no collect-stats** command, to disable writing of accounting records for the SAP.
2. Change the counter mode to the desired option by executing the **config service ies interface sap statistics ingress counter-mode {in-out-profile-count | forward-drop-count}** command.
3. Execute the **config service ies interface sap collect-stats** command, to enable writing of accounting records for the SAP.

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

Default

in-out-profile-count

Parameters

in-out-profile-count

If the counter mode is specified as **in-out-profile-count**, one counter counts the total in-profile packets and octets received on ingress of a SAP and another counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

forward-drop-count

If the counter mode is specified as **forward-drop-count**, one counter counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

7.4.2.2 Show commands

customer

Syntax

customer [*customer-id*] [**site** *customer-site-name*]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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

site customer-site-name

Specifies the customer site which is an anchor point for an ingress and egress virtual scheduler hierarchy.

Output

The following output is an example of service customer information, and [Table 99: Output Fields: Customer](#) describes the output fields.

Sample output

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

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

```

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

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

Customer-ID : 7
Contact      : Lucy
Description  : ABC Customer
Phone       : (567) 555-1212

Customer-ID : 8
Contact      : Customer Service
Description  : IES Customer
Phone       : (678) 555-1212

Customer-ID : 274
Contact      : Mssrs. Beaucoup
Description  : ABC Company
Phone       : 650 123-4567

Customer-ID : 94043
Contact      : Test Engineer on Duty
Description  : TEST Customer
Phone       : (789) 555-1212

```

```

-----
Total Customers : 8
-----

```

```

*A:ALA-12#

```

```

*A:ALA-12# show service customer 274

```

```

=====
Customer  274
=====

```

```

Customer-ID : 274
Contact      : Mssrs. Beaucoup
Description  : ABC Company
Phone       : 650 123-4567

```

```

-----
Multi Service Site
-----

```

```

Site        : west
Description  : (Not Specified)

```

```

=====
*A:ALA-12#

```

```

*A:ALA-12# show service customer 274 site west

```

```

=====
Customer  274
=====

```

```

Customer-ID : 274
Contact      : Mssrs. Beaucoup
Description  : ABC Company
Phone       : 650 123-4567

```

```

-----
Multi Service Site
-----

```

```

Site      : west
Description : (Not Specified)
Assignment : Card 5
I. Sched Pol: SLA1
E. Sched Pol: (Not Specified)

```

```

-----
Service Association
-----

```

```

No Service Association Found.
=====

```

```

*A:ALA-12#

```

Table 99: Output Fields: Customer

Label	Description
Customer-ID	The ID that uniquely identifies a customer.
Contact	The name of the primary contact person.
Description	Generic information about the customer.
Phone	The phone/pager number to reach the primary contact person.
Total Customers	The total number of customers configured.
Multi-service site	
Site	Multi-service site name. A multi-service customer site is a group of SAPs with common origination and termination points.
Description	Information about a specific customer's multi-service site.
Assignment	The port ID, MDA, or card number, where the SAPs that are members of this multi-service site are defined.
I. Sched Pol	The ingress QoS scheduler policy assigned to this multi-service site.
E. Sched Pol	The egress QoS scheduler policy assigned to this multi-service site.
Service Association	
Service-ID	The ID that uniquely identifies a service.
SAP	Specifies the SAP assigned to the service.

sap-using

Syntax

sap-using [sap sap-id]

sap-using interface [ip-address | ip-int-name]

sap-using [ingress | egress] **filter** *filter-id*
sap-using [ingress] **qos-policy** *qos-policy-id*

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

Displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs. The optional parameters restrict output to only SAPs matching the specified properties.

Parameters

sap *sap-id*

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

ingress

Specifies matching an ingress policy.

egress

Specifies matching an egress policy.

filter *filter-id*

Specifies the ingress or egress filter policy ID for which to display matching SAPs.

Values 1 to 65535

interface

Specifies matching SAPs with the specified IP interface.

ip-addr

Specifies the IP address of the interface for which to display matching SAPs.

Values a.b.c.d

ip-int-name

Specifies the IP interface name for which to display matching SAPs.

Output

The following output is an example of SAP information, and [Table 100: Output Fields: SAP-using](#) describes the output fields.

Sample output

```
*A:DUT-B# show service sap-using sap 1/1/3:100.*
=====
Service Access Points
```

PortId	SvcId	Ing. QoS	Ing. Fltr	Egr. Fltr	Adm	Opr
1/1/1	6	1	none	none	Up	Down
1/1/2	700	1	none	none	Up	Down
Number of SAPs : 2						
*A:DUT-B#						

Table 100: Output Fields: SAP-using

Label	Description
Port ID	The ID of the access port where the SAP is defined.
Svc ID	The value that identifies the service.
SapMTU	The SAP MTU value.
Igr.QoS	The SAP ingress QoS policy number specified on the ingress SAP.
Ing.Fltr	The MAC or IP filter policy ID applied to the ingress SAP.
E.QoS	The SAP egress QoS policy number specified on the egress SAP.
Egr.Fltr	The MAC or IP filter policy ID applied to the egress SAP.
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The administrative state of the SAP.
Opr	The actual state of the SAP.

service-using

Syntax

service-using [ies] [customer *customer-id*]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays the services matching certain usage properties. If no optional parameters are specified, all services defined on the system are displayed.

Parameters

ies

Displays matching IES services.

customer *customer-id*

Displays services only associated with the specified customer ID.

Values 1 to 2147483647

Default Services associated with an customer.

Output

The following output is an example of service information, and [Table 101: Output Fields: Service-using](#) describes the output fields.

Sample output

```
A:ALA-48# show service service-using ies
=====
Services [ies]
=====
ServiceId   Type      Adm      Opr      CustomerId  Last Mgmt Change
-----
88          IES       Up       Down     8            07/25/2006 15:46:28
89          IES       Up       Down     8            07/25/2006 15:46:28
104         IES       Up       Down     1            07/25/2006 15:46:28
200         IES       Up       Down     1            07/25/2006 15:46:28
214         IES       Up       Down     1            07/25/2006 15:46:28
321         IES       Up       Down     1            07/25/2006 15:46:28
322         IES       Down    Down     1            07/25/2006 15:46:28
1001        IES       Up       Down     1730         07/25/2006 15:46:28
-----
Matching Services : 8
-----
A:ALA-48#
```

Table 101: Output Fields: Service-using

Label	Description
Service Id	The value that identifies the service.
Type	Specifies the service type configured for the service ID.
Adm	The administrative state of the service.
Opr	The operating state of the service.
CustomerId	The ID of the customer who owns this service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

id

Syntax

id *service-id* {**all** | **arp** | **base** | **sap** | **interface** | **mstp-configuration**}

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays information for a particular service-id.

Parameters

service-id

The unique service identification number to identify the service in the service domain.

all

Display detailed information about the service.

arp

Display ARP entries for the service.

base

Display basic service information.

interface

Display service interfaces.

mstp-configuration

Display MSTP information.

sap

Display SAPs associated to the service.

split-horizon-group

Display split horizon group information.

all

Syntax

all

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays detailed information for all aspects of the service.

Output

The following output is an example of detailed service information, and [Table 102: Output Fields: Service ID All](#) describes the output fields.

Sample output (split horizon group)

```
*A:SAS>show>service# id 10 all

=====
Service Detailed Information
=====
Service Id       : 10                Vpn Id          : 0
Service Type     : VPLS
Description      : (Not Specified)
Customer Id      : 1
Last Status Change: 07/22/2011 11:06:02
Last Mgmt Change : 07/22/2011 11:04:51
Admin State      : Up                Oper State       : Up
MTU              : 1450
MTU Check        : Enabled
SAP Count        : 2                SDP Bind Count   : 2
Snd Flush on Fail : Disabled
Uplink Type      : MPLS

-----
Split Horizon Group specifics
-----
Split Horizon Group : test
-----
Description        : test
Instance Id        : 1                Last Change      : 07/23/2011 11:40:50
-----
Service Destination Points(SDPs)
-----
Sdp Id 2:10  -(10.20.1.6)
-----
Description       : (Not Specified)
SDP Id            : 2:10                Type             : Spoke
Split Horiz Grp   : (Not Specified)
VC Type          : VLAN                VC Tag           : 10
Admin Path MTU    : 0                  Oper Path MTU    : 9186
Far End           : 10.20.1.6          Delivery         : MPLS

Admin State       : Up                Oper State       : Up
Acct. Pol         : None              Collect Stats    : Disabled
Ingress Label     : 131063            Egress Label     : 131067
Admin ControlWord : Preferred          Oper ControlWord : True
Last Status Change: 07/22/2011 11:07:26
Last Mgmt Change  : 07/22/2011 11:04:51
Flags             : None              Signaling        : TLDP
Force Vlan-Vc     : Disabled
```

```

Peer Pw Bits      : None
Peer Fault Ip     : None
Max Nbr of MAC Addr: No Limit
Learned MAC Addr  : 0
Total MAC Addr    : 0
Static MAC Addr   : 0

MAC Learning      : Enabled
BPDU Translation  : Disabled
L2PT Termination  : Disabled
MAC Pinning       : Disabled
MAC Pinning       : Disabled
Discard Unkwn Srce: Disabled

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

Statistics        :
I. Fwd. Pkts.     : 0
E. Fwd. Pkts.     : 1
Extra-Tag-Drop-Pkts: n/a
I. Fwd. Octets    : 0
E. Fwd. Octets    : 98
Extra-Tag-Drop-Octets*: n/a

Associated LSP LIST :
Lsp Name          : toF
Admin State       : Up
Oper State        : Up
-----
Stp Service Destination Point specifics
-----
Stp Admin State   : Up
Core Connectivity : Down
Port Role         : Designated
Port Number       : 2049
Port Path Cost    : 10
Admin Edge        : Disabled
Link Type         : Pt-pt
Root Guard        : Disabled
Last BPDU from    : N/A
Designated Bridge : This Bridge
Stp Oper State    : Up
Port State        : Forwarding
Port Priority      : 128
Auto Edge         : Enabled
Oper Edge         : True
BPDU Encap        : Dot1d
Active Protocol   : Rstp
Designated Port Id: 34817

Fwd Transitions   : 1
Cfg BPDUs rcvd    : 0
TCN BPDUs rcvd    : 0
TC bit BPDUs rcvd : 0
RST BPDUs rcvd    : 0
Bad BPDUs rcvd    : 0
Cfg BPDUs tx      : 0
TCN BPDUs tx      : 0
TC bit BPDUs tx   : 0
RST BPDUs tx      : 44265
-----
Sdp Id 4:10  -(10.20.1.3)
-----
Description       : (Not Specified)
SDP Id           : 4:10
Split Horiz Grp   : (Not Specified)
VC Type          : VLAN
Admin Path MTU    : 0
Far End          : 10.20.1.3
Type             : Spoke
VC Tag           : 10
Oper Path MTU     : 9182
Delivery         : MPLS

Admin State       : Up
Acct. Pol         : None
Ingress Label     : 131059
Admin ControlWord : Preferred
Last Status Change : 07/22/2011 11:07:26
Last Mgmt Change  : 07/22/2011 11:04:51
Flags            : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Max Nbr of MAC Addr: No Limit
Learned MAC Addr  : 0
Oper State        : Up
Collect Stats     : Disabled
Egress Label      : 131065
Oper ControlWord  : True
Signaling         : TLDP
Force Vlan-Vc     : Disabled
Total MAC Addr    : 0
Static MAC Addr   : 0

```

```

MAC Learning      : Enabled          Discard Unkwn Srce: Disabled
BPDU Translation  : Disabled
L2PT Termination  : Disabled
MAC Pinning       : Disabled
MAC Pinning       : Disabled          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.     : 44285              I. Fwd. Octs.    : 3852802
E. Fwd. Pkts.     : 0                  E. Fwd. Octets   : 0
Extra-Tag-Drop-Pkts: n/a              Extra-Tag-Drop-0c*: n/a

Associated LSP LIST :
Lsp Name          : toh2_facility
Admin State       : Up                Oper State       : Up
Time Since Last Tr*: 01d00h37m

```

Stp Service Destination Point specifics

```

Stp Admin State   : Up                Stp Oper State    : Up
Core Connectivity : Down
Port Role         : Root              Port State        : Forwarding
Port Number       : 2050              Port Priority      : 128
Port Path Cost    : 10                Auto Edge         : Enabled
Admin Edge        : Disabled           Oper Edge         : False
Link Type         : Pt-pt             BPDU Encap        : Dot1d
Root Guard        : Disabled           Active Protocol    : Rstp
Last BPDU from    : 80:01.00:25:ba:02:de:90
Designated Bridge : 80:01.00:25:ba:02:de:90 Designated Port Id: 34817

Fwd Transitions   : 1                 Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd    : 0                 Cfg BPDUs tx      : 0
TCN BPDUs rcvd    : 0                 TCN BPDUs tx      : 0
TC bit BPDUs rcvd : 2                 TC bit BPDUs tx   : 2
RST BPDUs rcvd    : 44284             RST BPDUs tx      : 3

```

Number of SDPs : 2

Service Access Points

SAP 1/1/2

```

Service Id        : 10
SAP               : 1/1/2              Encap             : null
Description       : (Not Specified)
Admin State       : Up                Oper State        : Down
Flags            : PortOperDown
Last Status Change : 07/22/2011 11:04:50
Last Mgmt Change  : 07/23/2011 11:42:22
Dot1Q Ethertype   : 0x8100           QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Max Nbr of MAC Addr: No Limit          Total MAC Addr    : 0
Learned MAC Addr   : 0                 Static MAC Addr    : 0

```

Admin MTU	: 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		
Mac Learning	: Enabled	Discard Unkwn Srce:	Disabled
Mac Aging	: Enabled	Mac Pinning	: Disabled
BPDU Translation	: Disabled		
L2PT Termination	: Disabled		

Acct. Pol	: None	Collect Stats	: Disabled

----- Stp Service Access Point specifics -----

Stp Admin State	: Up	Stp Oper State	: Up
Core Connectivity	: Down		
Port Role	: Disabled	Port State	: Discarding
Port Number	: 2051	Port Priority	: 128
Port Path Cost	: 10	Auto Edge	: Enabled
Admin Edge	: Disabled	Oper Edge	: False
Link Type	: Pt-pt	BPDU Encap	: Dot1d
Root Guard	: Disabled	Active Protocol	: Rstp
Last BPDU from	: N/A		
CIST Desig Bridge	: N/A	Designated Port	: 0
Forward transitions:	0	Bad BPDUs rcvd	: 0
Cfg BPDUs rcvd	: 0	Cfg BPDUs tx	: 0
TCN BPDUs rcvd	: 0	TCN BPDUs tx	: 0
TC bit BPDUs rcvd	: 0	TC bit BPDUs tx	: 0
RST BPDUs rcvd	: 0	RST BPDUs tx	: 0
MST BPDUs rcvd	: 0	MST BPDUs tx	: 0

----- ARP host -----

Admin State	: outOfService		
Host Limit	: 1	Min Auth Interval	: 15 minutes

----- QoS -----

Ingress qos-policy : 1

----- Aggregate Policer -----

rate	: n/a	burst	: n/a

----- Ingress QoS Classifier Usage -----

Classifiers Allocated:	4	Meters Allocated	: 2
Classifiers Used	: 2	Meters Used	: 2

----- Sap Statistics -----

	Packets	Octets
Ingress Stats:	0	0
Egress Stats:	0	0
Ingress Drop Stats:	0	0

Extra-Tag Drop Stats:	n/a	n/a

----- Sap per Meter stats -----

Packets	Octets

```

Ingress Meter 1 (Unicast)
For. InProf      : 0
For. OutProf     : 0

Ingress Meter 11 (Multipoint)
For. InProf      : 0
For. OutProf     : 0

-----
SAP 1/1/7:10
-----
Service Id       : 10
SAP              : 1/1/7:10
Description      : (Not Specified)
Admin State      : Up
Flags           : None
Last Status Change : 07/22/2011 11:06:02
Last Mgmt Change  : 07/22/2011 11:04:51
Dot1Q Ethertype  : 0x8100
Split Horizon Group: (Not Specified)

Max Nbr of MAC Addr: No Limit
Learned MAC Addr   : 0
Admin MTU          : 1518
Ingr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a
tod-suite         : None
Mac Learning       : Enabled
Mac Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled

Total MAC Addr    : 2
Static MAC Addr   : 2
Oper MTU          : 1518
Egr IP Fltr-Id    : n/a
Egr Mac Fltr-Id   : n/a
Egr IPv6 Fltr-Id  : n/a
Discard Unkwn Srce: Disabled
Mac Pinning       : Disabled

Acct. Pol         : None
Collect Stats     : Disabled

-----
Stp Service Access Point specifics
-----
Stp Admin State   : Up
Core Connectivity : Down
Port Role         : Designated
Port Number       : 2048
Port Path Cost    : 10
Admin Edge        : Disabled
Link Type         : Pt-pt
Root Guard        : Disabled
Last BPDU from    : N/A
CIST Desig Bridge : This Bridge

Stp Oper State    : Up
Port State        : Forwarding
Port Priority      : 128
Auto Edge         : Enabled
Oper Edge         : True
BPDU Encap        : Dot1d
Active Protocol    : Rstp
Designated Port   : 34816

Forward transitions: 1
Cfg BPDUs rcvd    : 0
TCN BPDUs rcvd    : 0
TC bit BPDUs rcvd : 0
RST BPDUs rcvd    : 0
MST BPDUs rcvd    : 0

Bad BPDUs rcvd    : 0
Cfg BPDUs tx      : 0
TCN BPDUs tx      : 0
TC bit BPDUs tx   : 0
RST BPDUs tx      : 44379
MST BPDUs tx      : 0

-----
ARP host
-----
Admin State       : outOfService
Host Limit        : 1
Min Auth Interval : 15 minutes

-----
QOS

```

```

-----
Ingress qos-policy : 1
-----
Aggregate Policer
-----
rate           : n/a           burst           : n/a
-----
Ingress QoS Classifier Usage
-----
Classifiers Allocated: 4           Meters Allocated : 2
Classifiers Used    : 2           Meters Used      : 2
-----
Sap Statistics
-----
Ingress Stats:      Packets      Octets
Egress Stats:      0             0
Ingress Drop Stats: 1             68
Extra-Tag Drop Stats: n/a         n/a
-----
Sap per Meter stats
-----
Ingress Meter 1 (Unicast)
For. InProf        : 0           0
For. OutProf       : 0           0
Ingress Meter 11 (Multipoint)
For. InProf        : 0           0
For. OutProf       : 0           0
-----
VPLS Spanning Tree Information
-----
VPLS oper state   : Up           Core Connectivity : Down
Stp Admin State   : Up           Stp Oper State    : Up
Mode              : Rstp         Vcp Active Prot.  : N/A

Bridge Id         : 80:02:00:25:ba:04:37:10 Bridge Instance Id: 2
Bridge Priority    : 32768         Tx Hold Count     : 6
Topology Change   : Inactive      Bridge Hello Time  : 2
Last Top. Change  : 1d 00:38:51   Bridge Max Age    : 20
Top. Change Count : 1             Bridge Fwd Delay   : 15

Root Bridge       : 80:01:00:25:ba:02:de:90
Primary Bridge    : N/A

Root Path Cost    : 10             Root Forward Delay: 15
Rcvd Hello Time   : 2             Root Max Age      : 20
Root Priority      : 32769         Root Port         : 2050
-----
Forwarding Database specifics
-----
Service Id        : 10           Mac Move          : Disabled
Mac Move Rate     : 2           Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250         Total Count       : 2
Learned Count     : 0           Static Count      : 2
Remote Age        : 900         Local Age         : 300
High Watermark    : 95%        Low Watermark     : 90%
Mac Learning      : Enabled     Discard Unknown   : Disabled

```

Mac Aging : Enabled Relearn Only : False

Service Endpoints

Endpoint name : e1
Description : (Not Specified)
Revert time : 0
Act Hold Delay : 0
Ignore Standby Signaling : false
Suppress Standby Signaling : false
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 : 07/22/2011 11:04:50

Members

No members found.
=====

Endpoint name : e2
Description : (Not Specified)
Revert time : 0
Act Hold Delay : 0
Ignore Standby Signaling : false
Suppress Standby Signaling : false
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 : 07/22/2011 11:04:50

Members

No members found.
=====

Table 102: Output Fields: Service ID All

Label	Description
Service Detailed Information	
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.
Service Type	Specifies the type of service.
SDP Id	The SDP identifier.
Description	Generic information about the service.
Customer Id	The customer identifier.

Label	Description
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.
Admin Path MTU	The largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the service.
Oper State	The current status of the service.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.

Label	Description
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far-end field.
Number of SDPs	The total number SDPs applied to this service ID.
Service Access Points	
Service Id	The service identifier.
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Encap	The value of the label used to identify this SAP on the access port.
Admin State	The desired state of the SAP.
Oper State	The operating state of the SAP.
Last Changed	The date and time of the last change.
Admin MTU	The largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The SAP ingress QoS policy ID.
Egress qos-policy	The SAP egress QoS policy ID.
Ingress Filter-Id	The SAP ingress filter policy ID.
Egress Filter-Id	The SAP egress filter policy ID.
Multi Svc Site	Indicates the multi-service site that the SAP is a member.
Ingress sched-policy	Indicates the ingress QoS scheduler for the SAP.

Label	Description
Egress sched-policy	Indicates the egress QoS scheduler for the SAP.
Table-based	Indicates the use of table-based resource classification: Enabled (table-based) or Disabled (CAM-based)
Dscp Class Pol Id	Indicates the DSCP classification policy ID.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
SAP Statistics	
Dropped	The number of packets or octets dropped.
Offered Hi Priority	The number of high priority packets, as determined by the SAP ingress QoS policy.
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Split Horizon Group Specifics	
Split Horizon Group	Displays the name of the split horizon group.
Description	Displays the description of the split horizon group.
Instance Id	Displays the Instance identifier of the split horizon group.
Last Change	Displays the date and time of most recent change to the split horizon group.
Split Horizon Group	Displays the name of the split horizon group the SAP or spoke-SDP is associated.

arp

Syntax

arp [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays the ARP table for the IES instance. The ARP entries for a subscriber interface are displayed uniquely. Each MAC associated with the subscriber interface child group-interfaces are displayed with each subscriber interface ARP entry. They do not reflect actual ARP entries but are displayed along the interfaces ARP entry for easy lookup.

Parameters

ip-address

Displays only ARP entries in the ARP table with the specified IP address.

Default All IP addresses.

mac ieee-address

Displays only ARP entries in the ARP table with the specified 48-bit MAC address. The MAC address can be expressed in the form **aa:bb:cc:dd:ee:ff** or **aa-bb-cc-dd-ee-ff** where **aa**, **bb**, **cc**, **dd**, **ee** and **ff** are hexadecimal numbers.

Default All MAC addresses.

sap sap-id

Displays SAP information for the specified SAP ID. See [Common CLI command descriptions](#) for command syntax.

port-id

Specifies the port ID.

interface

Specifies matching service ARP entries associated with the IP interface.

ip-address

The IP address of the interface for which to display matching ARP entries.

Values 1.0.0.0 to 223.255.255.255

ip-int-name

Specifies the IP interface name for which to display matching ARPs.

Output

The following output is an example of ARP table information, and [Table 103: Output Fields: ARP](#) describes the output fields.

Sample output

```
*A:DUT-B# show service id 100 arp
=====
ARP Table
=====
IP Address      MAC Address      Type      Expiry      Interface      SAP
-----
```

```
192.168.1.2      00:00:01:00:00:01 Other  00h00m00s  HW      1/1/1:10*
195.168.1.1      32:67:01:01:00:03 Other  00h00m00s  to7x    1/1/3:10*
195.168.1.2      32:68:01:01:00:02 Dynamic 03h59m58s  to7x    1/1/3:10*
=====
*A:DUT-B#
```

Table 103: Output Fields: ARP

Label	Description
IP Address	The IP address.
MAC Address	The specified MAC address.
Type	Static — FDB entries created by management. Learned — Dynamic entries created by the learning process. Other — Local entries for the IP interfaces created.
Expiry	The age of the ARP entry.
Interface	The interface applied to the service.
SAP	The SAP ID.

base

Syntax
base

Context
show>service>id

Platforms
Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description
This command displays basic information about this IES service.

Output
The following output is an example of basic IES service information, and [Table 104: Output Fields: Base](#) describes the output fields.

Sample output

```
*A:ALA-A# show service id 100 base
-----
Service Basic Information
```

```

-----
Service Id       : 100                Vpn Id       : 100
Service Type    : IES
Description     : Default Ies description for service id 100
Customer Id     : 1
Last Status Change: 08/29/2006 17:44:28
Last Mgmt Change  : 08/29/2006 17:44:28
Admin State     : Up                  Oper State    : Up
SAP Count       : 2
-----

```

Service Access & Destination Points

```

-----
Identifier      Type      AdmMTU  OprMTU  Adm    Opr
-----
sap:1/1/3      null      1514   1514   Up     Up
sap:1/1/4      null      1514   1514   Up     Up
=====

```

*A:ALA-A#

Table 104: Output Fields: Base

Label	Description
Service Id	The service identifier.
VPN Id	The VPN identifier.
Service Type	The type of service.
Description	Displays generic information about the service.
Customer Id	The customer identifier.
Last Status Change	The date and time of the most recent status change.
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.
SAP Count	The number of SAPs defined on the service.
Identifier	Specifies the service access point (SAP).
Type	The type of SAPs allowed in the service. It also describes the applied processing by the node to the packets received on these SAPs.
AdminMTU	The largest frame size (in octets) that the SAP can handle.
OprMTU	The actual largest service frame size (in octets) that can be transmitted through this port, without requiring the packet to be fragmented.
Admin	The administrative state of the SAP.

Label	Description
Opr	The operating state of the SAP.

interface

Syntax

interface [*ip-address* | *ip-int-name*] [**detail**]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays information for the IP interfaces associated with the IES service. If no optional parameters are specified, a summary of all IP interfaces associated to the service are displayed.

Parameters

ip-address

Specifies the IP address of the interface for which to display information.

Values ipv4-address: a.b.c.d (host bits must be 0)

ip-int-name

Specifies the IP interface name for which to display information.

Values 32 characters maximum

detail

Displays detailed IP interface information.

Default IP interface summary output.

Output

The following outputs are example of service ID interface information, and the associated tables describes the output fields.

- [Sample output — Standard, Table 105: Output Fields: Service ID Interface.](#)
- [Sample output for 7210 SAS-Mxp.](#)

Sample output — Standard

```
A:ALA-49# show service id 88 interface
```

```

Interface Table
=====
Interface-Name      Adm      Opr(v4/v6)  Type      Port/SapId
IP-Address          PfxState
-----
Sector A            Up        Down/Down   IES        1/1/1.2.2
-
test                Up        Down/Down   IES        1/1/2:0
10.1.1.1/31         n/a
10.1.1.1/31         n/a
10.1.2.1/31         n/a
test27              Up        Up/--       IES Sub    subscriber
192.168.10.21/24    n/a
grp-if              Up        Down/--     IES Grp    1/2/2
Interfaces : 4
=====
A:ALA-49#
A:ALA-49# show service id 88 interface
=====
Interface Table
=====
Interface-Name Adm Opr(v4/v6) Type Port/SapId
IP-Address PfxState
-----
Sector A Up Down/Down IES 1/1/1.2.2
-
test Up Down/Down IES 1/1/2:0
10.1.1.1/31 n/a
10.1.1.1/31 n/a
10.1.2.1/31 n/a
test27 Up Up/-- IES Sub subscriber
192.168.10.21/24 n/a
grp-if Up Down/-- IES Grp 1/2/2
Interfaces : 4
=====
A:ALA-49#

```

Table 105: Output Fields: Service ID Interface

Label	Description
If Name	The name used to refer to the IES interface.
Type	Specifies the interface type.
IP-Address	Specifies the IP address/IP subnet/broadcast address of the interface.
Adm	The administrative state of the interface.
Opr	The operational state of the interface.
Admin State	The administrative state of the interface.
Oper State	The operational state of the interface.
IP Addr/mask	Specifies the IP address/IP subnet/broadcast address of the interface.

Label	Description
If Index	The index corresponding to this IES interface. The primary index is 1; all IES interfaces are defined in the base virtual router context.
If Type	Specifies the interface type.
SAP Id	Specifies the SAP port ID.
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled.
Arp Timeout	Specifies the timeout for an ARP entry learned on the interface.
MAC Address	Specifies the 48-bit IEEE 802.3 MAC address.
ICMP Mask Reply	Specifies whether ICMP mask reply is enabled or disabled.
Cflowd	Specifies whether cflowd collection and analysis on the interface is enabled or disabled.
Redirects	Specifies the rate for ICMP redirect messages.
Unreachables	Specifies the rate for ICMP unreachable messages.
TTL Expired	Specifies the rate for ICMP TTL messages.

Sample output for 7210 SAS-Mxp

This output is an example of IES routed VPLS interface override on the 7210 SAS-Mxp.

```
*A:Dut-A# show service id 2000 interface "iesRvplsIngr" detail
=====
Interface Table
=====
-----
Interface
-----
If Name       : iesRvplsIngr
Admin State   : Up                               Oper (v4/v6)   : Down/Down
Down Reason C: assocObjNotReady
Protocols     : None
IP Addr/mask  : 10.55.55.2/24                     Address Type   : Primary
IGP Inhibit   : Disabled                           Broadcast Address: Host-ones
HoldUp-Time   : 0                                Track Srrp Inst : 0
-----
Details
-----
Description   : (Not Specified)
If Index      : 14                               Virt. If Index : 14
Last Oper Chg: 11/07/2017 04:48:25                Global If Index: 206
Mon Oper Grp  : None
Srrp En Rtng  : Disabled                           Hold time      : N/A
Port Id       : rvpls
TOS Marking   : Untrusted                           If Type        : IES
SNTP B.Cast   : False                               IES ID         : 2000
MAC Address   : c4:08:4a:45:c0:e4                   Mac Accounting : Disabled
Ingress stats: Disabled                             IPv6 DAD       : Enabled
ARP Timeout   : 14400s                               IPv6 Nbr ReachTi*: 30s
```

```

ARP Retry Ti*: 5000ms
ARP Limit : Disabled
ARP Threshold: Disabled
ARP Limit Lo*: Disabled
IP MTU : (default)
IP Oper MTU : 9198
LdpSyncTimer : None
LSR Load Bal*: system
EGR Load Bal*: both
Vas If Type : none
TEID Load Ba*: Disabled
SPI Load Bal*: Disabled
uRPF Chk : disabled
uRPF Ipv6 Chk: disabled
Mpls Rx Pkts : 0
Mpls Tx Pkts : 0

IPv6 stale time : 14400s
IPv6 Nbr Limit : Disabled
IPv6 Nbr Thresho*: Disabled
IPv6 Nbr Log Only: Disabled

Mpls Rx Bytes : 0
Mpls Tx Bytes : 0

DHCP6 Relay Details
Description : (Not Specified)
Admin State : Down
Oper State : Down
If-Id Option : None
Src Addr : Not configured
Python plcy : (Not Specified)

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

Routed VPLS Details
VPLS Name : ingRvpls
Binding Status : Up
Reason : (Not Specified)
Egr Reclass Plcy : 0
Ing Filter : none
Ingr IPv6 Flt : none
EVPN Tunnel : false
Table-based : enabled
Dscp Class Pol Id: 1

Network Domains Associated
default

-----
Admin Groups
-----
No Matching Entries
-----

-----
Srlg Groups
-----
No Matching Entries
-----

Interfaces : 1
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-A#

```

Table 106: Output Fields: ID Interface on 7210 SAS-Mxp

Label	Description
If Name	The name used to refer to the IES interface.
Type	Specifies the interface type.
IP-Address	Specifies the IP address/IP subnet/broadcast address of the interface.
Adm	The administrative state of the interface.
Opr	The operational state of the interface.
Admin State	The administrative state of the interface.
Oper State	The operational state of the interface.
IP Addr/mask	Specifies the IP address/IP subnet/broadcast address of the interface.
If Index	The index corresponding to this IES interface. The primary index is 1; all IES interfaces are defined in the base virtual router context.
If Type	Specifies the interface type.
SAP Id	Specifies the SAP port ID.
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled.
Arp Timeout	Specifies the timeout for an ARP entry learned on the interface.
MAC Address	Specifies the 48-bit IEEE 802.3 MAC address.
ICMP Mask Reply	Specifies whether ICMP mask reply is enabled or disabled.
Cflowd	Specifies whether cflowd collection and analysis on the interface is enabled or disabled.
Redirects	Specifies the rate for ICMP redirect messages.
Unreachables	Specifies the rate for ICMP unreachable messages.
TTL Expired	Specifies the rate for ICMP TTL messages.

8 Virtual Private Routed Network service

This chapter provides information about the Virtual Private Routed Network (VPN) service and implementation notes. VPRN services are supported only in network mode. It is not supported in access-uplink mode.

8.1 VPRN service overview

RFC2547b is an extension to the original RFC 2547, which describes 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 which ensures that routes from different VPNs remain distinct and separate, even if two VPNs have an overlapping address space. The PE routers distribute routes from other CE routers in that VPN to the CE routers in a particular VPN. 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 SR-series, the label distributed with a VPN route depends on the configured label-mode of the VPRN that is originating the route.

Before a customer data packet travels across the service provider's backbone, it is encapsulated with the MPLS label that corresponds, in the customer's VPN, to the route which best matches the packet's destination address. The MPLS packet is further encapsulated with either another MPLS label header, so that it gets tunneled across the backbone to the correct PE router. Each route exchanged by the MP-BGP protocol includes a route distinguisher (RD), which identifies the VPRN association. Therefore the backbone core routers do not need to know the VPN routes. The following figure shows a VPRN network diagram.

Figure 90: Virtual Private Routed Network



OSSG024

**Note:**

VPN services are only supported on 7210 SAS platforms operating in network mode.

8.1.1 Routing prerequisites

RFC2547bis requires the following features:

- multi-protocol extensions
- extended BGP community support
- BGP capability negotiation
- parameters defined in RFC 2918

Tunneling protocol options are as follows:

- Label Distribution Protocol (LDP)
- MPLS RSVP-TE tunnels

8.1.2 BGP support

BGP is used with BGP extensions mentioned in [Routing prerequisites](#) to distribute VPRN routing information across the service provider network.


BGP was initially designed to distribute IPv4 routing information. Therefore, multi-protocol extensions and the use of a VPN-IPv4 address were created to extend BGP ability to carry overlapping routing information. A VPN-IPv4 address is a 12-byte value consisting of the 8-byte route distinguisher (RD) and the 4-byte IPv4 IP address prefix. The RD must be unique within the scope of the VPRN. This allows the IP address prefixes within different VRFs to overlap.

A VPN-IPv6 address is a 24-byte value consisting of the 8-byte RD and 16-byte IPv6 address prefix. Service providers typically assign one or a small number of RDs per VPN service network-wide.

8.1.3 Route distinguishers

The route distinguisher (RD) is an 8-byte value consisting of 2 major fields, the Type field and value field, as shown in the following figure. The type field determines how the value field should be interpreted. The 7210 SAS implementation supports the three (3) type values as defined in the Internet draft.

Figure 91: Route distinguisher



The three Type values are:

- **Type 0: Value Field - Administrator subfield (2 bytes)**
Assigned number subfield (4 bytes)
The administrator field must contain an ASN (using private AS numbers is discouraged). The Assigned field contains a number assigned by the service provider.
- **Type 1: Value Field - Administrator subfield (4 bytes)**
Assigned number subfield (2 bytes)
The administrator field must contain an IP address (using private IP address space is discouraged). The Assigned field contains a number assigned by the service provider.
- **Type 2: Value Field - Administrator subfield (4 bytes)**
Assigned number subfield (2 bytes)
The administrator field must contain a 4-byte ASN (using private AS numbers is discouraged). The Assigned field contains a number assigned by the service provider.

8.1.3.1 Route reflector

Per RFC2547bis the use of Route Reflectors is supported in the service provider core. Multiple sets of route reflectors can be used for different types of BGP routes, including IPv4 and VPN-IPv6. 7210 can only be used a route reflector client. It cannot be used as a route reflector ("server").

8.1.3.2 Customer Edge to Provider Edge route exchange

Routing information between the Customer Edge (CE) and Provider Edge (PE) can be exchanged by the following methods:

- Static Routes (with both IPv4 and IPv6)
- E-BGP (with both IPv4 and IPv6 VPNs)
- OSPF (v2 IPv4)

Each protocol provides controls to limit the number of routes learned from each CE router.

8.1.3.2.1 Route redistribution

Routing information learned from the CE-to-PE routing protocols and configured static routes should be injected in the associated local VPN routing/forwarding (VRF). In the case of dynamic routing protocols, there may be protocol specific route policies that modify or reject certain routes before they are injected into the local VRF.

Route redistribution from the local VRF to CE-to-PE routing protocols is to be controlled via the route policies in each routing protocol instance, in the same manner that is used by the base router instance.

The advertisement or redistribution of routing information from the local VRF to or from the MP-BGP instance is specified per VRF and is controlled by VRF route target associations or by VRF route policies.

VPN-IP routes imported into a VPRN, have the protocol **type bgp-vpn** to denote that it is an VPRN route. This can be used within the route policy match criteria.

8.1.3.2.2 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 will be removed from the VPRN routing tables dynamically and minimize wasted bandwidth.

[Figure 92: Directly connected IP target](#) and [Figure 93: Multiple hops to IP target](#) show static routes.

Figure 92: Directly connected IP target




Fig_18

Figure 93: Multiple hops to IP target




Fig_19

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 de-activated, the 7210 SAS router will continue to send polls and reactivate any routes that are restored.

8.1.4 Constrained Route Distribution

This section describes constrained route distribution or RT constraint (RTC).

8.1.4.1 Constrained VPN route distribution based on route targets

The RTC is a mechanism 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. After receiving this information, peers restrict the advertised VPN routes to only those requested, which minimizes the control plane load in terms of protocol traffic and possibly routing information base (RIB) memory.

MP-BGP carries the route target membership information, using an address family identifier (AFI) value of 1 and subsequent address family identifier (SAFI) value of 132. For two routers to exchange RT membership network layer reachability information (NLRI), they must advertise the corresponding AFI/SAFI to each other during capability negotiation. MP-BGP allows RT membership NLRI to be propagated, loop-free, within an AS and between ASs using well-known BGP route selection and advertisement rules.

Outbound route filtering (ORF) can also be used for RT-based route filtering, but ORF messages have a limited scope of distribution (to direct peers or neighbors), and, therefore, do not automatically create pruned inter-cluster and inter-AS route distribution trees.

8.1.4.2 Configuring the route target address family

RTC is supported only by the base router BGP instance. When the **family** command at the **bgp**, **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 is mutually exclusive with RTC on a specific BGP session. The CLI will not attempt to block this configuration, but if both capabilities are enabled on a session, the ORF capability is not included in the OPEN message sent to the peer.

8.1.4.3 Originating RT constraint 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 RT extended community imported into a locally-configured Layer-2 VPN or Layer-3 VPN service. These imported route targets are configured in the following contexts:

- **config>service>vpls>bgp**
- **config>service>vprn**
- **config>service>vprn>mvpn**



Note:

See the *7210 SAS-Mxp, R6, R12, S, Sx, T Routing Protocols Guide* for more information about BGP address families that support RTC.

By default, these 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, prefix length, and path attributes. The prefix value is the concatenation of the origin AS (a 4-byte value representing the 2- or 4-octet AS of the originating router, as configured using the **configure 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 any RTC peer using the new **default-route-target group** or **neighbor** CLI command. The default RTC route is a special type of RTC route that has zero prefix length. Sending the default RTC route to a peer conveys a request to receive all VPN routes (regardless of route target extended community) from that peer. The default RTC route is typically advertised by a route reflector to its clients. The advertisement of the default RTC route to a peer does not suppress other, more specific, RTC routes from being sent to that peer.

8.1.4.4 Receiving and re-advertising RT constraint routes

All received RTC routes that are deemed valid are stored in the RIB-IN. An RTC route is considered invalid and treated as withdrawn if any of the following conditions apply:

- The prefix length is 1 to 31.
- The prefix length is 33 to 47.
- The prefix length is 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, standard BGP best path selection procedures are used to determine the best of these routes.

The best RTC route per prefix is re-advertised 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 an RT type with prefix length 48) is never propagated to another peer.
- A PE with only iBGP RTC peers that is neither a route reflector nor an AS boundary router (ASBR) does not re-advertise the best RTC route to any RTC peer because of standard iBGP split horizon rules.
- A route reflector that receives its best RTC route for a prefix from a client peer re-advertises that route (subject to export policies) to all of its client and non-client iBGP peers (including the originator), per standard RR operation. When the route is re-advertised to client peers, the RR 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 re-advertises that route (subject to export policies) to all of its client peers, per standard RR operation. If the RR 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 neither a PE nor a route reflector that receives its best RTC route for a prefix from an iBGP peer re-advertises that route (subject to export policies) to its eBGP peers. It modifies the NEXT_HOP and AS_PATH of the re-advertised route per standard BGP rules. The aggregation of RTC routes is not supported.
- An ASBR that is neither a PE nor a route reflector that receives its best RTC route for a prefix from an eBGP peer re-advertises that route (subject to export policies) to its eBGP and iBGP peers. When re-advertised routes are sent to eBGP peers, the ASBR modifies the NEXT_HOP and AS_PATH per standard BGP rules. The aggregation of RTC routes is not supported.



Note:

These advertisement rules do not handle hierarchical RR topologies properly. This is a limitation of the current RT constraint standard.

8.1.4.5 Using RT constraint routes

In general (ignoring iBGP-to-iBGP rules, add-path, best-external, and so on), the best VPN route for every prefix/NLRI in the RIB is sent to every peer supporting the VPN address family, but export policies may be used to prevent the advertisement of some prefix/NLRIs to specific peers. These export policies may be configured statically or created dynamically based on use of ORF or RTC with a peer. ORF and RTC are mutually exclusive on a session.

When RTC is configured on a session that also supports VPN address families using route targets (vpn-ipv4, vpn-ipv6, and so on), the advertisement of the VPN routes is affected as follows:

- When the session comes up, the advertisement of the VPN routes is delayed briefly 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, the following applies:
 - The set of routes in S1 but not in S2 should be withdrawn immediately (subject to the minimum route advertisement interval (MRAI)).
 - The set of routes in S2 but not in S1 should be advertised immediately (subject to MRAI).

- If a default RTC route is received from a peer P1, the VPN routes that are advertised to P1 are the set 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

**Note:**

This applies whether P1 advertised the best route for the default RTC prefix.

In this context, a default RTC route is any of the following:

- a route with NLRI length = zero
- a route with NLRI value = origin AS and NLRI length = 32
- a route with 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 VPN routes that are advertised to I1 is the set 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

**Note:**

This applies whether 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 VPN routes that are advertised to I1 is the set 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

**Note:**

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 VPN routes that are advertised to E1 is the set 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

**Note:**

This applies only if E1 advertised the best route for A.

8.1.5 BGP fast reroute in a VPRN

BGP fast reroute is a feature that brings together indirection techniques in the forwarding plane and precomputation of BGP backup paths in the control plane to support fast reroute of BGP traffic around unreachable/failed next-hops. In a VPRN context BGP fast reroute is supported using VPN-IPv4 and VPN-IPv6 VPN routes. The supported VPRN scenarios are described in the following table.

Table 107: BGP fast reroute scenarios (VPRN context)

Ingress packet	Primary route	Backup route	Prefix independent convergence
IPv4 (ingress PE)	VPN-IPv4 route with next-hop A resolved by a LDP, RSVP, or BGP tunnel	VPN-IPv4 route with next-hop A resolved by a LDP, RSVP, or BGP tunnel	Yes
IPv6 (ingress PE)	VPN-IPv6 route with next-hop A resolved by a LDP, RSVP, or BGP tunnel	VPN-IPv6 route with next-hop B resolved by a LDP, RSVP, or BGP tunnel	Yes

8.1.5.1 BGP fast reroute in a VPRN configuration

Configuring the **config>service>vprn>enable-bgp-vpn-backup** command causes only imported BGP-VPN routes to be considered when selecting the primary and backup paths.

This command is required to support fast failover of ingress traffic from one remote PE to another remote PE.

**Note:**

7210 SAS platforms do not support BGP backup path commands that are used to enable consideration of multiple paths learned from CE BGP peers when selecting primary and backup paths to reach the CE.

8.2 VPRN features

This section describes VPRN features and special capabilities or considerations as they relate to VPRN services.

8.2.1 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:

- DHCPv4 and DHCPv6 relay
- VRRP
- Secondary IP addresses
- ICMP options

Configuration options found on core IP interfaces that are not supported on VPRN IP interfaces are:

- NTP broadcast receipt

8.2.1.1 DHCP and DHCPv6



Note:

Unless otherwise stated, DHCP is equivalent to "DHCP for IPv4," or DHCPv4.

The DHCP protocol is 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 capability to allocate dynamic network addresses. DHCP 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 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.

Service providers use the DHCP protocol to assign IP addresses and provide other configuration parameters.

IP routers do not forward broadcast or multicast packets, which may suggest that the DHCP client and server must reside on the same IP network segment. However, this configuration is not required because when the 7210 SAS is acting as a DHCP relay agent, it processes these DHCP broadcast or multicast packets and relays them to a preconfigured DHCP server. As a result, DHCP clients and servers do not need to reside on the same IP network segment.

For DHCP relay, the 7210 SAS supports a maximum of eight DHCP servers for each VPRN or IES instance and eight DHCP servers for each node.

For DHCPv6 relay, the 7210 SAS supports a maximum of eight DHCPv6 servers for each VPRN instance and eight DHCPv6 servers for each node.

8.2.1.1.1 DHCP relay and DHCPv6 relay

The 7210 SAS provides DHCP 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.

In the stateful autoconfiguration model, hosts obtain interface addresses and configuration information and parameters from a server. The server maintains a database that tracks which addresses are assigned to which hosts.

**Note:**

The 7210 SAS acts as a relay agent for DHCP and DHCPv6 requests and responses. DHCPv6 functionality is only supported on VPRN access IP interfaces.

8.2.1.1.1.1 DHCP relay

The 7210 SAS provides DHCP relay agent services for DHCP clients.

Service providers use the DHCP protocol to assign IP addresses and provide other configuration parameters. The DHCP protocol requires the client to transmit a request packet with a destination broadcast address of 255.255.255.255, which is processed by the DHCP server.

When the 7210 SAS is acting as a DHCP relay agent, it processes DHCP broadcast packets and relays them to a preconfigured DHCP server, ensuring that DHCP clients and servers do not need to reside on the same network segment.

DHCP offer messages are not dropped if they contain a **giaddr** that does not match the local configured subnets on the DHCP relay interface.

8.2.1.1.1.1 DHCP options

DHCP options are codes that the 7210 SAS inserts in packets being forwarded from a DHCP client to a DHCP server. Some options have additional information stored in suboptions.

The 7210 SAS supports the Relay Agent Information Option 82, as described in RFC 3046. The following suboptions are supported:

- circuit ID
- remote ID
- vendor-specific options

8.2.1.1.1.2 DHCPv6 Relay

**Note:**

DHCPv6 relay is only supported on 7210 SAS-Mxp.

DHCPv6 relay operation is similar to DHCP in that servers send configuration parameters, such as IPv6 network addresses, to IPv6 nodes; however, DHCPv6 relay is not based on the DHCP or BOOTP protocol. DHCPv6 can be used instead of, or in conjunction with, stateless autoconfiguration.

DHCPv6 uses 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 is assigned an IP address. When a client has an address and knows the identity of a server, the client 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 (which addresses 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. Servers that meet the requirements of the client respond with an Advertise message. The client chooses one of the servers and sends a Request message; 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 another way, the client only needs to obtain configuration information. In this case, exchanges are done using a two-message process. The client sends a Request message for only configuration information. A DHCPv6 server that has configuration information for the client sends back a Reply message with the information.

The 7210 SAS supports the DHCPv6 relay agent option in the same way that it supports the DHCP relay agent option: when the 7210 SAS is acting as a DHCPv6 relay agent, it relays messages between clients and servers that are not connected to the same link.

The DHCP relay agent uses one of its interfaces as an IP source address in the DHCP relay-forward message. The DHCPv6 server uses the same source IP address as the destination IP address in the DHCP relay-reply message. There are restrictions for the source IP address used by the DHCP relay agent, which depend on whether the relay agent is a few hops away or is directly connected.

In the case where the relay agent is a few hops away from the DHCPv6 server, the source address used by the relay agent must not fall under the subnet or prefix range configured on the IP interface on which the client is connected. For example, the loopback interface address of the DHCP relay agent can be used instead. To forward the DHCPv6 relay-reply message back to the relay agent, add a static route for the relay agent source IP address.

In the case where the relay agent is directly connected, there are two options. In the first option, the relay agent use the address of the directly connected interface as the relay-forward source address, and no additional configuration is required for the DHCP server to forward the relay-reply message back to the relay-agent. The other option is to use an interface address on the relay agent that does not fall under the subnet or prefix range configured on the IP interface on which the client is connected. Similar to the scenario where the relay-agent is a few hops away, a static route is required to forward the DHCP relay-reply message back to the relay agent.

8.2.1.1.2.1 DHCPv6 options

DHCPv6 options are codes that the 7210 SAS inserts in packets being forwarded from a DHCPv6 client to a DHCPv6 server. DHCPv6 supports interface ID and remote ID options, as described in RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)* and RFC 4649, *DHCPv6 Relay Agent Remote-ID Option*.

8.2.2 SAPs

8.2.2.1 IPv6 support for VPRN IP interfaces



Note:

IPv6 VPRN IP interfaces are supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode.

VPRN IPv6 access interfaces are allowed to be configured to provide IPv6 VPN connectivity to customers.

IPv4 and IPv6 route table lookup entries are shared. Before adding routes for IPv6 destinations, route entries in the routed lookup table needs to be allocated for IPv6 addresses. This can be done using the

CLI **config system resource-profile router max-ipv6-routes** command. This command allocates route entries for /64 IPv6 prefix route lookups. The system does not allocate any IPv6 route entries by default and user needs to allocate some resources before using IPv6. For the command to take effect the node must be rebooted after making the change. For more information, see the following example and see the *7210 SAS-Mxp, R6, R12, S, Sx, T Basic System Configuration Guide*.

A separate route table (or a block in the route table) is used for IPv6 /128-bit prefix route lookup. A limited amount of IPv6 /128 prefixes route lookup entries is supported. The software enables lookups in this table by default (that is no user configuration is required to enable Ipv6 /128-bit route lookup).

In addition, the number IP subnets can be configured by the user using the **configure system resource-profile router max-ip-subnets** command. Suitable defaults are assigned to this parameter. Users can increase the number of subnets if they plan to add more IPv6 addresses for each IPv6 interface.

Following features and restrictions is applicable for IPv6 VPRN IP interfaces:

- PE-CE routing - static routing and EBGp is supported
- A limited amount of IPv6 /128 prefixes route lookup entries is supported on 7210 SAS platforms.
- VRRP for VPRN IPv6 interfaces is not supported.

8.2.2.2 Encapsulations

The following SAP encapsulations are supported on the 7210 SAS VPRN service:

- Ethernet null
- Ethernet dot1q
- QinQ
- LAG

8.2.3 QoS policies

When applied to a VPRN SAP, service ingress QoS policies creates the unicast meter defined in the policy. QoS policies only create the unicast meters defined in the policy if PIM is not configured on the associated IP interface, if PIM is configured, the multipoint meters are applied as well.

For 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-T (network mode) (with VPRN services), access egress policies are available for use on access ports. On 7210 SAS-Mxp Service egress QoS policies are supported.

Both Layer 2 (dot1p only) and Layer 3 criteria can be used in the QoS policies for traffic classification in an VPRN.

8.2.4 Filter policies

Ingress and egress IPv4 and IPv6 filter policies can be applied to VPRN SAPs.

8.2.4.1 CPU QoS for VPRN interfaces

Traffic bound to CPU received on VPRN access interfaces are policed/rate-limited and queued into CPU queues. The software allocates a policer per IP application or a set of IP applications, for rate-limiting CPU bound IP traffic from all VPRN access SAPs. The policers CIR/PIR values are set to appropriate values based on feature scaling and these values are not user configurable. The software allocates a set of queues for CPU bound IP traffic from all VPRN access SAPs. The queues are either shared by a set of IP applications or in some cases allocated to an IP application. The queues are shaped to appropriate rate based on feature scaling. The shaper rate is not user configurable.

**Note:**

- The instance of queues and policers used for traffic received on network port IP interfaces is different for traffic received from access port IP interfaces. Additionally, the network CPU queues are accorded higher priority than the access CPU queues. This is done to provide better security and mitigate the risk of access traffic affecting network traffic.
- The 7210 SAS-Mxp allows the user to configure the IP differentiated services code point (DSCP) value for self-generated traffic. On the 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE, IP DSCP marking of self-generated traffic is not user-configurable and is assigned by software.

8.2.5 CE to PE routing protocols

The 7210 SAS VPRN supports the following CE to PE routing protocols:

- eBGP (for both IPv4 and IPv6)
- static with both IPv4 and IPv6)
- OSPF v2 (IPv4)

8.2.5.1 PE to PE tunneling mechanisms

The 7210 SAS supports multiple mechanisms to provide transport tunnels for the forwarding of traffic between PE routers within the 2547bis network.

The 7210 SAS VPRN implementation supports the use of:

- RSVP-TE protocol to create tunnel LSP's between PE routers
- LDP protocol to create tunnel LSP's between PE routers

These transport tunnel mechanisms provide the flexibility to use dynamically created LSPs, where the "autobind" feature is used to automatically bind service tunnels, and there is the ability to provide specific 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 or the ability to configure sets of LSPs with bandwidth reservations for specific customers, as is available with explicit SDPs for the service.

8.2.5.2 Per-VRF route limiting

The 7210 SAS allows setting the maximum number of routes that can be accepted in the VRF for a VPRN service. There are options to specify a percentage threshold at which to generate an event to indicate that the VRF table is nearly full, and an option to disable additional route learning when the VRF is full or only generate an event.

8.2.6 Exporting MP-BGP VPN routes

To reduce the number of MP-BGP VPN tunnels in a group of IP/MPLS PE routers that are part of the same L3 VPN instance, a hierarchy can be established by reexporting the VPN IP routes on a PE aggregation router (which can be an ABR node). In the case of VPRN service labels, reexporting VPN IP routes reduces the required MPLS FIB resources to the scale available on smaller access routers.

Use the **config>service>vprn>allow-export-bgp-vpn** command to configure the feature. This command enables the **vrf-export** and **vrf-target** export functions to include BGP-VPN routes that are installed in the VPRN route table.

When a route is installed in the VPRN route table, the route is exported as a new VPN-IP route to an MP-IBGP peer only; that is, the route is accepted by the VRF export policy but may be rejected by a BGP export policy. Assuming that the export policies have simple accept and reject actions, the new VPN-IP route is the same as the original VPN-IP route, except in the following cases:

- The RD is changed to the value of the advertising VPRN.
- The BGP next-hop is changed to a local address of the PE.
- The label value is changed to the per-VRF label value of the advertising VPRN.

8.2.6.1 Configuration guidelines

The following configuration guidelines apply to this feature:

- You must shut down and restart the VPRN context for any changes to the **allow-export-bgp-vpn** command to take effect.
- You must configure the VPRN service with a loopback IP interface for the command to take effect.
- SAPs cannot be configured in a VPRN service in which the **allow-export-bgp-vpn** command is enabled.

8.2.7 Spoke-SDPs

Spoke-SDP termination into a Layer 3 service is not supported on 7210 SAS platforms.

8.2.8 Using OSPF in IP-VPNs



Note:

OSPF used as a PE-CE routing protocol is supported only for IPv4 VPNs.

Using OSPF as a CE to PE routing protocol allows OSPF that is currently running as the IGP routing protocol to migrate to an IP-VPN backbone without changing the IGP routing protocol, introducing BGP as the CE-PE or relying on static routes for the distribution of routes into the service providers IP-VPN. The following features are supported:

- Advertisement/redistribution of BGP-VPN routes as summary (type 3) LSAs flooded to CE neighbors of the VPRN OSPF instance. This occurs if the OSPF route type (in the OSPF route type BGP extended community attribute carried with the VPN route) is not external (or NSSA) and the locally configured domain-id matches the domain-id carried in the OSPF domain ID BGP extended community attribute carried with the VPN route.
- OSPF sham links. A sham link is a logical PE-to-PE unnumbered point-to-point interface that essentially rides over the PE-to-PE transport tunnel. A sham link can be associated with any area and can therefore appear as an intra-area link to CE routers attached to different PEs in the VPN.

8.2.9 Service label mode of a VPRN

The 7210 SAS allocates one unique (platform-wide) service label per VRF. All VPN-IP routes exported by the PE from a particular VPRN service with that configuration have the same service label. When the PE receives a terminating MPLS packet, the service label value determines the VRF to which the packet belongs. A lookup of the IP packet DA in the forwarding table of the selected VRF determines the next-hop interface.

8.2.10 Multicast in IP-VPN applications




Note:

Multicast in IP-VPN services using NG-MVPN mechanisms is supported on the 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx 1/10GE operating in the standalone mode, and 7210 SAS-S 1/10GE operating in the standalone-VC mode.

Applications for this feature include enterprise customers implementing a VPRN solution for their WAN networking needs, customer applications including stock-ticker information, financial institutions for stock and other types of trading data, and video delivery systems.

The following figure shows an example of multicast in an IP-VPN application. The provider domain encompasses the core routers (1 through 4) and the edge routers (5 through 10). The various IP-VPN customers each have their own multicast domain, VPN-1 (CE routers 12, 13, and 16) and VPN-2 (CE Routers 11, 14, 15, 17, and 18). In this VPRN example, the VPN-1 data generated by the customer behind router 16 is multicast only by PE router 9 to PE routers 6 and 7 for delivery to CE routers 12 and 13 respectively. Data for VPN-2 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 18, 11, 14, and 17 respectively.

Figure 94: Multicast in IP-VPN applications



The demarcation of these domains is in the PE router (routers 5 through 10). The PE router participates in both the customer multicast domain and the provider multicast domain. The customer CE routers are limited to a multicast adjacency with the multicast instance on the PE created to support that specific customer IP-VPN. This way, 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 router for a specific customer multicast domain becomes adjacent to the CE routers attached to that PE and to all other PE routers that participate in the IP-VPN (or customer) multicast domain. This is achieved by the PE router, which encapsulates the customer multicast control data and multicast streams inside the provider multicast packets. These encapsulated packets are forwarded only to the PE routers that are participating in the same MVPN domain and are part of the same customer VPRN. This prunes the distribution of the multicast control and data traffic to the PE routers that do not participate in the customer multicast domain.

8.2.10.1 Multicast protocols supported in the provider network

An MVPN is defined by two sets of sites: the sender sites set and receiver sites set, with the following properties:

- Hosts within the sender sites set could originate multicast traffic for receivers in the receiver sites set.
- Receivers not in the receiver sites set should not be able to receive this traffic.
- Hosts within the receiver sites set could receive multicast traffic originated by any host in the sender sites set.

- Hosts within the receiver sites set should not be able to receive multicast traffic originated by any host that is not in the sender sites set.

A site could be both in the sender sites set and receiver sites set, which implies that hosts within such a site could both originate and receive multicast traffic. An extreme case is when the sender sites set is the same as the receiver sites set, in which case all sites could originate and receive multicast traffic from each other.

Sites within a specific MVPN can only be within the same organizations, which implies that an MVPN can be an intranet. A specific site may be in more than one MVPN, which implies that MVPNs may overlap. Not all sites of a specific MVPN have to be connected to the same service provider, which implies that an MVPN can span multiple service providers.

Another way to look at MVPN is to say that an MVPN is defined by a set of administrative policies. These policies determine the sender sites set and receiver site set. These policies are established by MVPN customers, but implemented by MVPN service providers using the existing BGP/MPLS VPN mechanisms, such as route targets, with extensions, as necessary.

8.2.10.1.1 MVPN Using BGP control plane

The 7210 SAS supports next generation MVPN with MLDP and RSVP P2MP provider tunnels.

The Nokia implementation supports the following features:

- MVPN is supported on the 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx 1/10GE operating in standalone mode, and 7210 SAS-S 1/10GE operating in standalone-VC mode.
- MVPN membership auto-discovery using BGP is supported.
- PE-PE transmission of C-multicast routing using BGP is supported.
- IPv4 is supported.
- Inter-AS MVPN with option A is supported. This does not require any additional control or data plane implementations.

8.2.10.1.2 MVPN membership auto-discovery using BGP

BGP-based auto-discovery (AD) is performed by using a multicast VPN address family. Any PE router that attaches to an MVPN must issue a BGP update message containing an NLRI in this address family, along with a specific set of attributes.

The PE router uses route targets to specify MVPN route import and export. The route target may be the same as the one used for the corresponding unicast VPN, or it may be different. The PE router can specify separate import route targets for sender sites and receiver sites for a specific MVPN.

The route distinguisher (RD) that is used for the corresponding unicast VPN can also be used for the MVPN.

When BGP AD is enabled, PIM peering on the I-PMSI is disabled, so no PIM hellos are sent on the I-PMSI. C-tree to P-tunnel bindings are also discovered using BGP S-PMSI AD routes, instead of PIM join TLVs.

For example, if AD is disabled, the **c-mcast-signaling bgp** command fails and the following error message displays:

C-multicast signaling in BGP requires auto-discovery to be enabled

AD is enabled by default on the 7210 SAS-T, 7210 SAS-Mxp, 7210 SAS-Sx 10/100GE operating in standalone mode, and 7210 SAS-S 1/10GE operating in standalone-VC mode.

If **c-mcast-signaling** is set to **bgp**, the **no auto-discovery** command fails and the following error message displays:

C-multicast signaling in BGP requires auto-discovery to be enabled

When **c-mcast-signaling** is set to **bgp**, S-PMSI AD is always enabled (configuration is ignored).

8.2.10.2 Provider tunnel support

The following provider tunnels are supported:

- mLDP inclusive provider tunnel
- mLDP selective provider tunnel
- RSVP P2MP LSPs inclusive provider tunnel
- RSVP P2MP LSPs selective provider tunnel

8.2.11 Inter-AS VPRNs

Inter-AS IP-VPN services have been driven by the popularity of IP services and service provider expansion beyond the borders of a single Autonomous System (AS) or the requirement for IP VPN services to cross the AS boundaries of multiple providers. Three options for supporting inter-AS IP-VPNs are described in RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*.




Note:

7210 SAS platforms support only Option-A and Option-C. Option-B is not supported and is included in this description only for completeness.

The first option, referred to as Option-A (shown in the following figure), is considered inherent in any implementation. This method uses a back-to-back connection between separate VPRN instances in each AS. As a result, each VPRN instance views the inter-AS connection as an external interface to a remote VPRN customer site. The back-to-back VRF connections between the ASBR nodes require individual sub-interfaces, one per VRF.


Figure 95: Inter-AS Option-A: VRF-to-VRF model



OSSG255

The second option, referred to as Option-B (shown in the following figure), relies heavily on the AS Boundary Routers (ASBRs) as the interface between the autonomous systems. This approach enhances the scalability of the eBGP VRF-to-VRF solution by eliminating the need for per-VPN configuration on the ASBRs. However it requires that the ASBRs provide a control plan and forwarding plane connection between the autonomous systems. The ASBRs are connected to the PE nodes in its local autonomous system using iBGP either directly or through route reflectors. This means the ASBRs receive all the VPRN information and will forward these VPRN updates, VPN-IPv4, to all its EBGP peers, ASBRs, using itself as the next-hop. It also changes the label associated with the route. This means the ASBRs must maintain an associate mapping of labels received and labels issued for those routes. The peer ASBRs will in turn forward those updates to all local IBGP peers.

Figure 96: Inter-AS Option-B




OSSG256

This form of inter-AS VPRNs does not require instances of the VPRN to be created on the ASBR, as in option-A, as a result there is less management overhead. This is also the most common form of Inter-AS VPRNs used between different service providers as all routes advertised between autonomous systems can be controlled by route policies on the ASBRs.

The third option, referred to as Option-C (shown in the following figure), allows for a higher scale of VPRNs across AS boundaries but also expands the trust model between ASNs. As a result this model is typically used within a single company that may have multiple ASNs for various reasons.

This model differs from Option-B, in that in Option-B all direct knowledge of the remote AS is contained and limited to the ASBR. As a result, in option-B the ASBR performs all necessary mapping functions and the PE routers do not need perform any additional functions then in a non-Inter-AS VPRN.

Figure 97: Option C example



OSSG257

With Option-C, knowledge from the remote AS is distributed throughout the local AS. This distribution allows for higher scalability but also requires all PEs and ASBRs involved in the Inter-AS VPRNs to participate in the exchange of inter-AS routing information.

In Option-C, the ASBRs distribute reachability information for remote PE system IP addresses only. This is done between the ASBRs by exchanging MP-eBGP labeled routes, using RFC 3107, *Carrying Label Information in BGP-4*.

Distribution of VPRN routing information is handled by either direct MP-BGP peering between PEs in the different ASNs or more likely by one or more route reflectors in ASN.

8.2.12 Node management using VPRN with GRT leaking



Note:

Node management using VPRN with Global Route Table (GRT) leaking is only supported on the 7210 SAS-Sx/S 1/10GE (standalone mode).

On the 7210 SAS, management traffic can target the system address in the Base router instance using GRT leaking. Management traffic is received and sent in the VPRN router instance.

Node management is achieved using GRT leaking. In this case, the management traffic uses an IP address in the Base routing context.

8.2.12.1 Management with VPRN using GRT leaking

Traffic leaking to GRT for the 7210 SAS allows service providers to manage the node using a VPRN instance.

For packets entering a local VRF interface, route processing results are derived from the VPRN forwarding table or the GRT. The leaking and preferred lookup results are configured on a per VPRN basis. Only specific configurations are supported (for example, specific routes should only be looked up in the GRT and ignored in the VPRN). To provide operational simplicity and improve streamlining, the CLI configuration is contained within the context of the VPRN service.

Use commands in the **config>service>vprn>grt-lookup** context to configure the traffic leaking to GRT feature.

This is an administrative context and provides the container under which the user can enter all specific commands, except policy definition. Policy definitions remain unchanged but are referenced from this context.

Users must import specific routes from GRT to VPRN to resolve the destination configured in the GRT. This is achieved by configuring route policies using the **config>service>vprn>grt-lookup>import-grt** command. By itself, this only provides part of the solution. Packet forwarding within GRT must route packets back to the correct node and to the specific VPRN from which the destination exists. Destination prefixes must be leaked from the VPRN to the GRT through the use of the route policy configured with the **config>service>vprn>grt-lookup>export-grt** command. Use the commands in the **config>router>policy-options** context to create the route policies.

By default, the number of prefixes leaked from the VPRN to the GRT is limited to five. Use the **config>service>vprn>grt-lookup>export-limit** command to override the default or remove the limit.

The following type of routes are not leaked from VPRN into the GRT:

- aggregate routes
- BGP VPN extranet routes

Management using a VPRN is achieved using IP addresses in the Base routing instance and GRT leaking.

If the destination IP address is not a VRF IP, GRT leaking is enabled, the IP address belongs to a local interface in GRT, and **allow-local-management** is enabled under the **config>service>vprn>grt-lookup>enable-grt** context, the packet is extracted using GRT leaking to the CPM.

8.2.12.2 Route leaking from GRT to VPRN instances

The GRT to VPRN route leaking feature allows routes from the global route table to be exported into specific VPRN instances. The exported routes can be forwarded and readvertised within the VPRN context.

There are two stages of route leaking:

- The first stage requires the configuration of a set of leak-export route policies that identify which GRT routes are subject to being exported into VPRN services. Use the **leak-export** command in the **config>router** context to configure between one and four route policies. The GRT routes must match the policy entries configured with the **config>router>policy-options>policy-statement>entry>action** command value set to **accept**.

In addition, the **leak-export-limit** command is used to specify the maximum number of GRT routes that can be included in the GRT leak pool.

- The second stage requires the configuration of an import GRT policy that specifies which routes within the GRT leak pool are leaked into the associated VPRN instances route table. The **config>service>vprn>grt-lookup>import-grt** command is used and accepts one route policy.

For the GRT route to be leaked into the local VPRN, the route must match a policy entry with the **config>router>policy-options>policy-statement>entry>action** command set to **accept**.

If a GRT route passes both stages, it is added into the VPRN route table that is used for IP forwarding as well as readvertisement within other routing protocols within the VPRN context.

This process can be used to leak IPv4 routes from the GRT into one or more VPRN instances. The GRT route types that can be leaked using this process are:

- RIP, OSPF, and IS-IS routes
- direct routes
- static routes

8.3 Configuring a VPRN service with CLI

This section provides information to configure Virtual Private Routed Network (VPRN) services using the command line interface.

8.3.1 Basic configuration

The following fields require specific input (there are no defaults) to configure a basic VPRN service:

- customer ID (see [Configuring customers accounts](#))
- specify interface parameters

Example: VPRN service configuration

```
*A:ALA-1>config>service>vprn# info
-----
vrf-import "vrfImpPolCust1"
vrf-export "vrfExpPolCust1"

autonomous-system 10000
route-distinguisher 10001:1
auto-bind ldp
vrf-target target:10001:1
interface "to-cel" create
    address 10.1.0.1/24

exit
sap 1/1/10:1 create
    ingress
        qos 100
    exit
    filter ip 10
    exit
exit
```

```
        exit
    exit
    static-route 10.5.0.0/24 next-hop 10.1.1.2
    bgp
        router-id 10.0.0.1
        group "to-cel"
            export "vprnBgpExpPolCust1"
            peer-as 65101
            neighbor 10.1.1.2
        exit
    exit
    exit
    no shutdown
-----
*A:ALA-1>config>service>vprn#
```

8.3.2 Common configuration tasks

About this task

This section provides a brief overview of the tasks that must be performed to configure a VPRN service and provides the CLI commands:

Procedure

- Step 1.** Associate a VPRN service with a customer ID.
- Step 2.** Define an autonomous system (optional).
- Step 3.** Define a route distinguisher (mandatory).
- Step 4.** Define VRF route-target associations or VRF import/export policies.
- Step 5.** Create an interface.
- Step 6.** Define SAP parameters on the interface:
 - Select nodes and ports.
 - Optional - select QoS policies other than the default (configured in `config>qos` context)
 - Optional - select filter policies (configured in `config>filter` context)
 - Optional - select accounting policy (configured in `config>log` context)
- Step 7.** Define BGP parameters (optional):

BGP must be enabled in the `config>router>bgp` context.
- Step 8.** Enable the service.

8.3.3 Configuring VPRN components

8.3.3.1 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.

```
config>service# vprn service-id [customer customer-id]
    route-distinguisher [ip-address:number1 | asn:number2]
    description description-string
    no shutdown
```

Example: VPRN service configuration

```
*A:ALA-1>config>service# info
-----
...
    vprn 1 customer 1 create
        route-distinguisher 10001:0
no shutdown
    exit
...
-----
*A:ALA-1>config>service>vprn#
```

8.3.3.2 Configuring global VPRN parameters

See [VPRN services command reference](#) for CLI syntax to configure VPRN parameters.

Example: VPRN service with configured parameters

```
*A:ALA-1>config>service# info
-----
...
    vprn 1 customer 1 create
        vrf-import "vrfImpPolCust1"
        vrf-export "vrfExpPolCust1"
        autonomous-system 10000
        route-distinguisher 10001:1

        exit
        no shutdown
    exit
...
-----
*A:ALA-1>config>service#
```

8.3.3.2.1 Configuring Router Interfaces

See the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide* for command descriptions and syntax information to configure router interfaces.

Example: Router interface configuration

```

ALA48>config>router# info
#-----
echo "IP Configuration"
#-----
...
    interface "if1"
        address 10.2.2.1/24

    exit
    interface "if2"
        address 10.49.1.46/24
        port 1/1/34
    exit
    interface "if3"
        address 10.11.11.1/24

    exit
...
#-----
ALA48>config>router#

```

8.3.3.2.2 Configuring VPRN protocols - BGP

The autonomous system number and router ID configured in the VPRN context only applies to that particular service.

The minimal parameters that should be configured for a VPRN BGP instance are:

- Specify an autonomous system number for the router. See [Configuring global VPRN parameters](#).
- Specify a router ID - Note that if a new or different router ID value is entered in the BGP context, then the new values takes precedence and overwrites the VPRN-level router ID. See [Configuring global VPRN parameters](#).
- Specify a VPRN BGP peer group.
- Specify a VPRN BGP neighbor with which to peer.
- Specify a VPRN BGP peer-AS that is associated with the preceding 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 ASN. There are no default VPRN BGP groups or neighbors. Each VPRN BGP group and neighbor must be explicitly configured.

All parameters configured for VPRN BGP are applied to the group and are inherited by each peer, but a group parameter can be overridden on a specific basis. VPRN BGP command hierarchy consists of three levels:

- the global level
- the group level
- the neighbor level

For example:

```

config>service>vprn>bgp#           (global level)
group                               (group level)

```

neighbor	(neighbor level)

Note that the local-address must be explicitly configured if two systems have multiple BGP peer sessions between them for the session to be established.

For more information about the BGP protocol, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.

8.3.3.2.1 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.

All parameters configured for a peer group are applied to the group and are inherited by each peer (neighbor), but a group parameter can be overridden on a specific neighbor-level basis.

After a group name is created and options are configured, neighbors can be added within the same autonomous system to create IBGP connections or neighbors in different autonomous systems to create EBGP peers. All parameters configured for the peer group level are applied to each neighbor, but a group parameter can be overridden on a specific neighbor basis.

8.3.3.2.2 VPRN BGP CLI syntax

Use the syntax to configure VPRN BGP parameters ([BGP configuration commands](#)).

Example: VPRN BGP configuration

```
*A:ALA-1>config>service# info
-----
...
    vprn 1 customer 1 create
      vrf-import "vrfImpPolCust1"
      vrf-export "vrfExpPolCust1"

      autonomous-system 10000
      route-distinguisher 10001:1
      auto-bind ldp
      vrf-target target:10001:1
      interface "to-cel" create
        address 10.1.0.1/24
        sap 1/1/10:1 create
        ingress

        qos 100
        exit

        filter ip 6
        exit
      exit
    exit
  static-route 10.5.0.0/24 next-hop 10.1.1.2
  bgp
    router-id 10.0.0.1
    group "to-cel"
      export "vprnBgpExpPolCust1"
      peer-as 65101
      neighbor 10.1.1.2
```

```

        exit
    exit
    exit
    spoke-sdp 2 create
    exit
    no shutdown
    exit
...
-----
*A:ALA-1>config>service#

```

8.3.3.2.3 Configuring a VPRN interface

Interface names associate an IP address to the interface, and then associate the IP interface with a physical port. The logical interface can associate attributes like an IP address, port, Link Aggregation Group (LAG) or the system.

There are no default interfaces.

Note that you can configure a VPRN interface as a loopback interface by issuing the **loopback** command instead of the **sap sap-id** 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.

When using mtrtrace/mstat in a Layer 3 VPN context then the configuration for the VPRN should have a loopback address configured which has the same address as the core instance's system address (BGP next-hop).

See [OSPF configuration commands \(IPv4 only\)](#) for CLI commands and syntax.

Example: VPRN interface configuration output

```

*A:7210 SAS>config>service>vprn>if# info detail
-----
        no description
        no address
        no mac
        arp-timeout 14400
        no allow-directed-broadcasts
        icmp
            mask-reply
            redirects 100 10
            unreachable 100 10
            ttl-expired 100 10
        exit
        no arp-populate
        dhcp
            shutdown
            no description
            proxy-server
                shutdown
                no emulated-server
                no lease-time
            exit
            no option
            no server
            no trusted
            no lease-populate
            no gi-address
            no relay-plain-bootp
            no use-arp
        exit

```

```

no authentication-policy
no ip-mtu
no host-connectivity-verify
no delayed-enable
no bfd
ipcp
    no peer-ip-address
    no dns
exit
no proxy-arp-policy
no local-proxy-arp
no remote-proxy-arp
no shutdown
-----
*A:7210 SAS>config>service>vprn>if#

```

8.3.3.2.4 Configuring a VPRN interface SAP

A SAP is a combination of a port and encapsulation parameters which identifies the service access point on the interface and within the 7210 SAS. Each SAP must be unique within a router. A SAP cannot be defined if the interface **loopback** command is enabled.

When configuring VPRN interface SAP parameters, a default QoS policy is applied to each ingress and egress SAP. Additional QoS policies and scheduler policies must be configured in the **config>qos** context. Filter policies are configured in the **config>filter** context and must be explicitly applied to a SAP. There are no default filter policies.

Example: VPRN interface SAP configuration output

```

*A:ALA-1>config>service# info
-----
...
    vprn 1 customer 1 create
        vrf-import "vrfImpPolCust1"
        vrf-export "vrfExpPolCust1"

        autonomous-system 10000
        route-distinguisher 10001:1
        auto-bind ldp
        vrf-target target:10001:1
        interface "to-cel" create
            address 10.1.0.1/24
            sap 1/1/10:1 create
                ingress

                qos 100
                exit

                filter ip 6
                exit
            exit
        exit
        static-route 10.5.0.0/24 next-hop 10.1.1.2
        spoke-sdp 2 create
        exit
        no shutdown
    exit
...
-----
*A:ALA-1>config>service#

```


8.3.4 Configuring VPRN protocols - OSPF

In a VPRN interface, each VPN routing instance is isolated from any other VPN routing instance, and from the routing used across the backbone. OSPF can be run with any VPRN, independently of the routing protocols used in other VPRNs, or in the backbone itself. For more information about the OSPF protocol, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Routing Protocols Guide*.

```
config>service>vprn>ospf#
```

8.3.4.1 VPRN OSPF CLI syntax

Example

The following is a sample output of the VPRN OSPF previous configuration.

```
A:duta>config>service>vprn# info
-----
router-id 10.10.10.1
autonomous-system 100
route-distinguisher 65510:1
auto-bind ldp
vrf-target target:65520:1
interface "to-ixia-1" create
  address 10.1.1.1/24
  sap 1/1/9:1 create
  exit
exit
interface "to-ixia-2" create
  address 10.1.2.1/24
  sap 1/1/9:12 create
  exit
exit
ospf
  super-backbone
  vpn-domain 0005 0000.0000.0001
  export "from_mbgp_to_ospf"
  area 0.0.0.0
    interface "to-ixia-2"
      mtu 1500
      no shutdown
    exit
    sham-link "to-ixia-1" 10.1.1.1
    exit
    sham-link "to-ixia-1" 10.11.1.1
    exit
  exit
exit
no shutdown
-----
A:duta>config>service>vprn#
```

For more information about the OSPF protocol, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Routing Protocols Guide*.

8.3.5 Service management tasks

This section describes the service management tasks.

8.3.5.1 Modifying VPRN service parameters

Use the CLI syntax to modify VPRN parameters ([VPRN services command reference](#)).

Example

The following is a sample VPRN service configuration output.

```
*A:ALA-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 10.1.1.1/24
            sap 1/1/10:1 create
            exit
        exit
        static-route 10.5.0.0/24 next-hop 10.1.1.2
        bgp
            router-id 10.0.0.1
            group "to-cel"
                export "vprnBgpExpPolCust1"
                peer-as 65101
                neighbor 10.1.1.2
            exit
        exit
    exit
    spoke-sdp 2 create
    exit
exit
...
-----
*A:ALA-1>config>service>vprn#
```

8.3.5.2 Deleting a VPRN service

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

Use the following syntax to delete a VPRN service.

```
config>service#
[no] vprn service-id [customer customer-id]
shutdown
[no] interface ip-int-name
shutdown
```

```
[no]sap sap-id]
[no] bgp
shutdown
[no] spoke-sdp sdp-id
[no] shutdown
```

8.3.5.3 Disabling a VPRN service

Use the following syntax to shut down a VPRN service without deleting any service parameters.

```
config>service#
  vprn service-id [customer customer-id]
  shutdown
```

Example:

```
config>service# vprn 1
config>service>vprn# shutdown
config>service>vprn# exit
```

```
*A:ALA-1>config>service# info
-----
...
  vprn 1 customer 1 create
    shutdown
    vrf-import "vrfImpPolCust1"
    vrf-export "vrfExpPolCust1"

    autonomous-system 10000
    route-distinguisher 10001:1
    auto-bind ldp
    vrf-target target:10001:1
    interface "to-cel" create
      address 10.1.0.1/24
      sap 1/1/10:1 create
      ingress

      qos 100
      exit
      filter ip 6
      exit
    exit
  exit
  static-route 10.5.0.0/24 next-hop 10.1.1.2
  bgp
    router-id 10.0.0.1
    group "to-cel"
      export "vprnBgpExpPolCust1"
      peer-as 65101
      neighbor 10.1.1.2
      exit
    exit
  exit
  spoke-sdp 2 create
  exit
exit
...
-----
```

```
*A:ALA-1>config>service#
```

8.3.5.4 Re-enabling a VPRN service

Use the following syntax to re-enable a VPRN service that was shut down.

```
config>service#  
  vprn service-id [customer customer-id]  
  no shutdown
```

8.4 VPRN services command reference

8.4.1 Command hierarchies

- [VPRN configuration commands](#)
 - [VPRN service configuration commands](#)
 - [Routed VPLS commands](#)
 - [IGMP commands](#)
 - [Multicast VPN commands](#)
 - [Interface commands](#)
 - [Interface VRRP commands](#)
 - [Interface SAP commands](#)
 - [VPRN SAP QoS and filter commands \(for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE\)](#)
 - [VPRN SAP QoS and filter commands \(for 7210 SAS-Mxp\)](#)
 - [BGP configuration commands](#)
 - [Router advertisement commands](#)
 - [OSPF configuration commands \(IPv4 only\)](#)
 - [PIM configuration commands \(for 7210 SAS-T \(network operating mode\) and 7210 SAS-Mxp\)](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

8.4.1.1 VPRN configuration commands

8.4.1.1.1 VPRN service configuration commands

```
config
```

```

- service
- vprn service-id [customer customer-id]
- no vprn service-id
- [no] allow-export-bgp-vpn
- auto-bind-tunnel
- resolution {any | filter | disabled}
- resolution-filter
- [no] ldp
- [no] rsvp
- [no] sr-isis
- [no] sr-ospf
- autonomous-system as-number
- no autonomous-system
- description description-string
- no description
- enable-bgp-vpn-backup [ipv4] [ipv6]
- no enable-bgp-vpn-backup
- grt-lookup
- [no] enable-grt
- [no] allow-local-management
- export-grt policy-name
- no export-grt
- export-limit num-routes
- no export-limit
- export-v6-limit num-routes
- no export-v6-limit
- import-grt plcy-or-long-expr [plcy-or-expr]
- no import-grt
- maximum-ipv6-routes number [log-only] [threshold percent]
- no maximum-ipv6-routes
- maximum-routes number [log-only] [threshold percent]
- no maximum-routes
- route-distinguisher [ip-address:number1 | asn:number2]
- no route-distinguisher
- router-id ip-address
- no router-id
- sgt-qos
- application dscp-app-name dscp {dscp-value | dscp-name}
- application dot1p-app-name dot1p dot1p-priority
- no application {dscp-app-name | dot1p-app-name}
- dscp dscp-name fc fc-name
- no dscp dscp-name
- [no] shutdown
- snmp-community community-name [version SNMP-version]
- no snmp-community community-name
- source-address
- application app [ip-int-name | ip-address]
- no application app

```

8.4.1.1.1 Spoke-SDP commands

```

config
- service
- vprn
- [no] spoke-sdp sdp-id
- description description-string
- no description
- [no] shutdown
- [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
[preference preference] [metric metric] [tag tag] [enable | disable] {next-hop ip-int-

```

```

name | ip-address | ipsec-tunnel ipsec-tunnel-name} [bfd-enable | {cpe-check cpe-ip-address
[interval seconds] [drop-count count] [log]]} {prefix-list prefix-list-name [all | none]]}
- [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
[preference preference] [metric metric] [tag tag] [enable | disable] indirect ip-address [cpe-
check cpe-ip-address [interval seconds][drop-count count] [log]] {prefix-list prefix-list-name
[all | none]]}
- [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
[preference preference] [metric metric] [tag tag] [enable | disable] black-hole {prefix-
list prefix-list-name [all | none]]}
- 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-comm | {[export ext-comm][import ext-comm]}}
- no vrf-target
- [no] shutdown

```

8.4.1.1.2 Routed VPLS commands

```

config
- service
- vprn service-id [customer customer-id]
- no vprn service-id
- interface ip-int-name [create]
- no interface ip-int-name
- vpls service-name
- no vpls
- ingress
- [no] enable-table-classification
- routed-override-qos-policy policy-id
- no routed-override-qos-policy
- v4-routed-override-filter ip-filter-id
- no v4-routed-override-filter

```

8.4.1.1.3 IGMP commands

```

config
- service
- vprn service-id [customer customer-id]
- no vprn service-id
- igmp
- [no] interface ip-int-name
- disable-router-alert-check
- no disable-router-alert-check
- import policy-name
- no import
- max-groups value
- no max-groups
- max-sources sources
- no max-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] starg

```

```

- [no] subnet-check
- version version
- no version
- [no] query-interval
- query-interval seconds
- [no] query-last-member-interval
- query-last-member-interval seconds
- [no] query-response-interval
- query-response-interval seconds
- [no] robust-count
- robust-count robust-count
- [no] shutdown
- ssm-translate
  - [no] grp-range start end
  - [no] source ip-address

```

8.4.1.1.4 Multicast VPN commands



Note:

Multicast VPN commands are only supported on the 7210 SAS-Mxp, 7210 SAS-T (network operating mode), and 7210 SAS-Sx/S 1/10GE (operating in standalone and standalone-VC mode).

```

config
- service
  - vprn service-id [customer customer-id]
  - no vprn service-id
  - mvpn
    - [no] auto-discovery [default]
    - c-mcast-signaling {bgp}
    - no c-mcast-signaling
    - [no] intersite-shared
    - mdt-type {sender-receiver | sender-only | receiver-only}
    - no mdt-type
    - provider-tunnel
      - inclusive
        - bsr {unicast | spmsi}
        - no bsr
        - [no] mldp
          - [no] shutdown
        - [no] rsvp
          - lsp-template lsp-template
          - no lsp-template
          - [no] shutdown
        - [no] wildcard-spmsi
      - selective
        - data-delay-interval value
        - no data-delay-interval
        - data-threshold {c-grp-ip-addr/mask | c-grp-ip-addr netmask}
        - no data-threshold {c-grp-ip-addr/mask | c-grp-ip-addr netmask}
        - maximum-p2mp-spmsi range
        - no maximum-p2mp-spmsi
        - [no] mldp
          - [no] shutdown
        - [no] rsvp
          - lsp-template lsp-template
          - no lsp-template
          - no shutdown
    - umh-selection {highest-ip}
    - no umh-selection

```

```

- vrf-export {unicast | policy-name [policy-name...(up to 15 max)]}
- no vrf-export
- vrf-import {unicast | 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}

```

8.4.1.1.5 Interface commands

```

config
- service
  - vprn service-id [customer customer-id]
  - no vprn service-id
    - [no] interface ip-int-name
      - address {ip-address/mask | ipaddress netmask} [broadcast all-ones | host-
ones]
      - no address
      - [no] allow-directed-broadcasts
      - arp-timeout [seconds]
      - no arp-timeout
      - bfd transmit-interval [receive receive-interval] [multiplier multiplier]
[echo-receive echo-interval]
      - no bfd
      - cflowd-parameters
        - sampling {unicast|multicast} type {interface} [direction {ingress-only}]
        - no sampling {unicast|multicast}
      - delayed-enable seconds
      - no delayed-enable
      - dhcp
        - description description-string
        - no description
        - gi-address ip-address [src-ip-addr]
        - no gi-address
        - [no] option
          - action {replace | drop | keep}
          - no action
          - [no] circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tuple]
          - [no] remote-id [mac | string string]
          - [no] vendor-specific-option
            - [no] client-mac-address
            - [no] sap-id
            - [no] service-id
            - string text
            - no string
            - [no] system-id
          - no relay-plain-bootp
          - relay-plain-bootp
          - no server
          - server server1 [server2 ... (up to 8 max)]
          - [no] shutdown
          - [no] trusted
        - description description-string
        - no description [description-string]
      - icmp
        - [no] mask-reply
        - redirects number seconds
        - no redirects [number seconds]
        - ttl-expired number seconds

```



```

- no ttl-expired [number seconds]
- unreachable number seconds
- no unreachable [number seconds]
- ip-mtu octets
- no ip-mtu
- ipv6
- no ipv6
- [no] address ipv6-address/prefix-length [eui-64] [preferred]
- [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
- icmp6
  - [no] packet-too-big number seconds
  - [no] param-problem number seconds
  - [no] redirects number seconds
  - [no] time-exceeded number seconds
  - [no] unreachable number seconds
- [no] link-local-address ipv6-address [preferred]
- [no] local-proxy-nd
- [no] neighbor ipv6-address mac-address
- [no] proxy-nd-policy policy-name [policy-name ... (up to 5 max)]
- [no] local-proxy-arp
- [no] loopback
- [no] proxy-arp-policy policy-name [policy-name ... (up to 5 max)]
- proxy-arp-policy ieee-address
- no proxy-arp-policy
- [no] remote-proxy-arp
- secondary {ip-address/mask | ip-address netmask} [broadcast {all-ones | host-ones}] [igp-inhibit]
- no secondary {ip-address/mask | ip-address netmask}
- static-arp ip-address ieee-address
- no static-arp ip-address [ieee-address]
- static-arp ieee-addr unnumbered
- no static-arp [ieee-addr] unnumbered
- [no] shutdown
- [no] vrrp virtual-router-id

```

8.4.1.1.6 Interface VRRP commands



Note:

- Interface VRRP commands support IPv6 addresses on 7210 SAS-Mxp only.
- Interface VRRP commands are only supported on 7210 SAS platforms operating in network mode.

```

config
- service
  - vprn service-id [customer customer-id]
  - no vprn service-id

```

```

- interface ip-int-name
-   ipv6
-     vrrp virtual-router-id [owner]
-     no vrrp virtual-router-id
-       [no] backup ip-address
-       init-delay seconds
-       no init-delay
-       [no] master-int-inherit
-       message-interval {[seconds] [milliseconds milliseconds]}
-       no message-interval
-       [no] ping-reply
-       policy vrrp-policy-id
-       no policy
-       [no] preempt
-       priority priority
-       no priority
-       [no] shutdown
-       [no] standby-forwarding
-       [no] telnet-reply
-       [no] traceroute-reply
-     vrrp virtual-router-id [owner]
-     no vrrp virtual-router-id
-       authentication-key {authentication-key | hash-key} [hash | hash2]
-       no authentication-key
-       [no] backup ip-address
-       [no] bfd-enable [service-id] interface interface-name dst-ip ip-address
-       init-delay seconds
-       no init-delay
-       [no] master-int-inherit
-       message-interval {[seconds] [milliseconds milliseconds]}
-       no message-interval
-       [no] ping-reply
-       policy vrrp-policy-id
-       no policy
-       [no] preempt
-       priority priority
-       no priority
-       [no] shutdown
-       [no] ssh-reply
-       [no] standby-forwarding
-       [no] telnet-reply
-       [no] traceroute-reply

```

8.4.1.1.7 Interface SAP commands

```

config
- service
-   vprn service-id [customer customer-id] [create]
-   no vprn service-id
-     [no] interface ip-int-name [create] [tunnel]
-     [no] sap sap-id [create]
-       accounting-policy acct-policy-id
-       no accounting-policy [acct-policy-id]
-       [no] collect-stats
-       description description-string
-       no description [description-string]
-       ingress
-         meter-override
-           meter meter-id [create]
-           no meter meter-id
-             adaptation-rule [pir adaptation-rule] [cir adaptation-rule]

```

```

- cbs size [kbits | bytes | kbytes]
- no cbs
- mbs size [kbits | bytes | kbytes]
- no mbs
- mode mode
- no mode
- rate cir cir-rate [pir pir-rate]
- [no] shutdown
- statistics
  - ingress
    - counter-mode {in-out-profile-count | forward-drop-count}

```

8.4.1.1.8 VPRN SAP QoS and filter commands (for 7210 SAS-T, 7210 SAS-Sx/S 1/10GE, and 7210 SAS-Sx 10/100GE)

```

config
- service
  - vprn service-id [customer customer-id] [create]
  - no vprn service-id
  - [no] interface ip-int-name [create] [tunnel]
    - [no] sap sap-id
      - egress
        - aggregate-meter-rate rate-in-kbps [burst burst-in-kbits] [enable-
stats]
        - no aggregate-meter-rate
        - filter ip ip-filter-id
        - filter ipv6 ipv6 -filter-id
        - no filter [ip ip-filter-id] [ipv6 ipv6 -filter-id]
        - qos policy-id
        - no qos
      - ingress
        - aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
        - no aggregate-meter-rate
        - filter ip ip-filter-id
        - filter [ipv6 ipv6-filter-id]
        - no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
        - qos policy-id
        - no qos

```

8.4.1.1.9 VPRN SAP QoS and filter commands (for 7210 SAS-Mxp)

```

config
- service
  - vprn service-id [customer customer-id] [create]
  - no vprn service-id
  - [no] interface ip-int-name [create] [tunnel]
    - [no] sap sap-id
      - egress
        - agg-rate-limit agg-rate
        - no agg-rate-limit
        - filter [ip ip-filter-id]
        - filter [ipv6 ipv6 -filter-id]
        - filter [mac mac-filter-id]
        - no filter [ip ip-filter-id] [ipv6 ipv6 -filter-id] [mac mac-filter-
id]
        - qos policy-id
        - no qos

```

```

- ingress
- aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
- no aggregate-meter-rate
- filter [ip ip-filter-id]
- filter [ipv6 ipv6-filter-id]
- filter [mac mac-filter-id]
- no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
- qos policy-id [enable-table-classification]
- no qos

```

8.4.1.1.10 BGP configuration commands

```

config
- service
- vprn service-id [customer customer-id]
- no vprn service-id
- [no] bgp
- [no] advertise-inactive
- [no] aggregator-id-zero
- always-compare-med {zero | infinity}
- no always-compare-med
- [no] as-path-ignore
- auth-keychain name
- authentication-key [authentication-key | hash-key] [hash | hash2]
- no authentication-key
- [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
- ignore-nh-metric
- no ignore-nh-metric
- ignore-router-id
- no ignore-router-id
- [no] connect-retry seconds
- [no] damping
- description description-string
- no description
- [no] disable-4byte-asn
- disable-capability-negotiation
- no disable-capability-negotiation
- disable-communities [standard] [extended]
- no disable-communities
- [no] disable-fast-external-failover
- [no] enable-peer-tracking
- export policy-name [policy-name...(up to 5 max)]
- no export
- family [ipv4] [ipv6]
- no family
- hold-time seconds [strict]
- no hold-time
- import policy-name [policy-name...(up to 5 max)]
- no import
- keepalive seconds
- no keepalive
- local-preference ip-address
- no local-preference
- local-as
- local-as as-number [private]

```

```

- no local-as
- local-preference local-preference
- no local-preference
- loop-detect {drop-peer | discard-route | ignore-loop | off}
- no loop-detect
- med-out {number | igp-cost}
- no med-out
- min-as-origination seconds
- no min-as-origination
- min-route-advertisement seconds
- no min-route-advertisement
- multihop ttl-value
- no multihop
- next-hop-self
- no next-hop-self
- preference preference
- no preference
- peer-as as number
- no peer-as
- [no] path-mtu-discovery
- [no] rapid-withdrawal
- [no] remove-private
- router-id ip-address
- no router-id
- [no] shutdown
- [no] group name [dynamic-peer]
  - [no] advertise-inactive
  - [no] aggregator-id-zero
  - [no] as-override
  - auth-keychain name
  - authentication-key [authentication-key | hash-key] [hash | hash2]
  - no authentication-key
  - connect-retry seconds
  - no connect-retry
  - [no] damping
  - description description-string
  - no description
  - [no] disable-4byte-asn
  - disable-communities [standard] [extended]
  - no disable-communities
  - [no] disable-fast-external-failover
  - [no] enable-peer-tracking
  - export policy-name [policy-name...(up to 5 max)]
  - no export
  - family [ipv4] [ipv6]
  - no family
  - hold-time seconds [strict]
  - no hold-time
  - import policy-name [policy-name...(up to 5 max)]
  - no import
  - keepalive seconds
  - no keepalive
  - local-address ip-address
  - no local-address
  - local-as as-number [private]
  - no local-as
  - local-preference local-preference
  - no local-preference
  - loop-detect {drop-peer | discard-route | ignore-loop | off}
  - no loop-detect
  - med-out {number | igp-cost}
  - no med-out
  - min-as-origination seconds
  - no min-as-origination

```

```

- min-route-advertisement seconds
- no min-route-advertisement
- multihop ttl-value
- no multihop
- [no] next-hop-self
- peer-as as-number
- no peer-as
- preference preference
- no preference
- [no] path-mtu-discovery
- prefix-limit limit [log-only] [threshold percent]
- no prefix-limit
- [no] remove-private
- [no] shutdown
- ttl-security min-ttl-value
- no ttl-security
- type {internal | external}
- no type
- [no] neighbor ip-address
  - [no] advertise-inactive
  - [no] aggregator-id-zero
  - [no] as-override
  - auth-keychain name
  - authentication-key [authentication-key | hash-key] [hash | hash2]
  - no authentication-key
  - connect-retry seconds
  - no connect-retry
  - [no] damping
  - description description-string
  - no description
  - [no] disable-4byte-asn
  - disable-communities [standard] [extended]
  - no disable-communities
  - [no] disable-fast-external-failover
  - [no] enable-peer-tracking
  - export policy-name [policy-name...(up to 5 max)]
  - no export
  - family [ipv4] [ipv6]
  - no family
  - hold-time seconds [strict]
  - no hold-time
  - import policy-name [policy-name...(up to 5 max)]
  - no import
  - keepalive seconds
  - no keepalive
  - local-address ip-address
  - no local-address
  - local-as as-number [private]
  - no local-as
  - local-preference local-preference
  - no local-preference
  - loop-detect {drop-peer | discard-route | ignore-loop | off}
  - no loop-detect
  - med-out {number | igp-cost}
  - no med-out
  - min-as-origination seconds
  - no min-as-origination
  - min-route-advertisement seconds
  - no min-route-advertisement
  - multihop ttl-value
  - no multihop
  - [no] next-hop-self
  - peer-as as-number
  - no peer-as

```

```

- preference preference
- no preference
- [no] path-mtu-discovery
- prefix-limit limit [log-only] [threshold percent]
- no prefix-limit
- [no] remove-private
- [no] shutdown
- ttl-security min-ttl-value
- no ttl-security
- type {internal | external}
- no type

```

8.4.1.1.11 Router advertisement commands

```

config
- service
  - vprn service-id [customer customer-id]
  - 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
          - [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

```

8.4.1.1.12 OSPF configuration commands (IPv4 only)



Note:

OSPF configuration commands only support IPv4 addresses.

```

config
- service
  - vprn service-id [customer customer-id]
  - no vprn service-id
    - [no] ospf
      - [no] area area-id
        - area-range ip-prefix/mask [advertise | not-advertise]

```

```

- no area-range ip-prefix/mask
- [no] blackhole-aggregate
- [no] interface ip-int-name [secondary]
  - [no] advertise-subnet
  - authentication-key [authentication-key | hash-key] [hash | hash2]
  - no authentication-key
  - authentication-type {password | message-digest}
  - no authentication-type
  - bfd-enable [remain-down-on-failure]
  - no bfd-enable
  - dead-interval seconds
  - no dead-interval
  - hello-interval seconds
  - no hello-interval
  - interface-type {broadcast | point-to-point}
  - no interface-type
  - message-digest-key key-id md5 [key | hash-key] [hash | hash2]
  - no message-digest-key key-id
  - metric metric
  - no metric
  - mtu bytes
  - no mtu
  - [no] passive
  - priority number
  - no priority
  - retransmit-interval seconds
  - no retransmit-interval
  - [no] shutdown
  - transit-delay seconds
  - no transit-delay
- [no] nssa
  - area-range ip-prefix/mask [advertise | not-advertise]
  - no area-range ip-prefix/mask
  - originate-default-route [type-7]
  - no originate-default-route
  - [no] redistribute-external
  - [no] summaries
- [no] sham-link ip-int-name ip-address
  - authentication-key [authentication-key | hash-key] [hash | hash2]
  - no authentication-key
  - authentication-type {password | message-digest}
  - no authentication-type
  - dead-interval seconds
  - no dead-interval
  - hello-interval seconds
  - no hello-interval
  - message-digest-key key-id md5 [key | hash-key] [hash | hash2]
  - no message-digest-key key-id
  - metric metric
  - no metric
  - retransmit-interval seconds
  - no retransmit-interval
  - [no] shutdown
  - transit-delay seconds
  - no transit-delay
- [no] stub
  - default-metric metric
  - no default-metric
  - [no] summaries
- [no] virtual-link router-id transit-area area-id
  - authentication-key [authentication-key | hash-key] [hash | hash2]
  - no authentication-key
  - authentication-type {password | message-digest}
  - no authentication-type

```



```

- dead-interval seconds
- no dead-interval
- hello-interval seconds
- no hello-interval
- message-digest-key key-id md5 [key | hash-key] [hash | hash2]
- no message-digest-key key-id
- retransmit-interval seconds
- no retransmit-interval
- [no] shutdown
- transit-delay seconds
- no transit-delay
- [no] compatible-rfc1583
- export policy-name [policy-name...(up to 5 max)]
- no export
- external-db-overflow limit seconds
- no external-db-overflow
- external-preference preference
- no external-preference
- [no] ignore-dn-bit
- import policy-name [policy-name...(up to 5 max)]
- no import policy-name [policy-name...(up to 5 max)]
- overload [timeout seconds]
- no overload
- [no] overload-include-stub
- overload-on-boot [timeout seconds]
- no overload-on-boot
- preference preference
- no preference
- reference-bandwidth bandwidth-in-kbps
- no reference-bandwidth
- router-id ip-address
- no router-id
- [no] shutdown
- [no] super-backbone
- [no] suppress-dn-bit
- timers
  - [no] lsa-arrival lsa-arrival-time
  - [no] lsa-generate max-lsa-wait [lsa-initial-wait [lsa-second-wait]]
  - [no] spf-wait max-spf-wait [spf-initial-wait [spf-second-wait]]
- vpn-domain id {0005 | 0105 | 0205 | 8005}
- no vpn-domain
- vpn-tag vpn-tag
- no vpn-tag

```

8.4.1.1.13 PIM configuration commands (for 7210 SAS-T (network operating mode) and 7210 SAS-Mxp)

```

config
- service
  - vpn
    - [no] pim
      - 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] bfd-enable [ipv4]
        - [no] bsm-check-rtr-alert
        - hello-interval hello-interval
        - no hello-interval

```

```

- hello-multiplier deci-units
- no hello-multiplier
- [no] improved-assert
- instant-prune-echo
- no instant-prune-echo
- max-groups value
- no max-groups
- multicast-senders {auto | always | never}
- no multicast-senders
- priority dr-priority
- no priority
- [no] shutdown
- sticky-dr [priority dr-priority]
- no sticky-dr
- three-way-hello
- no three-way-hello
- [no] tracking-support
- [no] non-dr-attract-traffic
- rp
  - [no] anycast rp-ip-address
    - [no] rp-set-peer ip-address
  - 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} spt-threshold
  - ssm-assert-compatible-mode [enable | disable]
  - ssm-default-range-disable ipv4
  - [no] ssm-groups
    - [no] group-range {grp-ip-address/mask | grp-ip-address netmask}

```

8.4.1.2 Show commands

```

show
- service
  - egress-label start-label [end-label]

```

```

- ingress-label start-label [[end-label]
- id service-id
  - all
  - base
  - dhcp
    - statistics [sap sap-id] [interface interface-name]
    - summary [interface interface-name | saps]
  - sap [sap-id [detail]]
  - sdp [sdp-id | far-end ip-address] [detail]
- labels
- sap-using [sap sap-id]
- sap-using interface [ip-address | ip-int-name]
- sap-using [ingress | egress] filter filter-id
- sap-using [ingress | qos-policy qos-policy-id]
- sdp-using [sdp-id | far-end ip-address] [detail | keep-alive-history]
- sdp-using [sdp-id[:vc-id]]
- service-using [vprn] [sdp sdp-id] [customer customer-id]

```

show

```

- router [vprn-service-id]
  - aggregate [family] [active]
  - arp [ip-int-name | ip-address[/mask] | macieee-mac address | summary] [local |
dynamic | static | managed]
  - bgp
    - auth-keychain [keychain]
    - damping [ip-prefix[/prefix-length]] [decayed | history | suppressed] [detail]
[ipv4]
  - damping [ip-prefix[/prefix-length]] [decayed | history | suppressed] [detail]
vpn-ipv4
  - group [name] [detail] inter-as-label
  - neighbor [ip-address [detail]]
  - neighbor [as-number [detail]]
  - neighbor [ip-address [[family family] filter1] [filter3]]
  - neighbor [as-number [[family family] filter2]]
  - next-hop [family] [ip-address [detail]]
  - paths
  - routes [family family] [prefix [detail | longer]]
  - routes [family family] [prefix [hunt | brief]]
  - routes [family family] [community comm-id]
  - routes [family family] [aspath-regex reg-ex1]
  - routes [family] [ipv6-prefix[/prefix-length] [detail | longer] | [hunt [brief]]]
  - summary [all]
  - interface [[ip-address | ip-int-name] [detail]] | summary [family family]
[neighbor ip-address]
  - mvpn
  - mvpn-list [type type] [auto-discovery auto-discovery] [signalling signalling]
[group group]
  - route-table [family][ip-address[/prefix-length] [longer | exact]] |
[protocol protocol-name] | [summary]
  - sgt-qos (See Note)
    - application
    - dscp-map
  - static-arp [ip-address | ip-int-name | mac ieee-mac-address]
  - static-route [ip-prefix /mask] | [preference preference] | [next-hop ip-address |
tag tag] [detail]
  - tunnel-table [ip-address[/mask] [protocol protocol | sdp sdp-id]
  - tunnel-table [summary]

```

**Note:**

For descriptions of the **show router sgt-qos** commands, refer to the *7210 SAS-Mxp, R6, R12, S, Sx, T Quality of Service Guide*, "Network QoS Policy Command Reference, Show Commands (for 7210 SAS-Mxp)".

8.4.1.3 Clear commands

```
clear
- router
  - bgp
    - damping [{prefix/mask [neighbor ip-address]} | {group name}]
    - flap-statistics [[ip-prefix/mask] [neighbor ip-address]] | [group group-name] |
      [regex reg-exp] | [policy policy-name]
    - neighbor {ip-address | as as-number | external | all} [soft | soft-inbound |
      statistics]
    - protocol
    - forwarding-table [slot-number]
    - interface [ip-int-name | ip-address] [icmp] [statistics]
clear
- service
  - id service-id
  - spoke-sdp sdp-id:vc-id ingress-vc-label
  - statistics
    - sap sap-id {all | counters | stp}
    - sdp sdp-id keep-alive
  - id service-id
    - counters
    - spoke-sdp sdp-id:vc-id {all | counters | stp}
    - spoke-sdp
```

8.4.1.4 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] sap sap-id
      - event-type {config-change | oper-status-change}
    - [no] sdp sdp-id:vc-id
      - event-type {config-change | oper-status-change}
  - stp
    - [no] all-events
    - [no] bpdu
    - [no] core-connectivity
    - [no] exception
    - [no] fsm-state-changes
    - [no] fsm-timers
    - [no] port-role
    - [no] port-state
    - [no] sap sap-id
    - [no] sdp sdp-id:vc-id
```

8.4.2 Command descriptions

- [VPRN service configuration commands](#)
- [Show commands](#)
- [Clear commands](#)
- [Debug commands](#)

8.4.2.1 VPRN service configuration commands

- [Generic commands](#)
- [Global commands](#)
- [Multicast VPN commands](#)
- [SDP commands](#)
- [Interface commands](#)
- [Router advertisement commands](#)
- [PIM commands](#)
- [Counter mode commands](#)
- [BGP commands](#)
- [OSPF commands](#)

8.4.2.1.1 Generic commands

description

Syntax

description *description-string*

no description

Context

config>service>vprn>bgp

config>service>vprn

config>service>vprn>if

config>service>vprn>if>sap

config>service>vprn>if>ipv6 (only supported on the 7210 SAS-Mxp)

config>service>vprn>if>ipv6>dhcp6-relay (only supported on the 7210 SAS-Mxp)

config>service>vprn>bgp

config>service>vprn>bgp>group

```
config>service>vpn>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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 configuration.

Parameters

description-string

Specifies the description character string. Allowed values are any string up to 80 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax

[no] shutdown

Context

```
config>service>vpn
```

```
config>service>vpn>if
```

```
config>service>vpn>if>sap
```

```
config>service>vpn>if>ipv6 (only supported on the 7210 SAS-Mxp)
```

```
config>service>vpn>if>ipv6>dhcp6-relay (only supported on the 7210 SAS-Mxp)
```

```
config>service>vpn>bgp
```

```
config>service>vpn>bgp>group
```

```
config>service>vpn>bgp>group>neighbor
```

```
config>service>vpn>spoke-sdp
```

```
config>service>vpn>if>vrrp
```

```
config>service>vpn>if>ipv6>vrrp
```

```
config>service>vpn>if>dhcp
```

```
config>service>vpn>if>ipv6>dhcp
```

```
config>service>vpn>if>ipv6>ospf
```

```
config>service>vpn>if>ipv6>pim
```

```
config>service>vpn>if>ipv6>igmp
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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

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

Default administrative states for services and service entities is described as follows in Special Cases.

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

If the AS number was previously changed, the BGP AS number inherits the new value.

Special Cases

Service Admin State

Bindings to an SDP within the service are 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.

VPNR BGP

Disables the BGP instance on the specific IP interface. Routes learned from a neighbor that is shut down are immediately removed from the BGP database and RTM. If BGP is globally shut down, all group and neighbor interfaces are operationally shut down. If a BGP group is shut down, all member neighbor interfaces are shutdown operationally. If a BGP neighbor is shut down, only that neighbor interface is operationally shutdown.

VRRP Protocol Handling

On all 7210 SAS platforms, VRRP is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure service vprn if vrrp** command instantiates the protocol in the **no shutdown** state, and resources are allocated to enable the node to process the protocol.

To deallocate resources, you must issue the **configure service vprn if vrrp shutdown** and **configure service vprn if no vrrp** commands to allow the node to start up correctly after the reboot. It is not sufficient to only issue a **configure service vprn if vrrp shutdown** command.

Resources for VRRP are allocated when the VRRP context is enabled, either in the base routing instance or the VPRN service instance. Resources are deallocated when the configuration of the last VRRP context under either base routing instances or VPRN service is removed or shutdown.

VRRPv3 Protocol Handling

On all 7210 SAS platforms, VRRPv3 is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure service vprn if ipv6 vrrp** command instantiates the protocol in the **no shutdown** state, and resources are allocated to enable the node to process the protocol.

To deallocate resources, you must issue the **configure service vprn if ipv6 vrrp shutdown** and **configure service vprn if ipv6 no vrrp** commands to allow the node to start up correctly after the reboot. It is not sufficient to only issue the **configure service vprn if ipv6 vrrp shutdown** command.

The resources for VRRPv3 are allocated when the VRRPv3 context is enabled either in the base routing instance or the VPRN service instance. Resources are deallocated when the configuration of the last VRRPv3 context under either base routing instances or VPRN service is removed or shutdown.

BGP Protocol Handling

On all 7210 SAS platforms, BGP is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure service vprn bgp** command instantiates the protocol in the **no shutdown** state, and resources are allocated to enable the node to process the protocol.

To deallocate resources, you must issue the **configure service vprn bgp shutdown** and **configure service vprn no bgp** commands to allow the node to start up correctly after the reboot. It is not sufficient to only issue a **configure service vprn bgp shutdown** command.

Resources for BGP are allocated when the BGP context is enabled either in the base routing instance or the VPRN service. Resources are deallocated when the configuration of the last BGP context under either the base routing instance or VPRN service instance is removed and shutdown.

IGMP Protocol Handling

On all 7210 SAS platforms, IGMP is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure service vprn igmp** command instantiates the protocol in the **no shutdown** state, and resources are allocated to enable the node to process the protocol.

To deallocate resources, you must issue the **configure service vprn igmp shutdown** and **configure service vprn no igmp** commands to allow the node to start up correctly after the reboot. It is not sufficient to issue only the **configure service vprn igmp shutdown** command.

Resources for IGMP are allocated when the IGMP protocol is enabled, either in the base routing instance or the VPRN service instance. Resources are deallocated when the configuration of the last IGMP context under either base routing instances or VPRN service is removed or shutdown.

PIM Protocol Handling

On all 7210 SAS platforms, PIM is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure service vprn pim** command instantiates the protocol in the **no shutdown** state and resources are allocated to enable the node to process the protocol.

To deallocate resources, you must issue the **configure service vprn pim shutdown** and **configure service vprn no pim** commands to allow the node to start up correctly after the reboot. It is not sufficient to issue only the **configure service vprn pim shutdown** command.

Resources for PIM are allocated when the PIM protocol is enabled, either in the base routing instance or the VPRN service instance. Resources are deallocated when the configuration of the last PIM context under either base routing instances or VPRN service is removed or shutdown.

OSPFv2 Protocol Handling

On all 7210 SAS platforms, OSPFv2 is created in the **no shutdown** state.

On the 7210 SAS-Mxp, the protocol is handled as follows.

The **configure service vprn ospf** command instantiates the protocol in the **no shutdown** state, and resources are allocated to enable the node to process the protocol.

To deallocate resource, you must issue the **configure service vprn ospf shutdown** and **configure service vprn no ospf** commands to allow the node to start up correctly after the reboot. It is not sufficient to issue only the **configure service vprn ospf shutdown** command.

Resources for OSPF are allocated when the OSPF protocol is enabled, either in the base routing instance or the VPRN service instance. Resources are deallocated when the configuration of the last OSPF context under either base routing instances or VPRN service is removed or shutdown.

DHCP Protocol Handling for 7210 SAS-Mxp

When the **no shutdown** command is issued in the **configure>service>vprn>if>dhcp** context under the first IPv4 interface configured, in either the base routing instance or VPRN service instance, resources are allocated to enable the node to process the protocol.

The resources are deallocated when you issue the **configure service vprn if dhcp shutdown** command for the last IPv4 interface configured, in either the base routing instance or VPRN service instance enabled to use DHCP relay (IPv4).

DHCP6-RELAY Protocol Handling for 7210 SAS-Mxp

When the **no shutdown** command is issued in the **configure>service>vprn>if>ipv6>dhcp6-relay** context under the first IPv6 interface configured in any VPRN service, resources are allocated to enable processing of the protocol by the node.

The resources are deallocated when you issue the **configure service vprn if ipv6 dhcp6-relay shutdown** command for the last IPv6 interface configured in either the base routing instance or VPRN service instance enabled to use DHCP6 relay.

8.4.2.1.2 Global commands

vprn

Syntax

vprn *service-id* [**customer** *customer-id*] [**create**]

no vprn *service-id*

Context

config>service

Platforms

Supported on all 7210 SAS platforms as described in this document

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 the same routing instance used for service network core routing connectivity. VPRN services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.

IP interfaces defined within the context of an VPRN service ID must have a SAP created as the access point to the subscriber network.

When a service is created, the **customer** keyword and *customer-id* must be specified to associate the service with a customer. The *customer-id* must already exist, having been created using the **customer** command in the service context. When a service is created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

When a service is created, the use of the **customer** *customer-id* is optional to navigate into the service configuration context. Attempting to edit a service with the incorrect *customer-id* results in an error.

Multiple VPRN services are created to separate customer-owned IP interfaces. More than one VPRN service can be created for a single customer ID. More than one IP interface can be created within a single VPRN service ID. All IP interfaces created within an VPRN service ID belong to the same customer.

The **no** form of this 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.

Parameters

service-id

Specifies the service ID number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7210 SAS on which this service is defined.

Values *service-id*: 1 to 2147483648
 svc-name: 64 characters maximum

customer customer-id

Specifies an existing 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

allow-export-bgp-vpn

Syntax

[no] **allow-export-bgp-vpn**

Context

config>service>vpn

Platforms

7210 SAS-Mxp

Description

This command causes the **vrf-export** and **vrf-target** commands to include BGP-VPN routes installed in the VPRN route table. These routes are usually not readvertisable as VPN-IP routes because of split horizon.

When a BGP-VPN route is reexported, the route distinguisher and label values are rewritten according to the configuration of the reexporting VPRN.



Note:

- This command requires the VPRN context to be shut down and restarted for changes to take effect.
- This command can be configured only with VPRN loopback interfaces.



Caution:

Ensure that routing updates do not loop back to the source when this command is used, otherwise the routes could become unstable.

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

Default

no allow-export-bgp-vpn

auto-bind-tunnel

Syntax

auto-bind-tunnel

Context

config>service>vpn

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure automatic binding of a VPRN service using tunnels to MP-BGP peers.

The user must configure the **resolution** option to enable auto-bind resolution to tunnels in the TTM. If the **resolution** option is explicitly set to disabled, auto-binding to tunnels is removed.

If the **resolution** is set to **any**, any supported tunnel type in a VPRN context is selected following the TTM preference. If one or more explicit tunnel types are specified using the **resolution-filter** option, only these tunnel types are selected again following the TTM preference.

The user must set the **resolution** command to **filter** to activate the list of tunnel types configured under **resolution-filter**.

When an explicit SDP to a BGP next hop is configured in a VPRN service (using the **configure service vprn spoke-sdp** command), it overrides the **auto-bind-tunnel** selection for that BGP next hop only. There is no support for reverting automatically to the **auto-bind-tunnel** selection if the explicit SDP goes down. The user must delete the explicit spoke-SDP in the VPRN service context to resume using the **auto-bind-tunnel** selection for the BGP next hop.

resolution

Syntax

resolution {**any** | **filter** | **disabled**}

Context

config>service>vprn>auto-bind-tunnel

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the resolution mode in the automatic binding of a VPRN service to tunnels to MP-BGP peers.

Parameters

any

Keyword that enables the binding to any supported tunnel type in the VPRN context following the TTM preference.

filter

Keyword that enables the binding to the subset of tunnel types configured under **resolution-filter**.

disabled

Keyword that 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

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure the subset of tunnel types that can be used in the resolution of VPRN prefixes within the automatic binding of VPRN services to tunnels to MP-BGP peers.

The following tunnel types are supported in a VPRN context: RSVP, LDP, and segment routing (SR). The BGP tunnel type is not explicitly configured and is therefore implicit. It is always preferred over any other tunnel type enabled in the **auto-bind-tunnel** context.

ldp

Syntax

[no] **ldp**

Context

config>service>vprn>auto-bind-tunnel>res-filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the use of LDP tunnel types for the resolution of VPRN prefixes within the automatic binding of VPRN services to tunnels to MP-BGP peers.

When **ldp** is specified, BGP searches for an LDP LSP with a FEC prefix corresponding to the address of the BGP next hop.

The **no** form of this command disables the use of LDP tunnel types for the resolution of VPRN prefixes within the automatic binding of VPRN services to tunnels to MP-BGP peers.

Default

no ldp

rsvp

Syntax

[no] rsvp

Context

config>service>vprn>auto-bind-tunnel>res-filter

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the use of RSVP tunnel types for the resolution of VPRN prefixes within the automatic binding of VPRN services to tunnels to MP-BGP peers.

When **rsvp** is specified, BGP searches for the best metric RSVP LSP to the address of the BGP next hop. This address can correspond to the system interface or to another loopback used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id.

The **no** form of this command disables the use of RSVP tunnel types for the resolution of VPRN prefixes within the automatic binding of VPRN service to tunnels to MP-BGP peers.

Default

no rsvp

sr-isis

Syntax

[no] sr-isis

Context

config>service>vprn>auto-bind-tunnel>res-filter

Platforms

7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

Description

This command configures the use of SR-ISIS tunnel types for the resolution of VPRN prefixes within the automatic binding of VPRN service to tunnels to MP-BGP peers.

When this command is specified, an SR tunnel to the BGP next hop is selected in the TTM from the lowest numbered IS-IS instance.

The **no** form of this command disables the use of SR-ISIS tunnel types for the resolution of VPRN prefixes within the automatic binding of VPRN service to tunnels to MP-BGP peers.

Default

no sr-isis

sr-ospf

Syntax

[no] **sr-ospf**

Context

config>service>vprn>auto-bind-tunnel>res-filter

Platforms

7210 SAS-Mxp, 7210 SAS-Sx/S 1/10GE (standalone mode), and 7210 SAS-Sx 10/100GE (standalone mode)

Description

This command configures the use of SR-OSPF tunnel types for the resolution of VPRN prefixes within the automatic binding of VPRN service to tunnels to MP-BGP peers.

When **sr-ospf** is specified, an SR tunnel to the BGP next hop is selected in the TTM from the lowest numbered OSPF instance.

The **no** form of this command disables the use of SR-OSPF tunnel types for the resolution of VPRN prefixes within the automatic binding of VPRN service to tunnels to MP-BGP peers.

Default

no sr-ospf

autonomous-system

Syntax

autonomous-system *as-number*

no autonomous-system

Context

config>service>vprn

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command defines the autonomous system (AS) to be used by this VPN routing/forwarding (VRF). The **no** form of this command removes the defined AS from this VPRN context.

Default

no autonomous-system

Parameters

as-number

Specifies the AS number for the VPRN service.

Values 1 to 4294967295

enable-bgp-vpn-backup

Syntax

enable-bgp-vpn-backup [ipv4] [ipv6]

no enable-bgp-vpn-backup

Context

config>service>vpn

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables only imported BGP-VPN routes from the remote PE to be considered when selecting the primary and backup paths. This command is required to support fast failover of ingress traffic from one remote PE to another remote PE.



Note:

7210 SAS platforms do not consider multiple paths learned from CE BGP peers when selecting primary and backup paths to reach the CE.

Default

no enable-bgp-vpn-backup

Parameters

ipv4

Keyword that allows BGP-VPN routes to be used as backup paths for IPv4 prefixes.

ipv6

Keyword that allows BGP-VPN routes to be used as backup paths for IPv6 prefixes.

grt-lookup

Syntax

grt-lookup

Context

config>service>vpn

Platforms

7210 SAS-Sx/S 1/10GE (standalone mode)

Description

Commands in this context configure all GRT leaking commands. If all supporting commands in this context are removed, this command is also removed.

enable-grt

Syntax

[no] enable-grt

Context

config>service>vpn>grt-lookup

Platforms

7210 SAS-Sx/S 1/10GE (standalone mode)

Description

Commands in this context configure the lookup of routes leaked into VPRN from the GRT using route policies.

The **no** form of this command disables the lookup of routes leaked.

Default

no enable-grt

allow-local-management

Syntax

[no] allow-local-management

Context

```
config>service>vprn>grt-lookup>enable-grt
```

Platforms

7210 SAS-Sx/S 1/10GE (standalone mode)

Description

This command enables the support of specific management protocols over VPRN interfaces that terminate on Base routing context IPv4 and IPv6 interface addresses, including Base loopback and system addresses. GRT leaking is used to enable access of the Base interface addresses in the VPRN. The supported protocols are Telnet, FTP, SNMP, TACACS+, RADIUS (IPv4 only, not IPv6), SSH (including applications that ride over the standard SSH TCP port 830 such as SCP and SFTP) and NETCONF (configured on port 830).

Ping and traceroute responses from the Base router interfaces are supported but are not configurable.

The **no** form of this command disables management protocols over VPRN interfaces.

Default

no allow-local-management

export-grt

Syntax

export-grt *policy-name*

no export-grt

Context

```
config>service>vprn>grt-lookup
```

Platforms

7210 SAS-Sx/S 1/10GE (standalone mode)

Description

This command uses the route policy to determine which routes are exported from the VRF to the GRT along with all the forwarding information. These entries are marked as BGP-VPN routes in the GRT. For correct routing to occur from the GRT to the VRF, the routes must be in the GRT.

The **no** form of this command disables exporting routes from the VRF to the GRT.

Default

no export-grt

Parameters

policy-name

Specifies the route policy name, up to 32 characters.

export-limit

Syntax

export-limit *num-routes*

no export-limit

Context

config>service>vpn>grt-lookup

Platforms

7210 SAS-Sx/S 1/10GE (standalone mode)

Description

This command limits the total number of routes exported from the VRF to the GRT. Configuring this limit to 0 disables the maximum limit for routes exported from the VRF to the GRT.

The **no** form of this command sets the export limit to the default value.

Default

export-limit 5

Parameters

num-routes

Specifies the maximum number of routes that can be exported. Configuring this value in a range of 1 to 1000 limits the number of routes to the specified value.

Values 0 to 1000

Default 5

export-v6-limit

Syntax

export-v6-limit *num-routes*

no export-v6-limit

Context

config>service>vpn>grt-lookup

Platforms

7210 SAS-Sx/S 1/10GE (standalone mode)

Description

This command limits the total number of IPv6 routes exported from the VPRN to the GRT. Configuring this limit to 0 disables the maximum limit for IPv6 routes exported from the VPRN to the GRT.

The **no** form of this command sets the export limit to the default value.

Default

export-limit 5

Parameters

num-routes

Specifies the maximum number of IPv6 routes that can be exported. Configuring this value in a range of 1 to 1000 limits the number of IPv6 routes to the specified value.

Values 0 to 1000

Default 5

import-grt

Syntax

import-grt *plcy-or-long-expr* [*plcy-or-expr*]

no import-grt

Context

config>service>vprn>grt-lookup

Platforms

7210 SAS-Sx/S 1/10GE (standalone mode)

Description

This command associates policies to control the leaking of GRT routes into the associated VPRN.

The GRT route must first be leaked by a **leak-export** policy defined under the **config>router** context. Then the route must match a route entry in the specified **import-grt** policy with an accept action.

The **no** form of this command removes route leaking policy associations and disables the leaking of GRT routes into the local VPRN.

Default

no import-grt

Parameters

plcy-or-long-expr

Specifies route policy names, up to 64 characters, or a policy logical expression, up to 255 characters.

Values *plcy-or-long-expr*: *policy-name* | *long-expr*
policy-name: up to 64 characters
long-expr: up to 255 characters

plcy-or-expr

Specifies up to four route policy names, up to 64 characters, or a policy logical expression, up to 64 characters.

Values *plcy-or-expr*: *policy-name* | *expr*
policy-name: up to 64 characters
expr: up to 64 characters

vpls**Syntax**

vpls *service-name*

Context

config>service

config>service>ies>if

config>service>vprn>if

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command, in the IP interface context, is used to bind the IP interface to the specified service name.

The system does not attempt to resolve the provided service name until the IP interface is placed into the administratively up state (**no shutdown**). When the IP interface is administratively up, the system scans the available VPLS services that have the **allow-ip-int-binding** flag set for a VPLS service associated with the name. If the service name is bound to the service name when the IP interface is already in the administratively up state, the system immediately attempts to resolve the specific name.

If a VPLS service is found associated with the name and has the **allow-ip-int-binding** flag set, the IP interface is attached to the VPLS service, allowing routing to and from the service virtual ports when the IP interface is operational.

A VPLS service associated with the specified name that does not have the **allow-ip-int-binding** flag set, or a non-VPLS service associated with the name, is ignored and is not attached to the IP interface.

If the service name is applied to a VPLS service after the service name is bound to an IP interface and the VPLS service **allow-ip-int-binding** flag is set at the time the name is applied, the VPLS service is automatically resolved to the IP interface if the interface is administratively up or when the interface is placed in the administratively up state.

If the service name is applied to a VPLS service without the **allow-ip-int-binding** flag set, the system does not attempt to resolve the applied service name to an existing IP interface bound to the name. To rectify

this condition, the flag must first be set, and then the IP interface must enter or reenter the administratively up state.

While the specified service name may be assigned to only one service context in the system, it is possible to bind the same service name to more than one IP interface. If two or more IP interfaces are bound to the same service name, the first IP interface to enter the administratively up state (if currently administratively down) or to reenter the administratively up state (if currently administratively up) when a VPLS service is configured with the name and has the **allow-ip-int-binding** flag set is attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

When an IP interface is attached to a VPLS service, the name associated with the service cannot be removed or changed until the IP interface name binding is removed. Also, the **allow-ip-int-binding** flag cannot be removed until the attached IP interface is unbound from the service name. Unbinding the service name from the IP interface causes the IP interface to detach from the VPLS service context. The IP interface may then be bound to another service name, or a SAP or SDP binding may be created for the interface using the SAP or spoke-SDP commands on the interface.

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 name. This parameter is required when using the IP interface **vpls** command. The specified name is expressed as an ASCII string comprised of up to 32 characters. It does not need to already be associated with a service and the system does not check to ensure that multiple IP interfaces are not bound to the same name.

interface

Syntax

interface *ip-int-name*

no interface *ip-int-name*

Context

config>service>ies

config>service>vprn

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a logical IP routing interface for a VPRN. When created, attributes such as an IP address and Service Access Point (SAP) can be associated with the IP interface.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within VPRN service IDs. The **interface** command can be executed in the context of an IES service ID.

The IP interface created is associated with the service core network routing instance and default routing.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config service vprn interface** (that is, the network core router instance). Interface names must not be in the dotted-decimal notation of an IP address. For example, the name "1.1.1.1" is not allowed but "int-1.1.1.1" is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All VPRN 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 configuration. The interface must be administratively shut down before issuing the **no interface** command.

For IES services, the IP interface must be shut down before the SAP on that interface may be removed.

Parameters

ip-int-name

Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service vprn interface** commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

If *ip-int-name* already exists within the service ID, the context is changed to maintain that IP interface. If *ip-int-name* already exists within another service ID, an error occurs, and the context is not changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

ingress

Syntax

ingress

Context

```
config>service>ies>if>vpls
```

```
config>service>vprn>if>vlps
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context define the VPLS routed *ip-filter-id* optional filter overrides.

enable-table-classification

Syntax

[no] **enable-table-classification**

Context

config>service>vpn>if>vpls>ingress

Platforms

7210 SAS-Mxp.

Description

This command enables and disables the use of IP DSCP table-based classification to assign forwarding class (FC) and profile on a per-interface ingress basis.

The match-criteria configured in the service ingress policy, which require CAM resources, are ignored. Only meters from the service ingress policy are used (and the meters still require CAM resources). If an IP DSCP classification policy is configured in the VPLS SAP ingress policy, it is not used to assign FC and profile.

The **no** form of this command disables table-based classification. When disabled, the IP ingress packets within a VPLS service attached to the IP interface use the SAP ingress QoS policy applied to the virtual port used by the packets, when defined.

Default

no enable-table-classification

routed-override-qos-policy

Syntax

routed-override-qos-policy *policy-id*

no routed-override-qos-policy

Context

config>service>vpn>if>vpls>ingress

Platforms

7210 SAS-Mxp.

Description

This command configures an IP DSCP classification policy that is applied to all ingress packets entering the VPLS service. The DSCP classification policy overrides any existing SAP ingress QoS policy applied to SAPs for packets associated with the routing IP interface. The routed override QoS policy is optional,

and when it is not defined or removed, the IP routed packets use the existing SAP ingress QoS policy configured on the VPLS virtual port.

The **no** form of this command removes the IP DSCP classification policy from the ingress IP interface. When removed, the IP ingress routed packets within a VPLS service attached to the IP interface use the SAP ingress QoS policy applied to the virtual port used by the packets, when defined.

Default

no routed-override-qos-policy

Parameters

policy-id

Specifies the ID for the routed override QoS policy. Allowed values are an integer that corresponds to an IP DSCP classification policy previously created in the **configure>qos>dscp-classification** context.

Values 1 to 65535

v4-routed-override-filter

Syntax

v4-routed-override-filter *ip-filter-id*

no v4-routed-override-filter

Context

config>service>ies>if>vpls>ingress

config>service>vprn>if>vpls>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies an IP filter ID that is applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IP filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional, and when it is not defined or is removed, the IP routed packets use the existing ingress IP filter on the VPLS virtual port.

The **no** form of this command is used to remove the IP routed override filter from the ingress IP interface. When removed, the IP ingress routed packets within a VPLS service attached to the IP interface use the IP ingress filter applied to the packets virtual port, when defined.

Parameters

ip-filter-id

Specifies the ID for the IP filter policy. Allowed values are an integer that corresponds to an IP filter policy previously created in the **configure>filter>ip-filter** context.

Values 1 to 65535

igmp

Syntax

igmp

Context

config>service>vpn

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

Commands in this context configure IGMP parameters.

interface

Syntax

interface *ip-int-name*

no interface

Context

config>service>vpn>igmp

Platforms

7210 SAS-T, 7210 SAS-Mxp

Description

This command configures 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 (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

import

Syntax

import *policy-name*

no import

Context

config>service>vprn>igmp>if

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command imports a policy to filter IGMP packets.

The **no** form of this command removes the policy association from the IGMP instance.

Default

no import

Parameters

policy-name

Specifies the import route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

The specified names must already be defined.

disable-router-alert-check

Syntax

disable-router-alert-check

no disable-router-alert-check

Context

config>service>vprn>igmp>if

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command enables the IGMP router alert check option.

The **no** form of this command disables the router alert check.

max-sources

Syntax

max-sources *sources*

no max-sources

Context

config>service>vprn>igmp>if

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command specifies the maximum number of 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 sources, the sources that are already accepted are not deleted. Only new sources are not allowed.

Parameters

sources

Specifies the maximum number of sources for this interface.

Values 1 to 1000

max-groups

Syntax

max-groups *value*

no max-groups

Context

config>service>vprn>igmp>if

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

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 are not allowed. By default, there is no limit to the number of groups.

Default

0

Parameters**value**

Specifies the maximum number of groups for this interface.

Values 1 to 1000

static**Syntax****static****Context**`config>service>vpn>igmp>if`**Platforms**

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command tests forwarding on an interface without a receiver host. When enabled, data is forwarded to an interface without receiving membership reports from host members.

group**Syntax**`[no] group grp-ip-address`**Context**`config>service>vpn>igmp>if>static`**Platforms**

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command adds a static multicast group either as a (*,G) or one or more (S,G) records. Use IGMP static group memberships to test multicast forwarding without a 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.

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>vpn>igmp>if>static>group

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command specifies the source address of the multicast group, which is an IPv4 unicast address. By specifying the source address, a multicast receiver host signals to the router that the multicast group only receives 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 this command removes the source from the configuration.

Parameters

ip-address

Specifies the IPv4 unicast address.

Values a.b.c.d

starg

Syntax

[no] **starg**

Context

config>service>vpn>igmp>if>static>group

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command adds a static (*,G) entry. This command can be enabled only if no existing source addresses for this group are specified.

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

subnet-check

Syntax

[no] **subnet-check**

Context

config>service>vprn>igmp>if

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command enables 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.

Default

enabled

version

Syntax

version *version*

no version

Context

config>service>vprn>igmp>if

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

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 on that LAN.

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, 3

query-interval

Syntax

query-interval *seconds*

no query-interval

Context

config>service>vpn>igmp

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command configures the frequency that the querier router transmits general host-query messages. The host-query messages solicit group membership information which are sent to the all-systems multicast group address, 224.0.0.1.

Default

125

Parameters

seconds

Specifies the time frequency, in seconds, that the router transmits general host-query messages.

Values 2 to 1024

query-last-member-interval

Syntax

query-last-member-interval *seconds*

Context

config>service>vpn>igmp

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

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 1024

query-response-interval

Syntax

query-response-interval *seconds*

Context

config>service>vpn>igmp

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command configures 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 to wait 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

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command configures the robust count. The *robust-count* parameter 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

ssm-translate

Syntax

igmp

Context

config>service>vprn>igmp

config>service>vprn>igmp>if

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command configures 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 be added only when the **starg** command is not enabled. An error message is generated when attempting to configure the **source** command while **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

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command configures group ranges that are translated to SSM (S,G) entries.

Parameters

start

Specifies an IP address for the start of the group range.

end

Specifies an IP address for the end of the group range. This value should always be greater than or equal to the value of the *start* value.

source

Syntax

[no] **source** *ip-address*

Context

config>service>vprn>igmp>ssm-translate

config>service>vprn>igmp>if>ssm-translate

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp

Description

This command specifies the source IP address for the group range. Whenever a (*,G) report is received in the range specified by **grp-range** *start* and *end* parameters, it is translated to an (S,G) report with the value of this object as the source address.

Parameters

ip-address

Specifies the IP address that is sending data.

maximum-ipv6-routes

Syntax

maximum-ipv6-routes *number* [**log-only**] [**threshold percent**]

no maximum-ipv6-routes

Context

config>service>vprn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the maximum number of remote IPv6 routes that can be held within a VRF context. Local, host, static, and aggregate routes are not counted.

The VPRN service ID must be in a shutdown state before the **maximum-ipv6-routes** command parameters can be modified.

If the **log-only** keyword is not specified, and the **maximum-ipv6-routes** value is set below the existing number of routes in a VRF, the offending RIP peer (if applicable) is brought down (but the VPRN instance remains up). BGP peering remains up, but the exceeding BGP routes are not added to the VRF.

The maximum route threshold can dynamically change to increase the number of supported routes, even when the maximum is already reached. Protocols resubmit the routes that were initially rejected.

The **no** form of this command disables any limit on the number of routes within a VRF context. Issue the **no** form of this command only when the VPRN instance is shut down.

Default

0 or disabled

Parameters

number

Specifies an integer that specifies the maximum 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 is logged. The **log-only** keyword does not disable the learning of new routes.

threshold percent

Specifies the percentage at which a warning log message and SNMP trap should be set. 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 ($(\text{max} + \text{mid}) / 2$).

Values 0 to 100

maximum-routes

Syntax

maximum-routes *number* [**log-only**] [**threshold percent**]

no maximum-routes

Context

config>service>vprn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the maximum number of remote routes that can be held within a 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 below the existing number of routes in a VRF, the offending RIP peer (if applicable) is brought down (but the VPRN instance remains up). BGP peering remains up, but the exceeding BGP routes are not added to the VRF.

The maximum route threshold can dynamically change to increase the number of supported routes even when the maximum has already been reached. Protocols resubmit the routes that were initially rejected.

The **no** form of this command disables any limit on the number of routes within a VRF context. Issue the **no** form of this command only when the VPRN instance is shut down.

Default

0 or disabled

Parameters***number***

Specifies the maximum number of routes to be held in a VRF context.

Values 1 to 2147483647

log-only

Specifies that if the maximum limit is reached, only log the event. The **log-only** keyword does not disable the learning of new routes.

threshold percent

Specifies the percentage at which a warning log message and SNMP trap should be set. 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 routes and the percent value $([\text{max} + \text{mid}] / 2)$.

Values 0 to 100

route-distinguisher

Syntax

route-distinguisher [*ip-address:number* | *asn:number*]

no route-distinguisher

Context

config>service>vpn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command sets the identifier attached to routes the VPN belongs to. Each routing instance must have a unique route distinguisher (within the carrier domain) associated with it. A route distinguisher must be defined for a VPRN to be operationally active.

Default

no route-distinguisher

Parameters

ip-address:number

Specifies the IP address in dotted-decimal notation. The assigned number must not be greater than 65535.

asn:number

Specifies the AS number 2-byte value less than or equal to 65535. The assigned number can be any 32-bit unsigned integer value.

router-id

Syntax

router-id *ip-address*

no router-id

Context

config>service>vprn

config>service>vprn>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 this command removes the router ID definition from the specific VPRN context.

Default

no router-id

Parameters

ip-address

Specifies the IP address in dotted-decimal notation.

service-name

Syntax

service-name *service-name*

no service-name

Context

config>service>vprn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures an optional service name, up to 64 characters, which adds a name identifier to a specific service to use in configuration references and in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the 7210 SAS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a specific service when 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 to 9).

sgt-qos

Syntax

sgt-qos

Context

config>service>vprn

Platforms

7210 SAS-Mxp

Description

Commands in this context configure DSCP/dot1p re-marking for select self-generated traffic.

application

Syntax

application *dscp-app-name* **dscp** {*dscp-value* | *dscp-name*}

application *dot1p-app-name* **dot1p** *dot1p-priority*

no application {*dscp-app-name* | *dot1p-app-name*}

Context

config>service>vprn>sgt-qos

Platforms

7210 SAS-Mxp

Description

This command configures DSCP/dot1p re-marking for self-generated application traffic. When an application is configured using this command, the specified DSCP name/value is used for all packets generated by this application within the router instance it is configured. The instances can be base router, VPRN service, or management.

Using the value configured in this command:

- Sets the DSCP bits in the IP packet
- Maps to the FC
- Based on this FC, the egress QoS policy sets the Ethernet 802.1p and MPLS EXP bits. This includes ARP and IS-IS packets that, due to their nature, do not carry DSCP bits.
- The DSCP value in the egress IP header is as configured in this command.

Only one DSCP name/value can be configured per application. If multiple entries are configured, the subsequent entry overrides the previously configured entry.

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

Parameters

dscp-app-name

Specifies the DSCP application name.

Values bgp, icmp, igmp, ndis, ospf, pim, ssh, telnet, traceroute, vrrp, arp

dscp-value

Specifies a value when this packet egresses the respective egress policy should provide the mapping for the DSCP value to either LSP-EXP bits or IEEE 802.1p (dot1p) bits as appropriate otherwise the default mapping applies.

Values 0 to 63

dscp-name

Specifies the DSCP name.

Values none, be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

dot1p-priority

Specifies the dot1p priority.

Values none, or 0 to 7

dot1p-app-name

Specifies the dot1p application name.

Values arp, isis

dscp

Syntax

dscp *dscp-name* **fc** *fc-name*

no dscp *dscp-name*

Context

config>service>vpn>sgt-qos

Platforms

7210 SAS-Mxp

Description

This command creates a mapping between the DiffServ Code Point (DSCP) of the self-generated traffic and the forwarding class.

Self-generated traffic for configured applications that matches the specified DSCP are assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all 64 DSCPs to a forwarding class.

All DSCP names that define a DSCP value must be explicitly defined.

The **no** form of this command removes the DSCP-to-forwarding class association.

Parameters

dscp-name

Specifies the name of the DSCP to be associated with the forwarding class. A DiffServ code point can only be specified by its name and only an existing DiffServ code point 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 *fc-name*

Specifies the forwarding class name. Applications and protocols that are configured under the **dscp** command use the configured IP DSCP value.

Values be, l2, af, l1, h2, ef, h1, nc

snmp-community

Syntax

snmp-community *community-name* [**version** *SNMP-version*]

no snmp-community [*community-name*]

Context

config>service>vprn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 this command removes the SNMP community name from the specific VPRN context.

Parameters

community-name

Specifies one or more SNMP community names.

version *SNMP-version*

Specifies the SNMP version.

Values **v1, v2c, both**

source-address

Syntax

source-address

Context

config>service>vprn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context specify the source address and application that should be used in all unsolicited packets.

application

Syntax

application *app* [*ip-int-name* | *ip-address*]

no application *app*

Context

config>service>vpn>source-address

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies the source address and application.

Parameters

app

Specifies the application name.

Values telnet, ssh, traceroute, ping

ip-int-name* | *ip-address

Specifies the name of the IP interface or IP address. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

static-route

Syntax

[no] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] {**next-hop** *ip-int-name* | *ip-address* | **ipsec-tunnel** *ipsec-tunnel-name*} [**bfd-enable** | {**cpe-check** *cpe-ip-address* [**interval** *seconds*] [**drop-count** *count*] [**log**]}

[no] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **indirect** *ip-address* [**cpe-check** *cpe-ip-address* [**interval** *seconds*]] [**drop-count** *count*] [**log**]

[no] **static-route** {*ip-prefix/prefix-length* | *ip-prefix netmask*} [**preference** *preference*] [**metric** *metric*] [**tag** *tag*] [**enable** | **disable**] **black-hole**

Context

config>service>vpn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates static route entries within the associated router instance. When configuring a static route, either **next-hop**, **indirect** or **black-hole** must be configured.

The **no** form of this 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 to uniquely identify the static route must be entered.

If a CPE connectivity check target address is already being used as the target address in a different static route, cpe-check parameters must match. If they do not, the new configuration command are rejected.

If a static-route command is issued with no cpe-check target but the destination prefix/netmask and next-hop matches a static route that did have an associated cpe-check, the cpe-check test is removed from the associated static route.

Parameters

ip-prefix

Specifies the destination address of the aggregate route in dotted-decimal notation.

Values	
ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32
ipv6-prefix	x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D
ipv6-prefix-length	0 to 128

netmask

Specifies the subnet mask in dotted-decimal notation.

Values	a.b.c.d (network bits all 1 and host bits all 0)

ip-int-name

Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service ies interface** commands. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed with

ip-address

Specifies the IP address of the IP interface. The *ip-addr* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted-decimal notation.

Values ipv4-address a.b.c.d (host bits must be 0)

enable

Static routes can be administratively enabled or disabled. Use the **enable** parameter to re-enable a disabled static route. To enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route.

The administrative state is maintained in the configuration file.

Default enable

disable

Static routes can be administratively enabled or disabled. Use the **disable** parameter to disable a static route while maintaining the static route in the configuration. To enable a static route, it must be uniquely identified by the IP address, mask, and any other parameter that is required to identify the exact static route.

The administrative state is maintained in the configuration file.

Default enable

interval seconds

Specifies the interval between ICMP pings to the target IP address.

Values 1 to 255 seconds

Default 1 seconds

drop-count count

Specifies the number of consecutive ping-replies that must be missed to declare the CPE down and to de-active the associated static route.

Values Value range: 1 to 255

Default 3

log

This optional parameter enables the ability to log transitions between active and in-active based on the CPE connectivity check. Events should be sent to the system log, syslog and SNMP traps.

next-hop [ip-address | ip-int-name]

Specifies the directly connected next hop IP address used to reach the destination. If the next hop is over an unnumbered interface, the *ip-int-name* of the unnumbered interface (on this node) can be configured.

The **next-hop** keyword and the **indirect** or **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of either the **indirect** or **black-hole** parameters), this static route is replaced with the newly entered command, and unless specified, the respective defaults for preference and metric are applied.

The *ip-addr* configured here can be either on the network side or the access side on this node. This address must be associated with a network directly connected to a network configured on this node.

ipsec-tunnel *ipsec-tunnel-name*

Specifies an IPSec tunnel name up to 32 characters.

indirect *ip-address*

Specifies that the route is indirect and specifies the next hop IP address used to reach the destination.

The configured *ip-addr* is not directly connected to a network configured on this node. The destination can be reachable via multiple paths. The static route remains valid as long as the address configured as the indirect address remains a valid entry in the routing table. Indirect static routes cannot use an ip-prefix/mask to another indirect static route.

The **indirect** keyword and the **next-hop** or **black-hole** keywords are mutually exclusive. If an identical command is entered (with the exception of either the **next-hop** or **black-hole** parameters), this static route is replaced with the newly entered command and unless specified the respective defaults for preference and metric are applied.

The *ip-addr* configured can be either on the network or the access side and is at least one hop away from this node.

black-hole

Specifies a blackhole route, meaning that if the destination address on a packet matches this static route it is silently discarded.

The **black-hole** keyword is mutually exclusive with either the **next-hop** or **indirect** keywords. If an identical command is entered, with exception of either the **next-hop** or **indirect** parameters, the static route is replaced with the new command, and unless specified, the respective defaults for **preference** and **metric** are applied.

preference preference

Specifies the preference of this static route (as opposed to the routes from different sources such as BGP or OSPF), expressed as a decimal integer. When modifying the **preference** value of an existing static route, unless specified, the metric does not change.

If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used. If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of which route to use is determined by the configuration of the ECMP command.

Values 1 to 255

Default 5

metric metric

Specifies the cost metric for the static route, expressed as a decimal integer. This value is used when importing this static route into other protocols such as OSPF. This value is also used to determine the static route to install in the forwarding table: When modifying the metrics of an existing static route, unless specified, the preference does not change.

If there are multiple static routes with the same preference but unequal metrics, the lower cost (metric) route is installed. If there are multiple static routes with equal preference and

metrics, ECMP rules apply. If there are multiple routes with unequal preferences, the lower preference route is installed.

Values 0 to 65535

Default 1

tag

Adds a 32-bit integer tag to the static route. The tag is used in route policies to control distribution of the route into other protocols.

Values 1 to 4294967295

bfd-enable

Associates the state of the static route to a BFD session between the local system and the configured nexthop. This keyword cannot be configured if the nexthop is **indirect** or a **blackhole** keywords are specified.

NOTE: For more information about the protocols and platforms that support BFD, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.

cpe-check target-ip-address

Specifies the IP address of the target CPE device. ICMP pings are sent to this target IP address. This parameter must be configured to enable the CPE connectivity feature for the associated static route. The target-ip-address cannot be in the same subnet as the static route subnet to avoid possible circular references. This option is mutually exclusive with BFD support on a specific static route.

Default no cpe-check enabled

vrf-export

Syntax

vrf-export *policy* [*policy...*]

no vrf-export

Context

config>service>vpn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies the export policies to control routes exported from the local VPN routing/forwarding (VRF) to other VRFs on the same or remote PE routers (via MP-BGP).

The **no** form of this command removes all route policy names from the export list.

Parameters

policy

Specifies the route policy statement name.

vrf-import

Syntax

vrf-import *policy* [*policy*...]

no vrf-import

Context

config>service>vpn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command sets the import policies to control routes imported to the local VPN routing/ forwarding (VRF) from other VRFs on the same or remote PE routers (via MP-BGP). BGP-VPN routes imported with a vrf-import policy 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 this command removes all route policy names from the import list.

Parameters

policy

Specifies the route policy statement name.

vrf-target

Syntax

vrf-target {**ext-community** | **export** *ext-community* | **import** *ext-community*}

no vrf-target

Context

config>service>vpn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 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 this command removes the vrf-target.

Default

no vrf-target

Parameters

ext-comm

Specifies an extended BGP community in the **type:x:y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. **x** and **y** are 16-bit integers.

Values *ip-addr:comm-val | 2byte-asnumber:ext-comm-val | 4byte-asnumber:comm-val*

where:

ip-addr — IP address in the form a.b.c.d.

comm-val — 0 to 65535

2byte-asnumber — 0 to 65535

ext-comm-val — 0 to 4294967295

4byte-asnumber — 0 to 4294967295

import ext-community

Specifies communities allowed to be accepted from remote PE neighbors.

export ext-community

Specifies communities allowed to be sent to remote PE neighbors.

8.4.2.1.3 Multicast VPN commands

```
mvpn
```

Syntax

```
mvpn
```

Context

```
config>service>vpn
```

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

Commands in this context configure MVPN-related parameters for the IP VPN.

auto-discovery

Syntax

[no] auto-discovery [default]

Context

config>service>vprn>mvpn

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command enables MVPN membership auto-discovery through BGP. When auto-discovery is enabled, PIM peering on the inclusive provider tunnel is disabled.

The **no** form of this command disables MVPN membership auto-discovery through BGP.

Default

enabled

Parameters

default

Enables AD route exchange based on format defined in draft-ietf-l3vpn-2547bis-mcast-10.

c-mcast-signaling

Syntax

c-mcast-signaling bgp

no c-mcast-signaling

Context

config>service>vprn>mvpn

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command specifies BGP or PIM, for PE-to-PE signaling of CE multicast states.

Changes may only be made to this command when the mvpn node is shutdown.

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

Default

mcast-signaling bgp

Parameters

bgp

Specifies to use BGP for PE-to-PE signaling of CE multicast states. Auto-discovery must be enabled.

intersite-shared

Syntax

intersite-shared

no intersite-shared

Context

config>service>vprn>mvpn

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command specifies whether to use inter-site shared C-trees.

Default

intersite-shared

mdt-type

Syntax

mdt-type {sender-receiver | sender-only | receiver-only}

no mdt-type

Context

```
config>service>vpn>mvpn
```

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command restricts MVPN instances per PE node to a specific role. By default, the MVPN instance on a specific PE node assumes the role of sender and 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 PE node to a specific role prevents the creation of full mesh of MDT/PMSI across all participating PE nodes in the MVPN instance.

The **auto-rp-discovery** command cannot be enabled together with the **mdt-type sender-only**, **mdt-type receiver-only**, or **wildcard-spmsi** configurations.

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

Default

```
mdt-type sender-receiver
```

Parameters

sender-receiver

Keyword to connect both senders and receivers to the PE node for MVPN.

sender-only

Keyword to connect only senders to the PE node for MVPN.

receiver-only

Keyword to connect only receivers to the PE node for MVPN.

provider-tunnel

Syntax

```
provider-tunnel
```

Context

```
config>service>vpn>mvpn
```

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

Commands in this context configure tunnel parameters for the MVPN.

inclusive

Syntax

inclusive

Context

config>service>vpn>mvpn>pt

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

Commands in this context specify inclusive provider tunnels

bsr

Syntax

bsr {unicast | spmsi}

no bsr

Context

config>service>vpn>mvpn>pt>inclusive

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command configures the type of bootstrap router (BSR) signaling used.

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

Default

no bsr

Parameters

unicast

Keyword to send or forward BSR PDUs using unicast PDUs (default).

spmsi

Keyword to send or forward BSR PDUs using S-PMSI full mesh.

mldp

Syntax

mldp

no mldp

Context

config>service>vprn>mvpn>provider-tunnel>inclusive

config>service>vprn>mvpn>provider-tunnel>selective

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command enables the use of mLDP LSP for the provider tunnel.

Default

no mldp

shutdown

Syntax

shutdown

no shutdown

Context

config>service>vprn>mvpn>provider-tunnel>inclusive>mldp

config>service>vprn>mvpn>provider-tunnel>selective>mldp

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command administratively disables and enables the use of mLDP LSP for the provider tunnel.

Default

no shutdown

rsvp

Syntax

rsvp
no rsvp

Context

config>service>vprn>mvpn>pt>inclusive
config>service>vprn>mvpn>pt>selective

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command enables use of P2MP RSVP as the inclusive or selective provider tunnel.
The **no** form of this command removes the **rsvp** context, including all the statements in the context.

Default

no rsvp

lsp-template

Syntax

lsp-template *lsp-template*
no lsp-template

Context

config>service>vprn>mvpn>pt>inclusive>rsvp
config>service>vprn>mvpn>pt>exclusive>rsvp

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command configures the use of an automatically created P2MP LSP as the inclusive or selective provider tunnel. The P2MP LSP is signaled using the parameters specified in the template, such as bandwidth constraints.

The **no** form of the command removes the configuration.

Default

no lsp-template

Parameters***lsp-template***

Specifies the LSP template name, up to 32 characters.

shutdown**Syntax**

shutdown

no shutdown

Context

config>service>vpn>mvpn>pt>inclusive>rsvp

config>service>vpn>mvpn>pt>selective>rsvp

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command administratively disables the use of RSVP P2MP LSP for the inclusive or selective provider tunnel.

The **no** form of this command administratively enables the use of RSVP P2MP LSP for the provider tunnel.

Default

no shutdown

wildcard-spmsi**Syntax**

wildcard-spmsi

no wildcard-spmsi

Context

config>service>vpn>mvpn>pt>inclusive

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command enables RFC 6625 (C-*, C-*) S-PMSI functionality for NG-MVPN. When enabled, (C-*, C-*) S-PMSI is used instead of I-PMSI for this MVPN. Wildcard S-PMSI uses the I-PMSI LSP template.

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 this command disables the (C-*, C-*) S-PMSI functionality.

Default

no wildcard-spmsi

selective

Syntax

selective

Context

config>service>vprn>mvpn>provider-tunnel

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

Commands in this context specify selective provider tunnel parameters.

data-delay-interval

Syntax

data-delay-interval *value*

no data-delay-interval

Context

config>service>vprn>mvpn>provider-tunnel>selective

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

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 this command reverts the value to the default.

Default

data-delay-interval 3

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*}
no data-threshold {*c-grp-ip-addr/mask* | *c-grp-ip-addr netmask*}

Context

config>service>vprn>mvpn>provider-tunnel>selective

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command provides an option to the user to specify the group range for which a switch from the inclusive provider tunnel to the selective provider tunnel for C-(S,G) must be triggered. This command provides an option to use selective provide tunnel, independent of the multicast data rate (that is, there is no rate-threshold configuration required). For C-(S,G) groups specified with this command, the selective provider tunnel is used.



Note:
For C-(S,G) groups not configured with the **data-threshold** command, the inclusive provider tunnel is used.

Multiple statements are allowed in the configuration to specify multiple group ranges.
The **no** form of this command removes the values from the configuration.

Parameters

c-grp-ip-addr/mask* | *c-grp-ip-addr netmask
Specifies an IPv4 multicast group address and mask or network mask.

Values	<i>c-grp-ip-addr</i>	multicast group address a.b.c.d
	<i>mask</i>	4 to 32

netmask

a.b.c.d (network bits all 1 and host bits all 0)

maximum-p2mp-spmsi

Syntax

[no] maximum-p2mp-spmsi

Context

config>service>vprn>mvpn>provider-tunnel>selective

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command specifies the maximum number of RSVP P2MP or LDP P2MP S-PMSI tunnels for the mVPN. When the limit is reached, no more RSVP P2MP S-PMSI or LDP P2MP S-PMSI is created and traffic over the data-threshold stays on I-PMSI.

Default

maximum-p2mp-spmsi 10

Parameters

number

Specifies the maximum number of RSVP P2MP or LDP P2MP S-PMSI tunnel for the mVPN.

Values 1 to 510

umh-selection

Syntax

umh-selection highest-ip**no umh-selection**

Context

config>service>vprn>mvpn

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command specifies which UMH selection mechanism to use, highest IP address. The **no** form of this command resets it back to default.

Default

umh-selection highest-ip

vrf-export

Syntax

vrf-export {**unicast** | *policy-name* [*policy-name*...(up to 15 max)]}

no vrf-export

Context

config>service>vpn>mvpn

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

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 unicast VRF export policy for the MVPN.

policy

Specifies a route policy name.

vrf-import

Syntax

vrf-import {**unicast** | *policy-name* [*policy-name*...(up to 15 max)]}

no vrf-import

Context

config>service>vpn>mvpn

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

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

Specifies a 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>vpn>mvpn

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

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. vrf-import or vrf-export policies override the vrf-target policy.

The **no** form of this command removes the vrf-target.

Default

no vrf-target

Parameters

unicast

Specifies to use unicast vrf-target ext-community for the multicast VPN.

ext-comm

Specifies an extended BGP community in the **type:x:y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. **x** and **y** are 16-bit integers.

Values *ip-addr:comm-val | 2byte-asnumber:ext-comm-val | 4byte-asnumber:comm-val*

where:

ip-addr — IP address in the form a.b.c.d.

comm-val — 0 to 65535

2byte-asnumber — 1 to 65535

ext-comm-val — 0 to 4294967295

4byte-asnumber — 0 to 4294967295

import ext-community

Specifies communities allowed to be accepted from remote PE neighbors.

export ext-community

Specifies communities allowed to be sent to remote PE neighbors.

export

Syntax

export {unicast | *ext-community*}

Context

config>service>vpn>mvpn>vrf-target

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command specifies communities to be sent to peers.

Parameters**unicast**

Specifies the use of unicast vrf-target ext-community for the multicast VPN.

ext-community

Specifies an extended BGP community in the **type:x:y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. **x** and **y** are 16-bit integers.

Values *ip-addr:comm-val | 2byte-asnumber:ext-comm-val | 4byte-asnumber:comm-val*

where:

ip-addr — IP address in the form a.b.c.d.

comm-val — 0 to 65535

2byte-asnumber — 1 to 65535

ext-comm-val — 0 to 4294967295

4byte-asnumber — 0 to 4294967295

import

Syntax

import {unicast | ext-community}

Context

config>service>vpn>mvpn>vrf-target

Platforms

7210 SAS-Mxp, 7210 SAS-T (network mode), and 7210 SAS-Sx/S 1/10GE (standalone mode and standalone-VC mode)

Description

This command specifies communities to be accepted from peers.

Parameters

unicast

Specifies that a unicast vrf-target ext-community is used for the multicast VPN.

ext-community

Specifies an extended BGP community in the **type:x:y** format. The value **x** can be an integer or IP address. The **type** can be the target or origin. **x** and **y** are 16-bit integers.

Values *ip-addr:comm-val | 2byte-asnumber:ext-comm-val | 4byte-asnumber:comm-val*

where:

ip-addr — IP address in the form a.b.c.d.

comm-val — 0 to 65535

2byte-asnumber — 1 to 65535

ext-comm-val — 0 to 4294967295

4byte-asnumber — 0 to 4294967295

8.4.2.1.4 SDP commands

spoke-sdp

Syntax

[no] **spoke-sdp** *sdp-id*

Context

config>service>vprn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command binds a service to an existing Service Distribution Point (SDP). The SDP defines the transport tunnel to which this VPRN service is bound.

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 exist in the **config>service>sdp** context before it can be associated 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 the specific *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service to allow far-end 7210 SAS 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. When the SDP binding is removed, no packets are forwarded to the far-end router.

Default

no *sdp-id* is bound to a service

Special Cases

VPRN

Several SDPs can be bound to a VPRN service. Each SDP must be destined to a different 7210 SAS router. If two *sdp-id* bindings terminate on the same 7210 SAS, an error occurs and the second SDP binding is rejected.

Parameters

sdp-id

Specifies the SDP identifier.

Values 1 to 17407

vc-id

Specifies the virtual circuit identifier.

Values 1 to 4294967295

8.4.2.1.5 Interface commands

interface

Syntax

interface *ip-int-name*

no interface *ip-int-name*

Context

config>service>vprn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates a logical IP routing interface for a VPRN. When created, attributes like an IP address and SAP can be associated with the IP interface.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within VPRN service IDs. The **interface** command can be executed in the context of an VPRN service ID. The IP interface created is associated with the service core network routing instance and default routing table. The typical use for IP interfaces created in this manner is for subscriber Internet access.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config router interface** and **config service vprn interface**. Interface names must not be in the dotted-decimal notation of an IP address. For example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

The available IP address space for local subnets and routes is controlled with the **config router service-prefix** command. The **service-prefix** command administers the allowed subnets that can be defined on service IP interfaces. It also controls the prefixes that may be learned or statically defined with the service IP interface as the egress interface. This allows segmenting the IP address space into **config router** and **config service** domains.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All VPRN IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The **no** form of this command removes IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For VPRN services, the IP interface must be shutdown before the SAP on that interface may be removed. VPRN services do not have the **shutdown** command in the SAP CLI context. VPRN service SAPs rely on the interface status to enable and disable them.

Parameters

ip-int-name

Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service vprn interface** commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

If *ip-int-name* already exists within the service ID, the context is changed to maintain that IP interface. If *ip-int-name* already exists within another service ID or is an IP interface defined within the **config router** commands, an error occurs and the context is not changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

address

Syntax

address {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**]

no address

Context

config>service>vprn>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command assigns an IP address, IP subnet, and broadcast address format to a VPRN IP router interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each VPRN IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7210 SAS.

The local subnet that the **address** command defines must be part of the services address space within the routing context using the **config router service-prefix** command. The default is to disallow the complete address space to services. When a portion of the address space is allocated as a service prefix, that portion can be made unavailable for IP interfaces defined within the **config router interface** CLI context for network core connectivity with the **exclude** option in the **config router service-prefix** command.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted-decimal notation. The show commands display CIDR notation and is stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

The **no** form of this command removes the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Address	Admin state	Oper state
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable. The address and admin states are the only controlling variables and can be set independently. If an address is assigned to an interface that is in an administratively up state, it becomes operationally up and the protocol interfaces and the MPLS LSPs associated with that IP interface are reinitialized.

Parameters

ip-address

Specifies the IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted-decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

Values a.b.c.d (no multicast/broadcast address)

/

The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "/" and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted-decimal mask must follow the prefix. The IPv6-prefix is x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d, where x: [0 to FFFF]H, d: [0 to 255]D and the ipv6-prefix-length is 0 to 128.

mask

Specifies the subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address.

Values 0 to 32

netmask

Specifies the subnet mask in dotted-decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted-decimal mask.

The *mask* parameter indicates the complete mask that is used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted-decimal addresses in the range 128.0.0.0 to 255.255.255.254. A mask of 255.255.255.255 is reserved for system IP addresses.

Values a.b.c.d (network bits all 1 and host bits all 0)

broadcast

The optional **broadcast** parameter overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) is received by the IP interface.

Default host-ones

all-ones

The **all-ones** keyword following the **broadcast** parameter 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

The **host-ones** keyword following the **broadcast** parameter 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 the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

allow-directed-broadcasts

Syntax

[no] **allow-directed-broadcasts**

Context

config>service>vpn>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 to the subnet broadcast address of the egress IP interface.

When enabled, a frame destined to the local subnet on this IP interface is 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 are counted in the normal discard counters for the egress SAP.

By default, directed broadcasts are not allowed and are 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

bfd

Syntax

bfd *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier*] [**echo-receive** *echo-interval*]

no bfd

Context

config>service>vpn>if

config>service>ies>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the BFD parameters for the associated IP interface. If no parameters are defined, the default value is 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 and the upper level protocols (OSPF, IS-IS, BGP) are notified of the fault.

Note: See the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide* for information about the list of routing and MPLS protocols and features that use BFD for protection on 7210 SAS platforms.

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

Default

no bfd

Parameters

transmit-interval

Specifies the transmit interval for the BFD session.

Values 10 to 100000

Default 100

receive receive-interval

Specifies the receive interval for the BFD session.

Values 10 to 100000

Default 100

multiplier multiplier

Specifies the multiplier for the BFD session.

Values 3 to 20

Default 3

echo-receive echo-interval

Specifies the minimum echo receive interval, in milliseconds, for the BFD session.

Values 100 to 100000

Default 100

cflowd-parameters

Syntax

cflowd-parameters

Context

config>service>vpn>interface

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

Commands in this context configure traffic sampling for the interface.

sampling

Syntax

sampling {unicast|multicast} type {interface} [direction {ingress-only}]
no sampling {unicast|multicast}

Context

config>service>vprn>interface>cflowd-parameters

Platforms

7210 SAS-Mxp and 7210 SAS-Sx/S 1/10GE (standalone)

Description

This command enables traffic sampling for the interface. See "Configuration Notes" in the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide* for more information.

The **no** form of this command disables traffic sampling for the interface.

Default

no sampling

Parameters

unicast

Keyword to enable unicast sampling.

multicast

Keyword to enable multicast sampling.

type

Keyword to configure the cflowd sampling type.

interface

Keyword to configure interface cflowd sampling type.

direction

keyword to configure the direction of the cflowd analysis.

ingress-only

Keyword to configure the ingress direction only for cflowd analysis.

dhcp6-relay

Syntax

[no] dhcp6-relay

Context

```
config>service>vprn>if>ipv6
```

Platforms

7210 SAS-Mxp

Description

This command enables DHCPv6 relay for the interface.

The **no** form of this command disables DHCPv6 relay.

option**Syntax**

[no] option

Context

```
config>service>vprn>if>ipv6>dhcp6-relay
```

Platforms

7210 SAS-Mxp

Description

This command configures DHCPv6 relay information options.

The **no** form of this 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
```

Platforms

7210 SAS-Mxp

Description

This command send interface ID options in the DHCPv6 relay packet.

The **no** form of this command disables the sending of interface ID options in the DHCPv6 relay packet.

Parameters

ascii-tuple

Specifies the use of the ASCII-encoded concatenated tuple, which consists of the *access-node-identifier*, *service-id*, and *interface-name*, separated by "|".

ifindex

Specifies the use of the interface index. Display the interface index of a router interface by using the **show router if detail** command.

sap-id

Specifies the use of the SAP identifier.

string

Specifies a string of up to 32 characters, composed of printable, seven-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

remote-id

Syntax

[no] **remote-id**

Context

config>service>vprn>if>ipv6>dhcp6-relay>option

Platforms

7210 SAS-Mxp

Description

This command sends 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 this command disables the sending of the remote ID option in the DHCPv6 relay packet.

server

Syntax

[no] **server** *ipv6z-address*

Context

config>service>vprn>if>ipv6>dhcp6-relay

d: [0 to 255]D

local-proxy-arp

Syntax

[no] local-proxy-arp

Context

config>service>vprn>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables local proxy ARP on an IP interface. When this command is enabled, the system responds with its own MAC address to all ARP requests for IP addresses belonging to the subnet, and therefore becomes the forwarding point for all traffic between hosts in that subnet.

When this command 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>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated interface cannot be bound to a SAP.

When using mtrace/mstat in an L3 VPN context, the configuration for the VPRN should have a loopback address configured that has the same address as the system address of the core instance (BGP next-hop).

proxy-arp-policy

Syntax

[no] proxy-arp-policy *policy-name* [*policy-name*...(up to 5 max)]

Context

config>service>vpn>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables a proxy ARP policy for the interface.

The **no** form of this command disables the proxy ARP capability.

Default

no proxy-arp

Parameters

policy-name

Specifies the export route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

remote-proxy-arp

Syntax

[no] remote-proxy-arp

Context

config>service>vpn>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables remote proxy ARP on the interface.

Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet

on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.

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>vpn>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command assigns up to 64 secondary IP addresses to the interface, including the primary IP address. Each address can be configured in an IP address, IP subnet, or broadcast address format.

Parameters

ip-address

Specifies the IP address of the IP interface. The *ip-address* portion of the address command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

/

The forward slash is a parameter delimiter that separates the *ip-address* portion of the IP address from the *mask* that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "/" and the *mask* parameter. If a forward slash does not immediately follow the *ip-address*, a dotted-decimal *netmask* must follow the prefix.

mask

Specifies the subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask* parameter. The *mask* parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 1 to 32. A mask length of 32 is reserved for system IP addresses.

Values 1 to 32

netmask

Specifies the subnet mask in dotted-decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted-decimal

mask. The *netmask* parameter indicates the complete mask that is used in a logical 'AND' function to derive the local subnet of the IP address. A netmask of 255.255.255.255 is reserved for system IP addresses.

Values a.b.c.d (network bits all 0 and host bits all 1)

broadcast {all-ones | host-ones}

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** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the broadcast type to **host-ones** after being configured as **all-ones**, the **address** command must be executed with the **broadcast** parameter defined. 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**) is received by the IP interface

Values **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* or *netmask* with all of the host bits set to binary 1. This is the default broadcast address used by an IP interface.

Default host-ones

igp-inhibit

Specifies that the secondary IP address should not be recognized as a local interface by the running IGP.

static-arp

Syntax

static-arp *ip-address* *ieee-address*

no static-arp *ip-address* [*ieee-address*]

static-arp [*ieee-addr*] **unnumbered**

no static-arp *ieee-addr* **unnumbered**

Context

config>service>vpn>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 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 with the new MAC address.

When the **unnumbered** keyword is used, this command configures a static ARP entry associating an unnumbered interface 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 unnumbered interface.

If an entry for a particular unnumbered interface already exists and a new MAC address is configured for the interface, the existing MAC address is replaced with the new MAC address.

The **no** form of this command removes a static ARP entry.

Parameters

ip-address

Specifies the IP address for the static ARP in IP address dotted-decimal notation.

ieee-address

Specifies the 48-bit MAC address for the static ARP in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

unnumbered

Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. When this command is configured, it overrides any dynamic ARP.

8.4.2.1.6 Router advertisement commands

router-advertisement

Syntax

[no] router-advertisement

Context

config>service>vpn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures router advertisement properties. By default, it is disabled for all IPv6 enabled interfaces.

The **no** form of this command disables all IPv6 interfaces. However, the **no interface *interface-name*** command disables a specific interface.

Default

disabled

interface

Syntax

[no] **interface** *ip-int-name*

Context

config>service>vpn>router-advertisement

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures router advertisement properties on a specific interface. The interface must already exist in the **config>router>interface** context.

Default

No interfaces are configured.

Parameters

ip-int-name

Specifies the interface name. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

current-hop-limit

Syntax

current-hop-limit *number*

no current-hop-limit

Context

config>service>vpn>router-advert>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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

Specifies the hop limit number.

Values 0 to 255. A value of zero means there is an unspecified number of hops.

managed-configuration

Syntax

[no] managed-configuration

Context

config>service>vpn>router-advert>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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

[no] max-advertisement-interval *seconds*

Context

```
config>service>vpn>router-advert>if
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the maximum interval between sending router advertisement messages.

Default

600

Parameters***seconds***

Specifies the maximum interval in seconds between sending router advertisement messages.

Values 4 to 1800

min-advertisement-interval**Syntax**

```
[no] min-advertisement-interval seconds
```

Context

```
config>service>vpn>router-advert>if
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the minimum interval between sending ICMPv6 neighbor discovery router advertisement messages.

Default

200

Parameters***seconds***

Specifies the minimum interval in seconds between sending ICMPv6 neighbor discovery router advertisement messages.

Values 3 to 1350

mtu

Syntax

[no] **mtu** *mtu-bytes*

Context

config>service>vpn>router-advert>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the MTU for the nodes to use to send packets on the link.

Default

no mtu

Parameters

mtu-bytes

Specifies the MTU for the nodes to use to send packets on the link.

Values 1280 to 9212

other-stateful-configuration

Syntax

[no] **other-stateful-configuration**

Context

config>service>vpn>router-advert>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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

[no] prefix [ipv6-prefix|prefix-length]

Context

config>service>vpn>router-advert>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 explicitly configured using prefix statements.

Parameters

ip-prefix

Specifies the IP prefix for prefix list entry in dotted-decimal notation.

Values	
ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	0 to 32
ipv6-prefix	x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 — FFFF]H d: [0 — 255]D
ipv6-prefix-length	0 to 128

prefix-length

Specifies a route must match the most significant bits and have a prefix length.

Values 1 to 128

autonomous

Syntax

[no] autonomous

Context

config>service>vpn>router-advert>if>prefix

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies whether the prefix can be used for stateless address autoconfiguration.

Default

enabled

on-link

Syntax

[no] on-link

Context

config>service>vpn>router-advert>if>prefix

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies whether the prefix can be used for onlink determination.

Default

enabled

preferred-lifetime

Syntax

[no] preferred-lifetime {seconds | infinite}

Context

config>service>vpn>router-advert>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the remaining length of time in seconds that this prefix continues to be preferred, such as, time until deprecation. The address generated from a deprecated prefix should not be used as a source address in new communications, but packets received on such an interface are processed as expected.

Default

604800

Parameters

seconds

Specifies the remaining length of time in seconds that this prefix continues to be preferred.

infinite

Specifies that the prefix is always preferred. A value of 4,294,967,295 represents infinity.

valid-lifetime

Syntax

valid-lifetime {*seconds* | *infinite*}

Context

config>service>vpn>router-advert>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies the length of time in seconds that the prefix is valid for the purpose of on-link determination. A value of all one bits (0xffffffff) represents infinity.

The address generated from an invalidated prefix should not appear as the destination or source address of a packet.

Default

2592000

Parameters

seconds

Specifies the remaining length of time, in seconds, that this prefix continues to be valid.

infinite

Specifies that the prefix is always valid. A value of 4,294,967,295 represents infinity.

reachable-time

Syntax

reachable-time *milli-seconds*

no reachable-time

Context

config>service>vpn>router-advert>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures how long this router should be considered reachable by other nodes on the link after receiving a reachability confirmation.

Default

no reachable-time

Parameters

milli-seconds

Specifies the length of time in milliseconds the router should be considered reachable.

Values 0 to 3600000

retransmit-time

Syntax

retransmit-timer *milli-seconds*

no retransmit-timer

Context

config>service>vpn>router-advert>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the retransmission frequency of neighbor solicitation messages.

Default

no retransmit-time

Parameters

milli-seconds

Specifies in milliseconds how often the retransmission should occur.

Values 0 to 1800000

router-lifetime

Syntax

router-lifetime *seconds*

no router-lifetime

Context

config>service>vpn>router-advert>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command sets the router lifetime.

Default

1800

Parameters

seconds

Specifies the length of time, in seconds, (relative to the time the packet is sent) that the prefix is valid for route determination.

Values 0, 4 to 9000 seconds. 0 means that the router is not a default router on this link.

ip-mtu

Syntax

ip-mtu *octets*

no ip-mtu

Context

config>service>vpn>if

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command configures the maximum IP transmit unit (packet) for the interface.

The MTU that is advertised from the VPRN size is:

MINIMUM((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))

By default (for Ethernet network interface) if **no ip-mtu** is configured, the packet size is $(1568 - 14) = 1554$.

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

Default

no ip-mtu

Parameters

octets

Specifies the number of octets in the IP-MTU.

Values 512 to 9000

8.4.2.1.7 Interface ICMP commands

icmp

Syntax

icmp

Context

config>service>vpn>if

config>service>vpn>sub-if>grp-if

```
config>service>vprn>nw-if
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures Internet Control Message Protocol (ICMP) parameters on a VPRN service.

mask-reply

Syntax

[no] mask-reply

Context

```
config>service>vprn>if>icmp
```

```
config>service>vprn>sub-if>grp-if>icmp
```

```
config>service>vprn>nw-if>icmp#
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance replies to mask requests.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

Default

mask-reply

redirects

Syntax

redirects [*number seconds*]

no redirects

Context

```
config>service>vpn>if>icmp
config>service>vpn>sub-if>grp-if>icmp
config>service>vpn>nw-if>icmp#
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.

When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

The rate at which ICMP redirects is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of redirect messages that can be issued on the interface for a specific time interval.

By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default

redirects 100 10

Parameters

number

Specifies the maximum number of ICMP redirect messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 to 1000

seconds

Specifies the time frame in seconds used to limit the *seconds* of ICMP redirect messages that can be issued.

Values 1 to 60

t1l-expired

Syntax

t1l-expired *number seconds*

no t1l-expired

Context

```
config>service>vprn>if>icmp  
config>service>vprn>sub-if>grp-if>icmp  
config>service>vprn>nw-if>icmp#
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the rate Internet Control Message Protocol (ICMP) TTL expired messages are issued by the IP interface.

By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default

```
ttl-expired 100 10
```

Parameters

number

Specifies the maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. This parameter must be specified with the *seconds* parameter.

Values 10 to 1000

seconds

Specifies 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>vprn>if>icmp  
config>service>vprn>sub-if>grp-if>icmp  
config>service>vprn>nw-if>icmp#
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables and configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a specific time interval.

By default, generation of ICMP destination unreachable messages is enabled at a maximum rate of 10 per 10 second time interval.

The **no** form of this command disables the generation of icmp destination unreachable messages on the router interface.

Default

unreachables 100 10

Parameters

number

Specifies the maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 to 1000

seconds

Specifies the time frame in seconds used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 to 60

8.4.2.1.8 Interface SAP commands

```
sap
```

Syntax

```
sap sap-id [create]
```

```
no sap sap-id
```

Context

```
config>service>vprn>if
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates a SAP within a service. A SAP is a combination of port and encapsulation parameters which identifies the SAP on the interface and within the 7210 SAS. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP does not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface port-type port-id mode access** command.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service is 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 are also deleted.

Default

No SAPs are defined.

Special Cases

VPRN

A VPRN SAP must be defined on an Ethernet interface.

sap ipsec-id.private | public:tag — This parameter associates an IPsec group SAP with this interface. This is the public side for an IPsec tunnel. Tunnels referencing this IPsec group in the private side may be created if their local IP is in the subnet of the interface subnet and the routing context specified matches with the one of the interface.

This context provides a SAP to the tunnel. The operator may associate an ingress and egress QoS policies as well as filters and virtual scheduling contexts. Internally, this creates an Ethernet SAP that is used to send and receive encrypted traffic to and from the MDA. Multiple tunnels can be associated with this SAP. The "tag" is a dot1q value. The operator may see it as an identifier. The range is limited to 1 to 4094.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

create

Keyword used to create a SAP instance.

tod-suite

Syntax

tod-suite *tod-suite-name*

no tod-suite

Context

config>service>vprn>if>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command applies a time-based policy (filter or QoS policy) to the SAP. The suite name must already exist in the **config>cron** context.

Default

no tod-suite

Parameters

tod-suite-name

Specifies collection of policies (ACLs, QoS) including time-ranges that define the full or partial behavior of a SAP or a subscriber. The suite can be applied to more than one SAP.

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

config>service>vprn>if>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates the accounting policy context that can be applied to an interface SAP or interface SAP spoke-SDP.

An accounting policy must be defined before it can be associated with a SAP. If the *policy-id* does not exist, an error message is generated.

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

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

Default

Default accounting policy.

Parameters

acct-policy-id

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

Values 1 to 99

collect-stats

Syntax

[no] collect-stats

Context

config>service>vprn>if>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables accounting and statistical data collection for either an interface SAP or interface SAP spoke-SDP, or network port. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU 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

no collect-stats

arp-timeout

Syntax

arp-timeout *seconds*

no arp-timeout

Context

config>service>vprn>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the minimum time in seconds 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 **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

The **no** form of this command restores **arp-timeout** to the default value.

Default

14400 seconds

Parameters

seconds

Specifies the minimum number of seconds a learned ARP entry is stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries is not aged.

Values 0 to 65535

delayed-enable

Syntax

delayed-enable *seconds* [init-only]

no delayed-enable

Context

config>service>vprn>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command delays making the interface operational by the specified number of seconds.

In environments with many subscribers, it can take time to synchronize the subscriber state between peers when the subscriber-interface is enabled (perhaps, after a reboot). To ensure that the state has time to be synchronized, the **delayed-enable** timer can be specified. The optional parameter **init-only** can be added to use this timer only after a reboot.

Default

no delayed-enable

Parameters

seconds

Specifies the number of seconds to delay before the interface is operational.

Values 1 to 1200

init-only

Delays the initialization of the subscriber-interface to give the rest of the system time to complete necessary tasks such as allowing routing protocols to converge and to allow MCS to sync the subscriber information. The delay only occurs immediately after a reboot.

8.4.2.1.9 Interface SAP filter and QoS policy commands

egress

Syntax

egress

Context

config>service>vprn>if>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure egress SAP Quality of Service (QoS) policies and filter policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

ingress

Syntax

ingress

Context

config>service>vpn>if>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure ingress SAP Quality of Service (QoS) policies and filter policies. If no SAP ingress QoS policy is defined, the system default SAP ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

aggregate-meter-rate

Syntax

aggregate-meter-rate *rate-in-kbps* [**burst** *burst-in-kbits*]
no aggregate-meter-rate

Context

config>service>vpn>if>sap>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, excluding platforms configured in the access-uplink operating mode; not supported on 7210 SAS-Sx 10/100GE

Description

This command allows the user to configure the SAP aggregate policer. The rate of the SAP aggregate policer must be specified by the user. The user can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the traffic on different FCs and determines the destination of the packet. The packet is either forwarded to an identified profile or dropped.



Note:
The sum of CIR of the individual FCs configured under the SAP cannot exceed the PIR rate configured for the SAP. Though the 7210 SAS software does not block this configuration, it is not recommended for use.

The following table provides information about the final disposition of the packet based on the operating rate of the per FC policer and the per SAP aggregate policer.

Table 108: Final disposition of the packet based on per-FC and per-SAP policer or meter

Per FC meter operating rate	Per FC assigned color	SAP aggregate meter operating rate	SAP aggregate meter color	Final packet color
Within CIR	Green	Within PIR	Green	Green or In-profile
Within CIR ¹⁶	Green	Above PIR	Red	Green or In-profile
Above CIR, Within PIR	Yellow	Within PIR	Green	Yellow or Out-of-Profile
Above CIR, Within PIR	Yellow	Above PIR	Red	Red or Dropped
Above PIR	Red	Within PIR	Green	Red or Dropped
Above PIR	Red	Above PIR	Red	Red or Dropped

When the SAP aggregate policer is configured, per FC policer can be only configured in "trtcm2" mode (RFC 4115).

**Note:**

The meter modes "srtcm" and "trtcm1" are used in the absence of an aggregate meter.

The SAP ingress meter counters increment the packet or octet counts based on the final disposition of the packet.

If ingress Frame-based accounting is used, the SAP aggregate meter rate accounts for the Ethernet frame overhead. The system accounts for 12 bytes of IFG and 8 bytes of start delimiter.

The **no** form of this command removes the aggregate policer from use.

Default

no aggregate-meter-rate

Parameters***rate-in-kbps***

Specifies the rate in kilobits per second.

¹⁶ This row is not recommended for use. For more information about this, see the Note in the [aggregate-meter-rate](#) description.

Values 01 to 20000000, max

Default max

burst *burst-in-kilobits*

Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 to 2146959

Default 512

filter

Syntax

filter [**ip** [*ip-filter-id* | **ipv6** *ipv6-filter-id*]]

filter [**mac** *mac-filter-id*]

no filter [**ip** [*ip-filter-id* | **ipv6** *ipv6-filter-id*]]

no filter [**mac** *mac-filter-id*]

no filter

Context

config>service>vprn>if>sap>egress

config>service>vprn>if>sap>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface. Filter policies control the forwarding and dropping of packets based on IP matching criteria.



Note:

SAP egress QoS policies are only supported on the 7210 SAS-Mxp.

The *ip-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 returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID is not removed from the system unless the scope of the created filter is set to local.

Parameters

ip *ip-filter-id*

Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

ipv6 *ipv6-filter-id*

Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

mac *mac-filter-id*

Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 to 65535

qos

Syntax

qos *policy-id*

qos *policy-id* [**enable-table-classification**] (for 7210 SAS-Mxp only)

no qos

Context

config>service>vprn>if>sap>egress

config>service>vprn>if>sap>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command associates a Quality of Service (QoS) policy with an ingress or egress Service Access Point (SAP) or IP interface. QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined before associating the policy with a SAP or IP interface. If the *policy-id* does not exist, an error is returned.

The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress, and only allows egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second policy of same or different type replaces the earlier one with the new policy.

**Note:**

SAP egress QoS policies are only supported on the 7210 SAS-Mxp.

On the 7210 SAS-Mxp (ingress), using the **enable-table-classification** keyword enables the use of IP DSCP tables to assign FC and profile on a per-SAP ingress basis. The match-criteria configured in the service ingress policy, which require CAM resources, are ignored. Only meters from the service ingress policy are used (and the meters still require CAM resources). The IP DSCP classification policy configured in the SAP ingress policy is used to assign FC and profile. The default FC is assigned from the SAP ingress policy.

**Note:**

On the 7210 SAS-Mxp, when the interface is associated with RVPLS, the behavior of the **qos** command is affected. See the **config>service>vprn>if>vpls>ingress>enable-table-classification** and **routed-override-qos-policy** commands for more information about classification behavior for RVPLS.

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

Default

No specific QoS policy is associated with the SAP or IP interface for ingress or egress, so the default QoS policy is used.

Parameters***policy-id***

Specifies the ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.

Values 1 to 65535

enable-table-classification

Enables the use of table-based classification instead of CAM-based classification at SAP ingress. The FC and profile are taken from the IP DSCP classification policy configured in the ingress policy, along with the meters from the SAP ingress policy. Match-criteria entries in the SAP ingress policy are ignored.

8.4.2.1.10 Interface VRRP commands

ipv6

Syntax

ipv6

Context

config>service>vprn>if

Platforms

7210 SAS-Mxp

Description

Commands in this context configure VPRN IPv6 parameters.

```
vrrp
```

Syntax

vrrp *virtual-router-id* [**owner**]

no vrrp *virtual-router-id*

Context

config>service>vprn>if

config>service>vprn>if>ipv6 (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates or edits a Virtual Router ID (VRID) on the service IP interface. A VRID is internally represented in conjunction with the IP interface name. This allows the VRID to be used on multiple IP interfaces while representing different virtual router instances.

Two VRRP nodes can be defined on an IP interface. One, both, or none may be defined as owner. The nodal context of **vrrp** *virtual-router-id* is used to define the configuration parameters for the VRID.

The **no** form of this command removes the specified VRID from the IP interface. This terminates VRRP participation for the virtual router and deletes all references to the VRID. The VRID does not need to be shutdown to remove the virtual router instance.

Parameters

virtual-router-id

Specifies a new virtual router ID or one that can be modified on the IP interface.

Values 1 to 255

```
authentication-key
```

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

```
config>service>vprn>if>vrrp
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command, within the **vrrp** *virtual-router-id* context, assigns a simple text password authentication key to generate master VRRP advertisement messages and validate received VRRP advertisement messages.

The **authentication-key** command is one of the few commands not affected by the presence of the **owner** keyword. If simple text password authentication is not required, this command is not required. 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, altering the simple text password used when **authentication-type** password authentication method is used by the virtual router instance. The **authentication-type password** command does not need to be executed before defining the **authentication-key** command.

To change the current in-use password key on multiple virtual router instances:

- Identify the current master
- Shutdown 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 key

The **no** form of this command restores the default null string to the value of key.

Default

No default. The authentication data field contains the value 0 in all 16 octets.

Parameters

authentication-key

Specifies 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 *key* parameter is expressed as a string consisting of up to eight alpha-numeric characters. Spaces must be contained in quotation marks (" "). The quotation marks are not considered part of the string.

The string is case sensitive and is left-justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet (starting with IETF RFC bit position 0) 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

The hash key. The key can be any combination of ASCII characters up to 22 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks ("").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

best-path-selection

Syntax

best-path-selection

Context

config>service>vpn>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure path selection.

always-compare-med

Syntax

always-compare-med {zero | infinity}

no always-compare-med strict-as {zero | infinity}

no always-compare-med

Context

config>service>vprn>bgp>best-path-selection

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the comparison of BGP routes based on the MED attribute. The default behavior of 7210 SAS (equivalent to the **no** form of this command) is to only compare two routes on the basis of MED if they have the same neighbor AS (the first non-confed AS in the received AS_PATH attribute). Also by default, a route without a MED attribute is handled the same as though it had a MED attribute with the value 0. The **always-compare-med** command without the **strict-as** keyword allows MED to be compared even if the paths have a different neighbor AS; in this case, if neither **zero** or **infinity** is specified, the **zero** option is inferred, meaning a route without a MED is handled the same as though it had a MED attribute with the value 0. When the **strict-as** keyword is present, MED is only compared between paths from the same neighbor AS, and in this case, **zero** or **infinity** is mandatory and tells BGP how to interpret paths without a MED attribute.

Default

no always-compare-med

Parameters

zero

Specifies that for routes learned without a MED attribute that a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.

infinity

Specifies for routes learned without a MED attribute that a value of infinity ($2^{32}-1$) is used in the MED comparison. This in effect makes these routes the least desirable.

strict-as

Specifies BGP paths to be compared even with different neighbor AS.

as-path-ignore

Syntax

as-path-ignore [ipv4] [ipv6]

no as-path-ignore

Context

config>service>vprn>bgp>best-path-selection

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command determines whether the AS path is used to determine the best BGP route.

If this option is present, the AS paths of incoming routes are not used in the route selection process.

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

Default

no as-path-ignore

Parameters

ipv4

Specifies that the AS-path length is ignored for all IPv4 routes.

ipv6

Specifies that the length AS-path is ignored for all IPv6 VPRN routes.

ignore-nh-metric

Syntax

ignore-nh-metric

no ignore-nh-metric

Context

config>service>vprn>bgp>best-path-selection

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command instructs BGP to disregard the resolved distance to the BGP next-hop in its decision process for selecting the best route to a destination. When configured in the **config>router>bgp>best-path-selection** context, this command applies to the comparison of two BGP routes with the same NLRI learned from base router BGP peers. When configured in the **config>service>vprn** context, this command applies to the comparison of two BGP-VPN routes for the same IP prefix imported into the VPRN from the base router BGP instance. When configured in the **config>service>vprn>bgp>best-path-selection** context, this command applies to the comparison of two BGP routes for the same IP prefix learned from VPRN BGP peers.

The **no** form of this command restores the default behavior whereby BGP factors distance to the next-hop into its decision process.

Default

no ignore-nh-metric

ignore-router-id

Syntax

ignore-router-id

no ignore-router-id

Context

config>service>vpn>bgp>best-path-selection

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command ensures that when the current best path to a destination is learned from eBGP peer X with BGP identifier x, and a new path is received from eBGP peer Y with BGP identifier y, the best path remains unchanged if the new path is equivalent to the current best path up to the BGP identifier comparison – even if y is less than x.

The **no** form of this command restores the default behavior of selecting the route with the lowest BGP identifier (y) as best.

Default

no ignore-router-id

backup

Syntax

[no] **backup** *ip-address*

Context

config>service>vpn>if>vrrp

config>service>vpn>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures virtual router IP addresses for the interface.

bfd-enable

Syntax

[no] **bfd-enable** [*service-id*] **interface** *interface-name* **dst-ip** *ip-address*

Context

config>service>vpn>if>vrrp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command assigns a bi-directional forwarding (BFD) session, providing a heart-beat mechanism for the VRRP instance. There can only be one BFD session assigned to a specified VRRP instance, but multiple VRRP instances can use the same BFD session. If the specified interface is configured with centralized BFD, the BFD transmit and receive intervals must be 300 ms or longer.

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 configured using the BFD command under the IP interface. The virtual router initiates the BFD session after the specified interface is configured with BFD.

Parameters

service-id

Specifies the service ID of the interface that is running BFD.

Values *service-id* — 1 to 2147483648

svc-name — Specifies an existing service name up to 64 characters in length.

interface-name

Specifies the name of the interface that is 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>vpn>if>vrrp

```
config>service>vpn>if>ipv6>vrrp (7210 SAS-Mxp only)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures a VRRP initialization delay timer.

Default

no init-delay

Parameters

seconds

Specifies the initialization delay timer, in seconds, for VRRP.

Values 1 to 65535

master-int-inherit

Syntax

[no] **master-int-inherit**

Context

```
config>service>vpn>if>vrrp
```

```
config>service>vpn>if>ipv6>vrrp (7210 SAS-Mxp only)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command allows the master instance to dictate the master down timer (non-owner context only).

Default

no master-int-inherit

message-interval

Syntax

message-interval {[seconds] [milliseconds milliseconds]}

no message-interval

Context

```
config>service>vprn>if
config>service>vprn>if>ipv6>vrrp (7210 SAS-Mxp only)
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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 participating as a virtual router. Any VRRP advertisement message received with an Advertisement Interval field different from the virtual router instance configured message-interval value is silently discarded.

The message-interval command is available in both non-owner and owner **vrrp** *virtual-router-id* nodal contexts. If the message-interval command is not executed, the default message interval of 1 second is used.

The **no** form of this command restores the default message interval value of 1 second to the virtual router instance.

Parameters***seconds***

Specifies the number of seconds that transpire before the advertisement timer expires.

Values 1 to 255

Default 1

milliseconds milliseconds

Specifies the time interval, in milliseconds, between sending advertisement messages. This parameter is not supported on single-slot chassis.

Values 100 to 900

ping-reply**Syntax**

[no] ping-reply

Context

```
config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp (7210 SAS-Mxp only)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the non-owner master to reply to ICMP echo requests directed at the virtual router instances IP addresses. The ping request can be received on any routed interface.

Ping must not have been disabled at the management security level (either on the parental IP interface or based on the Ping source host address). When ping-reply is not enabled, ICMP echo requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP echo requests regardless of the setting of ping-reply configuration.

The ping-reply command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the ping-reply command is not executed, ICMP echo requests to the virtual router instance IP addresses are silently discarded.

The **no** form of this command restores the default operation of discarding all ICMP echo request messages destined to 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 (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command associates a VRRP priority control policy with the virtual router instance (non-owner context only).

Parameters

vrrp-policy-id

Specifies a VRRP priority control policy.

Values 1 to 9999

preempt

Syntax

preempt
no preempt

Context

config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command overrides an existing non-owner master to the virtual router instance. Enabling preempt mode is recommended for correct operation of the base-priority and vrrp-policy-id definitions on the virtual router instance. If the virtual router cannot preempt an existing non-owner master, the affect of the dynamic changing of the in-use priority is greatly diminished.

The preempt command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. The owner may not be preempted due to the fact that 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 virtual router instances only preempt when preempt is set and the current master has an in-use message priority value less than the virtual router instances 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 the virtual router in-use priority value
- Equal to the in-use priority value and the source IP address (primary IP address) is greater than the virtual router instance primary IP address

The **no** form of this command prevents a non-owner virtual router instance from preempting another, less desirable virtual router. Use the preempt command to restore the default mode.

Default

preempt

priority

Syntax

priority *priority*
no priority

Context

```
config>service>vprn>if>vrrp
config>service>vprn>if>ipv6>vrrp (7210 SAS-Mxp only)
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command provides configures a specific priority value to the virtual router instance. In conjunction with an optional policy command, the base-priority is used to derive the in-use priority of the virtual router instance.

The priority command is only available in the non-owner **vrrp** *virtual-router-id* nodal context. 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 to base-priority.

Parameters

base-priority

Specifies the base priority used by the virtual router instance. If a VRRP priority control policy is not also defined, the base-priority is 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
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the non-owner master to reply to SSH Requests directed at the virtual router instance IP addresses. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parental IP interface or based on the SSH source host address). Correct login and CLI command authentication is still enforced.

When `ssh-reply` is not enabled, SSH packets to non-owner master virtual IP addresses are silently discarded. Non-owner backup virtual routers never respond to SSH regardless of the `ssh-reply` configuration.

The `ssh-reply` command is only available in non-owner **vrrp** *virtual-router-id* nodal context. If the `ssh-reply` command is not executed, SSH packets to the virtual router instance IP addresses are silently discarded.

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 (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command allows the forwarding of packets by a standby router.

The **no** form of this command specifies that a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router real MAC address.

Default

no `standby-forwarding`

telnet-reply

Syntax

[no] `telnet-reply`

Context

config>service>vprn>if>vrrp

config>service>vprn>if>ipv6>vrrp (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the non-owner master to reply to TCP port 23 Telnet Requests directed at the virtual router instance IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Correct login and CLI command authentication is still enforced.

When telnet-reply 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 in non-owner VRRP nodal context. If the telnet-reply command is not executed, Telnet packets to the virtual router instance IP addresses are silently discarded.

The **no** form of this command restores the default operation of discarding all Telnet packets destined to 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 (7210 SAS-Mxp only)

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.

A non-owner backup virtual router never responds to such traceroute requests regardless of the trace-route-reply status.

Default

no traceroute-reply

8.4.2.1.11 PIM commands

pim

Syntax

[no] pim

Context

config>service>vprn

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures a Protocol Independent Multicast (PIM) instance in the VPRN service. When an PIM instance is created, the protocol is enabled. PIM is used for multicast routing within the network. Devices in the network can receive the multicast feed requested and non-participating routers can be pruned. The supports PIM sparse mode (PIM-SM).

The **no** form of this command deletes the PIM protocol instance removing all associated configuration parameters.

import

Syntax

import {join-policy | register-policy} [policy-name [.. policy-name] policy-name...up to 5 max]

no import {join-policy | register-policy}

Context

config>service>vprn>pim

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command specifies the import route policy 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 this command removes the policy association from the IGMP instance.

Default

no import join-policy

no import register-policy

Parameters

join-policy

Filters PIM join messages which prevents unwanted multicast streams from traversing the network.

register-policy

Filters register messages. PIM register filters 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. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

interface

Syntax

[no] interface *ip-int-name*

Context

config>service>vpn>pim

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

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 this command deletes the PIM interface configuration for this interface. If the **no** command parameter is configured, the **no interface** form must be saved in the configuration to avoid automatic (re)creation after the next **no** 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. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

assert-period

Syntax

assert-period *assert-period*

no assert-period

Context

config>service>vpn>pim>if

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures the period in seconds for periodic refreshes of PIM Assert messages on an interface.

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

Default

60

Parameters

assert-period

Specifies the period, in seconds, for periodic refreshes of PIM Assert messages on an interface.

Values 1 to 300

bfd-enable

Syntax

[no] bfd-enable [ipv4]

Context

config>service>vpn>pim>if

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command enables the use of bidirectional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a specific protocol interface, the state of the protocol interface is

tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

For more information about the protocols and platforms that support BFD, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

Default

no bfd-enable

bsm-check-rtr-alert

Syntax

[no] bsm-check-rtr-alert

Context

config>service>vpn>pim>if

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command enables the checking of router alert option in the bootstrap messages received on this interface.

Default

no bsm-check-rtr-alert

hello-interval

Syntax

hello-interval *hello-interval*

no hello-interval

Context

config>service>vpn>pim>if

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures the frequency 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 (zero) value disables the sending of hello messages.

Values 0 to 255 seconds

hello-multiplier**Syntax**

hello-multiplier *deci-units*

no hello-multiplier

Context

config>service>vpn>pim>if

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures the multiplier to determine the hold time for a PIM neighbor.

The **hello-multiplier** in conjunction with the **hello-interval** determines the holdtime for a PIM neighbor.

Parameters***deci-units***

Specifies the value, specified in multiples of 0.1, for the formula used to calculate the hello-holdtime based on the hello-multiplier:

$(\text{hello-interval} * \text{hello-multiplier}) / 10$

This allows the PIMv2 default timeout of 3.5 seconds to be supported.

Values 20 to 100

Default 35

improved-assert**Syntax**

[no] improved-assert

Context

```
config>service>vpn>pim>if
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command enables improved assert processing on this interface. The PIM assert process establishes a forwarder 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. This could impact performance if data is continuously received on an outgoing interface.

When enabled, the PIM assert process is done entirely on the control-plane with no interaction between the control and forwarding plane.

Default

enabled

instant-prune-echo**Syntax**

```
[no] instant-prune-echo
```

Context

```
config>service>vpn>pim>if
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command enables PIM router to echo the PIM prune message received from a downstream router. It is typically used in a multi-access broadcast network (for example: Ethernet LAN) to reduce the probability of loss of PIM prune messages.

Default

no instant-prune-echo

max-groups**Syntax**

```
max-groups value
```

```
no max-groups
```

Context

```
config>service>vpn>pim>if
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures the maximum number of groups for which PIM can have downstream state based on received PIM Joins on this interface. This 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 are not allowed. When this object 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 16000

multicast-senders

Syntax

multicast-senders {auto | always | never}

no multicast-senders

Context

```
config>service>vpn>pim>if
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

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.

priority**Syntax**

priority *dr-priority*

no priority

Context

config>service>vpn>pim>if

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command sets the priority value to become the rendezvous point (RP) that is included in bootstrap messages sent by the router. The RP is sometimes called the bootstrap router. The **priority** command indicates whether the router is eligible to be a bootstrap router.

The **no** form of this command disqualifies the router to participate in the bootstrap election.

The default value means the router is the least likely to become the designated router.

Default

1

Parameters***dr-priority***

Specifies the priority to become the designated router. The higher the value, the higher the priority.

Values 1 to 4294967295

sticky-dr**Syntax**

sticky-dr [**priority** *dr-priority*]

no sticky-dr

Context

config>service>vpn>pim>if

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command enables sticky-dr operation on this interface. When enabled, the priority in PIM hellos sent on this interface when elected as the Designated Router (DR) is modified to the value configured in *dr-priority*. This is done to avoid the delays in forwarding caused by DR recovery, when switching back to the old DR on a LAN when it comes back up.

By enabling **sticky-dr**, this interface continues to act as the DR for the LAN even after the old DR comes back up.

The **no** form of this command disables sticky-dr operation on this interface.

Default

disabled

Parameters

priority *dr-priority*

Sets the DR priority to be sent in PIM Hello messages following the election of that interface as the DR, when sticky-dr operation is enabled.

Values 1 to 4294967295

three-way-hello

Syntax

three-way-hello

no three-way-hello

Context

config>service>vprn>pim>if

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures the compatibility mode for enabling the three-way hello.

tracking-support

Syntax

[no] tracking-support

Context

```
config>service>vpn>pim>if
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

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

non-dr-attract-traffic

Syntax

[no] non-dr-attract-traffic

Context

```
config>service>vpn>pim
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

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) to IGMP and PIM. The interface IGMP state is synchronized to the backup node if it is associated with the redundant peer port. The interface can be configured to use PIM which causes multicast streams to be sent to the elected DR only. The DR is also the router sending traffic to the DSLAM. Because it may be required to attract traffic to both routers a flag non-dr-attract-traffic can be used in the PIM context to have the router ignore the DR state and attract traffic when not DR. Note that while using this flag the router may not send the stream down to the DSLAM while not DR.

When enabled, the designated router state is ignored. When disabled, **no non-dr-attract-traffic**, the designated router value is honored.

Default

no non-dr-attract-traffic

rp

Syntax

rp

Context

config>service>vpn>pim

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command enables access to the context to configure the rendezvous point (RP) of a PIM protocol instance.

A Nokia PIM router acting as an RP must respond to a PIM register message specifying an SSM multicast group address by sending to the first hop router stop register messages. It does not build an (S, G) shortest path tree toward the first hop router. An SSM multicast group address can be 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>vpn>pim>rp

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

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 this command removes the anycast instance from the configuration.

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 another 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 simply used to enter the anycast CLI level.

Values Any valid loopback address configured on the node.

rp-set-peer

Syntax

[no] **rp-set-peer** *ip-address*

Context

config>service>vpn>pim>rp>anycast

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures a peer in the anycast rp-set. The address identifies the address used by the other node as the RP candidacy address for the same multicast group address range as configured on this node.

This is a manual procedure. Caution should be taken to produce a consistent configuration of an RP-set for a specific 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 of addresses that can be configured in an rp-set, up to 15 multicast addresses is recommended.

The **no** form of this command removes an entry from the list.

Parameters

ip-address

Specifies the address used by the other node as the RP candidacy address for the same multicast group address range as configured on this node.

bootstrap-export

Syntax

bootstrap-export *policy-name* [*policy-name... up to five*]

no bootstrap-export

Context

config>service>vprn>pim>rp

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

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.

Parameters

policy-name

Specifies the policy name. 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 *policy-name* [*policy-name*... up to 5 max]

Context

config>service>vprn>pim>rp

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command imports policies to control the flow of bootstrap messages into the RP. Up to five policies can be defined.

The **no** form of this command removes the specified policy names from the configuration.

Parameters

policy-name

Specifies the policy name. The policy statement must already be configured in the config>router>policy-options context.

bsr-candidate

Syntax

bsr-candidate

Context

```
config>service>vpn>pim>rp
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

Commands in this context configure a local rendezvous point (RP) of a PIM protocol instance.

Default

Enabled when PIM is enabled.

address

Syntax

[no] address *ip-address*

Context

```
config>service>vpn>pim>rp>bsr-candidate
```

```
config>service>vpn>pim>rp>rp-candidate
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures a static bootstrap or rendezvous point (RP) as long as the source is not directly attached to this router.

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

Default

No IP address is specified.

Parameters

ip-address

Specifies the static IP address of the RP. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

hash-mask-len

Syntax

hash-mask-len *hash-mask-length*

no hash-mask-len

Context

config>service>vpn>pim>rp>bsr-candidate

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command is used to configure the length of a mask that is to be combined with the group address before the hash function is called. All groups with the same hash map to the same RP. For example, if this 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.

Parameters

hash-mask-length

The hash mask length.

Values 0 to 32

priority

Syntax

priority *bootstrap-priority*

Context

config>service>vpn>pim>rp>bsr-candidate

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command defines the priority used to become 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 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>vpn>pim>rp

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

Commands in this context 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>vpn>pim>rp>rp-candidate

config>service>vpn>pim>ssm

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures the group address or range of group addresses for which this router can be the rendezvous point (RP).

The **no** form of this command removes the group address or range of group addresses for which this router can be the RP from the configuration.

Parameters

group-ip-address

Specifies the addresses or address ranges that this router can be an RP.

mask

Specifies the address mask with the address to define a range of addresses.

netmask

Specifies the subnet mask in dotted-decimal notation.

Values a.b.c.d (network bits all 1 and host bits all 0)

holdtime

Syntax

holdtime *holdtime*

no holdtime *holdtime*

Context

config>service>vpn>pim>rp>rp-candidate

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command to defines the length of time neighboring routers consider 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 0 to 255

priority

Syntax

priority *priority*

no priority *priority*

Context

```
config>router>pim>rp>local  
config>service>vpn>pim>rp>rp-candidate
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command defines the priority used to become the rendezvous point (RP). The higher the priority value, the more likely that this router becomes the RP.

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

Default

1

Parameters

priority

Specifies the priority to become the designated router. The higher the value the more likely the router becomes the RP.

Values 0 to 255

static

Syntax

static

Context

```
config>service>vpn>pim>rp
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command enables access to the context to configure a static rendezvous point (RP) of a PIM-SM protocol instance.

address

Syntax

[no] address *ip-address*

Context

```
config>service>vprn>pim>rp>static
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures the static rendezvous point (RP) address.

The override option specifies that dynamically learned RPs have less priority this static entry, by default dynamic learned RPs take preference over static configured RPs.

The **no** form of this command removes the static RP entry from the configuration.

group-prefix

Syntax

```
[no] group-prefix {grp-ip-address/mask | grp-ip-address netmask}
```

Context

```
config>service>vprn>pim>rp>static
```

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command defines a range of multicast-ip-addresses for which a certain RP is applicable.

The **no** form of this command removes the criterion.

Parameters

grp-ip-address

Specifies the multicast IP address.

mask

Specifies the mask of the multicast-ip-address.

Values 4 to 32

netmask

Specifies the subnet mask in dotted-decimal notation.

Values a.b.c.d (network bits all 1 and host bits all 0)

override

Syntax

[no] **override**

Context

config>service>vpn>pim>rp>static

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command changes the precedence of static RP over dynamically learned Rendezvous Point (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>vpn>pim

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures a shortest path tree (SPT tree) switchover threshold for a group prefix.

Parameters

grp-ip-address

Specifies the multicast group address.

mask

Specifies the mask of the multicast-ip-address.

Values 4 to 32

netmask

Specifies the subnet mask in dotted-decimal notation.

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

spt-threshold

Specifies the configured threshold in kilo-bits per second(kbps) 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)'s rate exceeds this configured threshold.

ssm-assert-compatible-mode

Syntax

ssm-assert-compatible-mode [enable | disable]

Context

config>service>vprn>pim

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command specifies whether SSM assert is enabled in compatibility mode for this PIM protocol instance. When enabled for SSM groups, PIM considers the SPT bit to be implicitly set to compute the value of CouldAssert (S,G,I) as defined in RFC 4601, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised). When disabled, for SSM groups, PIM does not assume the SPT bit to be set. The SPT bit is set by the Update_SPTbit(S,G,iif) macro defined in RFC 4601.

Default

disable

Parameters***enable***

Enables SSM assert in compatibility mode for this PIM protocol instance.

disable

Disables SSM assert in compatibility mode for this PIM protocol instance.

ssm-default-range-disable

Syntax

ssm-default-range-disable ipv4

Context

config>service>vpn>pim

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

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 default range is disabled for SSM and it can be used by ASM. When disabled, the SSM default range is enabled.

Default

disable

ssm-groups**Syntax**

[no] ssm-groups

Context

config>service>vpn

Platforms

7210 SAS-T (network operating mode) and 7210 SAS-Mxp

Description

This command configures a source-specific multicast (SSM) configuration instance.

8.4.2.1.12 Counter mode commands

statistics**Syntax**

statistics

Context

config>service>vpn>if>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure the counters associated with SAP ingress.

ingress

Syntax

ingress

Context

config>service>vprn>if>sap>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure the ingress SAP statistics counter.

counter-mode

Syntax

counter-mode {in-out-profile-count | forward-drop-count}

Context

config>service>vprn>if>sap>statistics>ingress

Platforms

7210 SAS-T (network operating mode), 7210 SAS-Mxp, and 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), and 7210 SAS-Sx 10/100GE

Description

This command sets the counter mode for the counters associated with sap ingress meters or policers. A pair of counters is available with each meter. These counters count different events based on the counter mode value.



Note:

The counter mode can be changed if an accounting policy is associated with a SAP. If the counter mode is changed, the counters associated with the meter are reset and the counts are cleared. If an accounting policy is in use when the counter-mode is changed, a new record is written into the current accounting file.

Execute the following sequence of commands on the specified SAP to ensure that the correct statistics are collected when the counter-mode is changed:

1. Execute the command **config>service>vprn>interface>sap>no collect-stats**, to disable writing of accounting records for the SAP.
2. Change the counter-mode to the desired option, execute the command **config>service>vprn>interface>sap>statistics>ingress>counter-mode {in-out-profile-count | forward-drop-count}**.
3. Execute the command **config>service>vprn>interface>sap>collect-stats**, to enable writing of accounting records for the SAP.

The **no** form of this command restores the counter mode to the default value.

Default

in-out-profile-count

Parameters

in-out-profile-count

If the counter mode is specified as **in-out-profile-count**, one counter counts the total in-profile packets and octets received on ingress of a SAP and another counts the total out-of-profile packets and octets received on ingress of a SAP. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. Dropped counts are not maintained in hardware when this mode is used. It is obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

forward-drop-count

If the counter mode is specified as **forward-drop-count**, one counter counts the forwarded packets and octets received on ingress of a SAP and another counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets/octets received on SAP ingress. The dropped count is count of packets/octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters configured. A packet is dropped by the policer if it exceeds the configured PIR rate. The in-profile count and out-of-profile count is not individually available when operating in this mode.

8.4.2.1.13 BGP commands

bgp

Syntax

[no] bgp

Context

config>service>vprn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the BGP protocol with the VPRN service.

The **no** form of this command disables the BGP protocol from the specific 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

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables or disables the advertising of inactive BGP routers to other BGP peers.

By default, BGP only advertises BGP routes to other BGP peers if a specific 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 specific destination.

Default

no advertise-inactive

aggregator-id-zero

Syntax

[no] aggregator-id-zero

Context

config>service>vprn>bgp

```
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command sets the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths.

When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute.

When this command is enabled, BGP adds the router ID to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level, while this command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of this command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute.

The **no** form of this command used at the group level reverts to the value defined at the group level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no aggregator-id-zero

always-compare-med

Syntax

```
always-compare-med {zero | infinity}
no always-compare-med
```

Context

```
config>service>vprn>bgp
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies how the Multi-Exit Discriminator (MED) path attribute is used in the BGP route selection process. The MED attribute is always used in the route selection process regardless of the peer AS that advertised the route. This parameter determines what MED value is inserted in the RIB-IN. If this parameter is not configured, only the MEDs of routes that have the same peer ASs are compared.

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

Default

no always-compare-med

Parameters**zero**

Specifies that for routes learned without a MED attribute that a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.

infinity

Specifies for routes learned without a MED attribute that a value of infinity (4294967295) is used in the MED comparison. This in effect makes these routes the least desirable.

as-path-ignore**Syntax**

[no] as-path-ignore

Context

config>service>vpn>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command determines whether the AS path is used to determine the best BGP route.

If this option is present, the AS paths of incoming routes are not used in the route selection process.

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

Default

no as-path-ignore

as-override**Syntax**

[no] as-override

Context

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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's loop detection mechanism. It should be used carefully.

Default

not enabled

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

config>service>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 MD-5 message-based digest. The authentication key can be any combination of letters or numbers from 1 to 16.

The **no** form of this 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

Specifies the authentication key. The key can be any combination of ASCII characters up to 255 characters (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 342 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

auth-keychain

Syntax

auth-keychain *name*

Context

config>service>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the BGP authentication key for all peers.

The keychain allows the rollover of authentication keys during the lifetime of a session.

Default

no auth-keychain

Parameters***name***

Specifies the name of an existing keychain, up to 32 characters, to use for the specified TCP session or sessions.

backup-path

Syntax

[no] backup-path [ipv4] [ipv6]

Context

config>service>vprn>bgp

Platforms

7210 SAS-T, 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC), 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp.

Description

This command enables the computation and use of a backup path for IPv4 and IPv6 BGP-learned prefixes belonging to the base router or a particular VPRN. Multiple paths must be received for a prefix 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 re-convergence to occur. When many prefixes share the same primary paths, and in some cases also the same backup path, the time to failover traffic to the backup path is independent of the number of prefixes.

By default, IPv4 and IPv6 prefixes do not have a backup path installed in the IOM.

Default

no backup-path

Parameters

ipv4

Enables the use of a backup path for BGP-learned unlabeled IPv4 prefixes.

ipv6

Enables the use of a backup path for BGP-learned unlabeled IPv6 prefixes.

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

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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), peer-group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

120

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

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables BGP route damping for learned routes which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route policy definition.

The **no** form of this command used at the global level disables route damping.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:

Half-life:	15 minutes
Max-suppress:	60 minutes
Suppress-threshold:	3000
Reuse-threshold	750

Default

no damping

disable-4byte-asn

Syntax

[no] **disable-4byte-asn**

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command disables the use of 4-byte AS numbers. It can be configured at all 3 level of the hierarchy so it can be specified down to the per peer basis.

If this command is enabled 4-byte AS number support should not be negotiated with the associated remote peers.

The **no** form of this command resets the behavior to the default which is to enable the use of 4-byte AS number.

disable-capability-negotiation

Syntax

[no] **disable-capability-negotiation**

Context

```
config>service>vpn>bgp>group  
config>service>vpn>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command disables the exchange of capabilities. When this command is enabled and after the peering is flapped, any new capabilities are not negotiated and strictly support IPv4 routing exchanges with that peer.

The **no** form of this command removes this command from the configuration and restores the normal behavior.

Default

no disable-capability-negotiation

disable-capability-negotiation

Syntax

[no] disable-capability-negotiation

Context

```
config>service>vpn>bgp  
config>service>vpn>bgp>group  
config>service>vpn>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command disables the exchange of capabilities. When command is enabled and after the peering is flapped, any new capabilities are not negotiated and strictly support IPv4 routing exchanges with that peer.

The **no** form of this command removes this command from the configuration and restores the normal behavior.

Default

no disable-capability-negotiation

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

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures BGP to disable sending communities.

Parameters

standard

Specifies standard communities that existed before VPRNs or 2547.

extended

Specifies BGP communities used were expanded after the concept of 2547 was introduced, to include handling the VRF target.

disable-fast-external-failover

Syntax

[no] disable-fast-external-failover

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures BGP fast external failover.

```
enable-peer-tracking
```

Syntax

[no] **enable-peer-tracking**

Context

config>service>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables BGP peer tracking.

Default

no enable-peer-tracking

```
export
```

Syntax

export *policy* [*policy*...]

no export

Context

config>service>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies the export policies to be used to control routes advertised to BGP neighbors.

When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be configured. The first policy that matches is applied.

Note that if a non-existent route policy is applied to a VPRN instance, the CLI generates a warning message. This message is only generated at an interactive CLI session and the route policy association is made. No warning message is generated when a non-existent route policy is applied to a VPRN instance in a configuration file or when SNMP is used.

The **no** form of this command removes all route policy names from the export list.

Default

no export

Parameters

policy

Specifies a route policy statement name.

family

Syntax

family [ipv4] [ipv6]

no family

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the IP family capability.

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

Default

no family

Parameters

ipv4

Specifies IPv4 support.

ipv6

Specifies IPv6 support.

group

Syntax

group *name* [**dynamic-peer**]

no group

Context

config>service>vpn>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures a BGP peer group.

The **no** form of this command deletes the specified peer group and all configurations associated with the peer group. The group must be **shutdown** before it can be deleted.

Parameters

name

Specifies the peer group name. Allowed values is a string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

dynamic-peer

Specifies that the specific BGP group is used by BGP peers created dynamically based on subscriber-hosts pointing to corresponding BGP peering policy. There can be only one BGP group with this flag set in any specific VPRN. No BGP neighbors can be manually configured in a BGP group with this flag set.

Default disabled

neighbor

Syntax

[**no**] **neighbor** *ip-address*

Context

config>service>vpn>bgp>group

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 configuration.

The **no** form of this command removes the specified neighbor and the entire configuration associated with the neighbor. The neighbor must be administratively **shutdown** before attempting to delete it. If the neighbor is not shutdown, the command does not result in any action except a warning message on the console indicating that neighbor is still administratively up.

Parameters

ip-address

Specifies the IP address of the BGP peer router in dotted-decimal notation.

Values ipv4-address : a.b.c.d

family

Syntax

family [ipv4] [ipv6]

no family

Context

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies the address family or families to be supported over BGP peerings in the base router. This command is additive so issuing the **family** command adds the specified address family to the list.

The **no** form of this command removes the specified address family from the associated BGP peerings. If an address family is not specified, reset the supported address family back to the default.

Default

ipv4

Parameters

ipv4

Specifies support for IPv4 routing information.

ipv6

Specifies support for IPv6 routing information.

hold-time

Syntax

hold-time *seconds* [**strict**]

no hold-time

Context

config>service>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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.

Even though the router OS implementation allows setting the [keepalive](#) time separately, the configured [keepalive](#) timer is overridden by the [hold-time](#) value under the following circumstances:

1. If the specified hold-time is less than the configured [keepalive](#) time, the operational [keepalive](#) time is set to a third of the [hold-time](#); the configured [keepalive](#) time is not changed.
2. If the [hold-time](#) is set to zero, the operational value of the [keepalive](#) time is set to zero; the configured [keepalive](#) time is not changed. This means that the connection with the peer is up permanently and no [keepalive](#) packets are sent to the peer.

The **no** form of this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

90

Parameters

seconds

The hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.

Values 0, 3 to 65535

strict

Specifies the advertised BGP hold-time from the far-end BGP peer must be greater than or equal to the specified value.

import**Syntax**

import *policy* [*policy*...]

no import

Context

config>service>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies the import policies to be used to control routes advertised to BGP neighbors. Route policies are configured in the **config>router>policy-options** context. When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be specified. The first policy that matches is applied.

The **no** form of this command removes all route policy names from the import list.

Default

no import

Parameters

policy

Specifies a route policy statement name.

keepalive**Syntax**

keepalive *seconds*

no keepalive

Context

config>service>vpn>bgp

```
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires. The *seconds* parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The *keepalive* value is generally one-third of the *hold-time* interval. Even though the OS implementation allows the *keepalive* value and the *hold-time* interval to be independently set, under the following circumstances, the configured *keepalive* value is overridden by the *hold-time* value:

If the specified *keepalive* value is greater than the configured *hold-time*, the specified value is ignored, and the *keepalive* is set to one third of the current *hold-time* value.

If the specified *hold-time* interval is less than the configured *keepalive* value, the *keepalive* value is reset to one third of the specified *hold-time* interval.

If the *hold-time* interval is set to zero, 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 this command used at the global level reverts to the default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

30

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>vpn>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Configures the local IP address used by the group or neighbor when communicating with BGP peers.

Outgoing connections use the **local-address** as the source of the TCP connection when initiating connections with a peer.

When a local address is not specified, the 7210 SAS uses the system IP address when communicating with iBGP peers and uses the interface address for directly connected eBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of this command removes the configured local-address for BGP.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Parameters

no local-address

Specifies the router ID is used when communicating with iBGP peers and the interface address is used for directly connected eBGP peers.

ip-address

Specifies the local address expressed in dotted-decimal notation. Allowed values are a valid routable IP address on the router, either an interface or system IP address.

local-as

Syntax

```
local-as as-number [private]
```

```
no local-as
```

Context

```
config>service>vpn>bgp
```

```
config>service>vpn>bgp>group
```

```
config>service>vpn>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 message before the router AS number makes the virtual AS the second AS in the as-path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). Therefore, by specifying this at each neighbor level, it is possible to have a separate as-number per eBGP session.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The **private** attribute can be added or removed dynamically by reissuing the command.

Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to reestablish 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 reestablish the peer relationship with the new local AS number.

This is an optional command and can be used in the following circumstance:

Provider router P is moved from AS1 to AS2. The customer router that is connected to P, however, is configured to belong to AS1. To avoid reconfiguring the customer router, the **local-as** value on router P can be set to AS1. Therefore, router P adds AS1 to the as-path message for routes it advertises to the customer router.

The **no** form of this command used at the global level removes any virtual AS number configured.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no local-as

Parameters

as-number

Specifies the virtual autonomous system number, expressed as a decimal integer.

Values 1 to 65535

private

Specifies the local-as is hidden in paths learned from the peering.

local-preference

Syntax

local-preference *local-preference*

no local-preference

Context

`config>service>vprn>bgp`

`config>service>vprn>bgp>group`

```
config>service>vpn>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command sets the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute. This value is used if the BGP route arrives from a BGP peer without the **local-preference** integer set.

The specified value can be overridden by any value set via a route policy. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

The default of no-local-preference does not override the local-preference value set in arriving routes and analyze routes without local preference with value of 100.

Default

no local-preference

Parameters

local-preference

Specifies 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>vpn>bgp
```

```
config>service>vpn>bgp>group
```

```
config>service>vpn>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

Note that dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of this command used at the global level reverts to default, which is **loop-detect ignore-loop**.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

loop-detect ignore-loop

Parameters

drop-peer

Sends a notification to the remote peer and drops the session.

discard-route

Discards routes received with loops in the AS path.

ignore-loop

Ignores routes with loops in the AS path but maintains peering.

off

Disables loop detection.

med-out

Syntax

med-out {number | igp-cost}

no med-out

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set via a route policy.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to default where the MED is not advertised.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no med-out

Parameters

number

Specifies the MED path attribute value, expressed as a decimal integer.

Values 0 to 4294967295

igp-cost

Specifies the MED is set to the IGP cost of the specific IP prefix.

min-as-origination

Syntax

min-as-origination *seconds*

no min-as-origination

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the minimum interval, in seconds, at which a path attribute, originated by the local router, can be advertised to a peer.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of this command used at the global level reverts to default.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

15

Parameters

seconds

Specifies the minimum path attribute advertising interval, in seconds, expressed as a decimal integer.

Values 2 to 255

min-route-advertisement

Syntax

min-route-advertisement *seconds*

no min-route-advertisement

Context

config>service>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

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

Default

30

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>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the Time To Live (TTL) value entered in the IP header of packets sent to an eBGP peer multiple hops away.

This parameter is meaningful only when configuring eBGP peers. It is ignored if set for an iBGP peer.

The **no** form of this command is used to convey to the BGP instance that the eBGP peers are directly connected.

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

Default

1 — eBGP peers are directly connected.

64 — iBGP

Parameters

ttl-value

Specifies the TTL value, expressed as a decimal integer.

Values 1 to 255

next-hop-self

Syntax

[no] **next-hop-self**

Context

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the group or neighbor to always set the NEXTHOP path attribute to its own physical interface when advertising to a peer.

This is primarily used to avoid third-party route advertisements when connected to a multi-access network.

The **no** form of this command used at the group level allows third-party route advertisements in a multi-access network.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

The default means that third-party route advertisements are allowed.

Default

no next-hop-self

peer-as

Syntax

peer-as *as-number*

Context

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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, because the peer is in a different autonomous system than this router.

For iBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level.

This is a required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.

Default

No AS numbers are defined.

Parameters

as-number

The autonomous system number, expressed as a decimal integer.

Values 1 to 65535

preference

Syntax

[no] **preference** *preference*

Context

config>service>vpn>bgp

config>service>vpn>bgp>group

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The lower the preference the higher the chance of the route being the active route. The OS assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF.

The **no** form of this command used at the global level reverts to default value.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

170

Parameters

preference

Specifies the route preference, expressed as a decimal integer.

Values 1 to 255

path-mtu-discovery

Syntax

[no] path-mtu-discovery

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables path MTU discovery for the associated TCP connections. In doing so, the MTU for the associated TCP session is initially set to the egress interface MTU. The DF bit is also set so that if a router along the path of the TCP connection cannot handle a packet of a particular size without fragmenting, it sends back an ICMP message to set the path MTU for the specific session to a lower value that can be forwarded without fragmenting.

The **no** form of this command disables path MTU discovery.

Default

no path-mtu-discovery

prefix-limit

Syntax

prefix-limit *limit* [**log-only**] [**threshold percent**]

no prefix-limit

Context

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the maximum number of routes BGP can learn from a peer.

When the number of routes reaches a certain percentage (default is 90% of this limit), an SNMP trap is sent. When the limit is exceeded, the BGP peering is dropped and disabled.

The **no** form of this command removes the **prefix-limit**.

Default

no prefix-limit

Parameters

limit

Specifies the number of routes that can be learned from a peer, expressed as a decimal integer.

Values 1 to 4294967295

log-only

Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, the BGP peering is not dropped.

percent

Specifies the threshold value (as a percentage) that triggers a warning message to be sent. The default value is 90%.

rapid-withdrawal

Syntax

[no] rapid-withdrawal

Context

config>service>vprn>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command disables the delay (Minimum Route Advertisement) on sending BGP withdrawals. Normal route withdrawals may be delayed up to the minimum route advertisement to allow for efficient packing of BGP updates.

The **no** form of this command removes this command from the configuration and returns withdrawal processing to the normal behavior.

Default

no rapid-withdrawal

remove-private

Syntax

[no] **remove-private**

Context

config>service>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command allows private AS numbers to be removed from the AS path before advertising them to BGP peers.

When the **remove-private** parameter is set at the global level, it applies to all peers regardless of group or neighbor configuration. When the parameter is set at the group level, it applies to all peers in the group regardless of the neighbor configuration.

The software recognizes the set of AS numbers that are defined by IANA as private. These are AS numbers in the range 64512 through 65535, inclusive.

The **no** form of this command used at the global level reverts to default value.

The **no** form of this command used at the group level reverts to the value defined at the global level. The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no remove-private — Private AS numbers are included in the AS path attribute.

type

Syntax

[no] **type** {internal | external}

Context

config>service>vpn>bgp>group

```
config>service>vprn>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command designates the BGP peer as type internal or external.

The type of **internal** indicates the peer is an iBGP peer while the type of external indicates that the peer is an eBGP peer.

By default, the software derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, the peer is considered **external**.

The **no** form of this command used at the group level reverts to the default value.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no type — Type of neighbor is derived on the local AS specified.

Parameters

internal

Configures the peer as internal.

external

Configures the peer as external.

ttl-security

Syntax

```
ttl-security min-ttl-value
```

```
no ttl-security
```

Context

```
config>service>vprn>bgp>group
```

```
config>service>vprn>bgp>group>neighbor
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures TTL security parameters for incoming packets.

Parameters

min-ttl-value

Specifies the minimum TTL value for an incoming BGP packet.

Values 1 to 255

Default 1

8.4.2.1.14 OSPF commands

```
ospf
```

Syntax

[no] ospf

Context

config>service>vpn

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables access to the context to enable an OSPF protocol instance.

When an OSPF instance is created, the protocol is enabled. To start or suspend execution of the OSPF protocol without affecting the configuration, use the **no shutdown** command.

The **no** form of this command deletes the OSPF protocol instance removing all associated configuration parameters.

Default

no ospf

```
area
```

Syntax

[no] area *area-id*

Context

config>service>vpn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates the context to configure an OSPF area. An area is a collection of network segments within an AS that have been administratively grouped together. The area ID can be specified in dotted-decimal notation or as a 32-bit decimal integer.

The **no** form of this command deletes the specified area from the configuration. Deleting the area also removes the OSPF configuration of all the interfaces, virtual-links, sham-links, address-ranges, and so on that are currently assigned to this area.

Default

no area

Parameters

area-id

Specifies the OSPF area ID expressed in dotted-decimal notation or as a 32-bit decimal integer.

Values a.b.c.d (dotted-decimal)
 0 to 4294967295 (decimal integer)

area-range

Syntax

area-range *ip-prefix/prefix-length* [**advertise** | **not-advertise**]

no area-range *ip-prefix/mask*

no area-range *ip-prefix/mask*

Context

config>service>vprn>ospf>area

ospf>service>vprn>nssa

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 into other areas. Multiple range commands may be used to summarize or hide different ranges. In the case of overlapping ranges, the most specific range command applies.

ABRs send summary link advertisements to describe routes to other areas. To minimize the number of advertisements that are flooded, you can summarize a range of IP addresses and send reachability information about these addresses in an LSA.

The **no** form of this command deletes the range (non) advertisement.

Default

no area-range

Special Cases

NSSA Context

In the NSSA context, the option specifies that the range applies to external routes (via type-7 LSAs) learned within the NSSA when the routes are advertised to other areas as type-5 LSAs.

Area Context

If this command is not entered under the NSSA context, the range applies to summary LSAs even if the area is an NSSA.

Parameters

ipv6-prefix/prefix-length

Specifies the IP prefix in dotted-decimal notation for the range used by the ABR to advertise that summarizes the area into another area.

Values ipv6-prefix - x:x:x:x:x:x:x (eight 16-bit pieces)
 x:x:x:x:x:d.d.d.d
 x [0..FFFF]H
 d [0..255]D
 prefix-length - [0..128]

mask

Specifies the subnet mask for the range expressed as a decimal integer mask length or in dotted-decimal notation.

Values 0 to 32 (mask length), a.b.c.d (dotted-decimal)

advertise | not-advertise

Specifies whether to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range is advertised, and the keyword **not-advertise** indicates the range is not advertised.

The default is **advertise**.

blackhole-aggregate

Syntax

[no] blackhole-aggregate

Context

```
config>service>vpn>ospf>area
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command installs a low priority blackhole route for the entire aggregate. Existing routes that make up the aggregate have a higher priority and only the components of the range for which no route exists are blackholed.

It is possible that when performing area aggregation, addresses may be included in the range for which no actual route exists. This can cause routing loops. To avoid this problem configure the blackhole aggregate option.

The **no** form of this command removes this option.

Default

blackhole-aggregate

interface

Syntax

```
[no] interface ip-int-name [secondary]
```

Context

```
config>service>vpn>ospf>area
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates a context to configure an OSPF interface.

By default interfaces are not activated in any interior gateway protocol such as OSPF unless explicitly configured.

The **no** form of this command deletes the OSPF interface configuration for this interface. The **shutdown** command in the **config>router>ospf>interface** context can be used to disable an interface without removing the configuration for the interface.

Default

no interface

Parameters

ip-int-name

Specifies 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 composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), 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 is returned.

If the IP interface exists in a different area it is moved to this area.

secondary

Allows multiple secondary adjacencies to be established over a single IP interface.

sham-link

Syntax

sham-link *ip-int-name* *ip-address*

Context

config>service>vprn>ospf>area

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command is similar to a virtual link with the exception that metric must be included to distinguish the cost between the MPLS-VP RN link and the backdoor.

Parameters

ip-int-name

Specifies the local interface name used for the sham-link. This is a mandatory parameter and interface names must be unique within the group of defined IP interfaces for **config>router>interface**, **config>service>ies>interface** and **config>service>vprn>interface** commands. An interface name cannot be in the form of an IP address. Interface names can be any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters, the entire string must be enclosed within double quotes. If the IP interface name does not exist or does not have an IP address configured, an error message is returned.

ip-address

Specifies the IP address of the SHAM-link neighbor in IP address dotted-decimal notation. This parameter is the remote peer of the sham link IP address used to set up the SHAM link. This is a mandatory parameter and must be a valid IP address.

advertise-subnet

Syntax

[no] advertise-subnet

Context

config>service>vpn>ospf>area>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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.



Note:

This command is not valid in the OSPF3 context.

The **no** form of this command disables advertising point-to-point interfaces as subnet routes meaning they are advertised as host routes.

Default

advertise-subnet

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

config>service>vpn>ospf>area>if

config>service>vpn>ospf>area>virtual-link

config>service>vpn>ospf>area>sham-link

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the password used by the OSPF interface or virtual-link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.

Note that this command is not valid in the OSPF3 context.

All neighboring routers must use the same type of authentication and password for correct protocol communication. If the **authentication-type** is configured as password, this key must be configured.

The **no** form of this command removes the authentication key.

Default

no authentication-key

Parameters

authentication-key

Specifies the authentication key. The key can be any combination of ASCII characters up to 8 characters (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

hash-key

Specifies the hash key. The key can be any combination of ASCII characters up to 22 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies 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 the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

authentication-type

Syntax

authentication-type {password | message-digest}

no authentication-type

Context

config>service>vprn>ospf>area>if

config>service>vprn>ospf>area>virtual-link

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables authentication and specifies the type of authentication to be used on the OSPF interface, virtual-link, and sham-link.

Note that this command is not valid in the OSPF3 context.

Both simple **password** and **message-digest** authentication are supported.

The **no** form of this command disables authentication on the interface.

Default

no authentication

Parameters

password

Enables simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple **password** authentication is enabled.

message-digest

Enables message digest MD5 authentication in accordance with RFC1321. 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>if

config>service>vprn>ospf>area>virtual-link

config>service>vprn>ospf>area>sham-link

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the use of bidirectional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a specific protocol interface, the state of the protocol interface is

ties to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

**Note:**

- BFD is not supported for IPv6 interfaces.
- For more information about the protocols and platforms that support BFD, see the *7210 SAS-Mxp, R6, R12, S, Sx, T Router Configuration Guide*.

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

Default

no bfd-enable

Parameters***remain-down-on-failure***

Forces adjacency down on BFD failure.

dead-interval

Syntax

dead-interval *seconds*

no dead-interval

Context

config>service>vpn>ospf>area>if

config>service>vpn>ospf>area>virtual-link

config>service>vpn>ospf>area>sham-link

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the time, in seconds, that OSPF waits before declaring a neighbor router down. If no hello packets are received from a neighbor for the duration of the dead interval, the router is assumed to be down. The minimum interval must be two times the hello interval.

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

Default

40

Special Cases

OSPF Interface

If the **dead-interval** configured applies to an interface, all nodes on the subnet must have the same dead interval.

Virtual Link

If the **dead-interval** configured applies to a virtual link, the interval on both termination points 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

Specifies the dead interval, in seconds, expressed as a decimal integer.

Values 2 to 2147483647

hello-interval

Syntax

hello-interval *seconds*

no hello-interval

Context

config>service>vprn>ospf>area>if

config>service>vprn>ospf>area>virtual-link

config>service>vprn>ospf>area>sham-link

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the interval between OSPF hellos issued on the interface, virtual link, or sham-link.

The hello interval, in combination with the dead-interval, is used to establish and maintain the adjacency. Use this parameter to edit the frequency that hello packets are sent.

Reducing the interval, in combination with an appropriate reduction in the associated **dead-interval**, allows for faster detection of link and/or router failures at the cost of higher processing costs.

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

Default

hello-interval 10

Special Cases

OSPF Interface

If the **hello-interval** configured applies to an interface, all nodes on the subnet must have the same hello interval.

Virtual Link

If the **hello-interval** configured applies to a virtual link, the interval on both termination points 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

Specifies 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>vpn>ospf>area>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 if the Ethernet link provided the link is used as a point-to-point.

If the interface type is not known at the time the interface is added to OSPF and subsequently the IP interface is bound (or moved) to a different interface type, this command must be entered manually.

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

Default

point-to-point — If the physical interface is SONET.

broadcast — If the physical interface is Ethernet or unknown.

Special Cases

Virtual-Link

A virtual link is always regarded as a point-to-point interface and not configurable.

Parameters

broadcast

Configures the interface to maintain this link as a broadcast network. To significantly improve adjacency forming and network convergence, a network should be configured as point-to-point if only two routers are connected, even if the network is a broadcast media such as Ethernet.

point-to-point

Configures the interface to maintain this link as a point-to-point link.

message-digest-key

Syntax

message-digest-key *keyid* **md5** [*key* | *hash-key*] [*hash*]

no message-digest-key *keyid*

Context

config>service>vpn>ospf>area>if

config>service>vpn>ospf>area>virtual-link

config>service>vpn>ospf>area>sham-link

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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.

Note that this command is not valid in the OSPF3 context.

The **no** form of this command removes the message digest key identified by the *key-id*.

Parameters

keyid

Specifies the *keyid* expressed as a decimal integer.

Values 1 to 255

md5 key

Specifies the MD5 key. The *key* can be any alphanumeric string up to 16 characters.

md5 hash-key

Specifies the MD5 hash key. The key can be any combination of ASCII characters up to 32 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Specifies 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.

metric

Syntax

metric *metric*

no metric

Context

config>service>vpn>ospf>area>if

config>service>vpn>ospf>area>sham-link

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures an explicit route cost metric for the OSPF interface that overrides the metrics calculated based on the speed of the underlying link.

The **no** form of this command deletes the manually configured interface metric, so the interface uses the computed metric based on the **reference-bandwidth** command setting and the speed of the underlying link.

Default

no metric

Parameters

metric

Specifies the metric to be applied to the interface expressed as a decimal integer.

Values 1 to 65535

mtu

Syntax

mtu *bytes*

no mtu

Context

config>service>vprn>ospf>area>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the OSPF packet size used on this interface. If this parameter is not configured OSPF derives the MTU value from the MTU configured (default or explicitly) in the **config>port>ethernet** context.

If this parameter is configured, the smaller value between the value configured here and the MTU configured (default or explicitly) in a previously-mentioned context is used.

To determine the actual packet size add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured in this command.

The **no** form of this command reverts to the default, which uses the value derived from the MTU configured in the **config>port** context.

Default

no mtu

Parameters

bytes

Specifies the MTU to be used by OSPF for this logical interface in bytes.

Values 512 to 9198 (9212-14) (Depends on the physical media)

passive

Syntax

[no] passive

Context

config>service>vprn>ospf>area>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command adds the passive property to the OSPF interface where passive interfaces are advertised as OSPF interfaces but do not run the OSPF protocol.

By default, only interface addresses that are configured for OSPF are advertised as OSPF interfaces. The **passive** parameter allows an interface to be advertised as an OSPF interface without running the OSPF protocol.

While in passive mode, the interface ignores ingress OSPF protocol packets and does not transmit any OSPF protocol packets.

By default, service interfaces defined in the **config>router>service-prefix** context are passive. All other interfaces are not passive.

The **no** form of this command removes the passive property from the OSPF interface.

priority

Syntax

priority *number*

no priority

Context

config>service>vpn>ospf>area>if

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the priority of the OSPF interface that is used an election of the designated router on the subnet.

This command is only used if the interface is of type broadcast. The router with the highest priority interface becomes the designated router. A router with priority 0 is not eligible to be Designated Router or Backup Designated Router.

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

Default

priority 1

Parameters

number

Specifies the interface priority expressed as a decimal integer. A value of 0 indicates the router is not eligible to be the Designated Router of Backup Designated Router on the interface subnet.

Values 0 to 255

retransmit-interval

Syntax

retransmit-interval *seconds*

no retransmit-interval

Context

config>service>vprn>ospf>area>if

config>service>vprn>ospf>area>virtual-link

config>service>vprn>ospf>area>sham-link

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies the length of time, in seconds, that OSPF waits before retransmitting an unacknowledged link state advertisement (LSA) to an OSPF neighbor.

The value should be longer than the expected round trip delay between any two routers on the attached network. When the retransmit-interval expires and no acknowledgment has been received, the LSA is retransmitted.

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

Default

retransmit-interval 5

Parameters

seconds

Specifies the retransmit interval in seconds expressed as a decimal integer.

Values 1 to 3600

transit-delay

Syntax

transit-delay *seconds*

no transit-delay

Context

config>service>vpn>ospf>area>if

config>service>vpn>ospf>area>virtual-link

config>service>vpn>ospf>area>sham-link

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the estimated time, in seconds, that it takes to transmit a link state advertisement (LSA) on the interface or virtual link or sham-link.

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

Default

transit-delay 1

Parameters

seconds

Specifies the transit delay in seconds expressed as a decimal integer.

Values 0 to 3600

nssa

Syntax

[no] nssa

Context

config>service>vpn>ospf>area

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command creates the context to configure an OSPF Not So Stubby Area (NSSA) and adds/removes the NSSA designation from the area.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF domain.

Existing virtual links of a non-stub or NSSA area 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 this command removes the NSSA designation and configuration context from the area.

Default

no nssa

originate-default-route

Syntax

originate-default-route [type-7]

no originate-default-route

Context

config>service>vpn>ospf>area>nssa

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the generation of a default route and its LSA type (3 or 7) into a Not So Stubby Area (NSSA) by an NSSA Area Border Router (ABR).

When configuring an NSSA with no summaries, the ABR injects a type 3 LSA default route into the NSSA area. Some older implementations expect a type 7 LSA default route.

The **no** form of this command disables origination of a default route.

Default

no originate-default-route

Parameters

type-7

Specifies a type 7 LSA should be used for the default route.

Configure this parameter to inject a type-7 LSA default route instead the type 3 LSA into the NSSA configured with no summaries. To revert to a type 3 LSA, enter **originate-default-route** without the **type-7** parameter.

Default Type 3 LSA for the default route.

redistribute-external

Syntax

[no] redistribute-external

Context

config>service>vprn>ospf>area>nssa

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables the redistribution of external routes into the Not So Stubby Area (NSSA) or an NSSA area border router (ABR) that is exporting the routes into non-NSSA areas.

NSSA are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an Area Border Router to the entire OSPF domain.

The **no** form of this command disables the default behavior to automatically redistribute external routes into the NSSA area from the NSSA ABR.

Default

redistribute-external

summaries

Syntax

[no] summaries

Context

config>service>vprn>ospf>area>nssa

config>service>vprn>ospf>area>stub

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables sending summary (type 3) advertisements into a stub area or Not So Stubby Area (NSSA) on an Area Border Router (ABR). This parameter is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub or nssa area. By default, summary route advertisements are sent into the stub area or NSSA.

The **no** form of this command disables sending summary route advertisements and, for stub areas, only the default route is advertised by the ABR.

Default

summaries

stub

Syntax

[no] stub

Context

config>service>vprn>ospf>area

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables access to the context to configure an OSPF stub area and adds/removes the stub designation from the area. External routing information is not flooded into stub areas. All routers in the stub area must be configured with the **stub** command. An OSPF area cannot be both an NSSA and a stub area. Existing virtual links of a non-stub area or NSSA are removed when its designation is changed to NSSA or STUB.

By default, an area is not a stub area.

The **no** form of this 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
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the metric used by the area border router (ABR) for the default route into a stub area. The default metric should only be configured on an ABR of a stub area. An ABR generates a default route if the area is a **stub** area.

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

Default

```
default-metric 1
```

Parameters

metric

Specifies the metric expressed as a decimal integer for the default route cost to be advertised into the stub area.

Values 1 to 16777215

virtual-link

Syntax

```
[no] virtual-link router-id transit-area area-id
```

Context

```
config>service>vpn>ospf>area
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures a virtual link to connect area border routers to the backbone via a virtual link. The backbone area (area 0.0.0.0) must be contiguous and all other areas must be connected to the backbone area. If it is not practical to connect an area to the backbone (see area 0.0.0.2 in [Figure 98: OSPF areas](#)), the area border routers (routers 1 and 2 in [Figure 98: OSPF areas](#)) must be connected via a virtual link. The two area border routers form a point-to-point like adjacency across the transit area (area 0.0.0.1 in [Figure 98: OSPF areas](#)). A virtual link can only be configured while in the area 0.0.0.0 context.

The *router-id* specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or a Not So Stubby Area (NSSA).

The **no** form of this command deletes the virtual link.

Parameters

router-id


Specifies the router ID of the virtual neighbor in IP address dotted-decimal notation.

transit-area area-id

Specifies the area-id specified identifies the transit area that links the backbone area with the area that has no physical connection with the backbone.

The OSPF backbone area, area 0.0.0.0, must be contiguous and all other areas must be connected to the backbone area. The backbone distributes routing information between areas. If it is not practical to connect an area to the backbone (see Area 0.0.0.5 in the following figure), the area border routers (such as routers Y and Z) must be connected via a virtual link. The two area border routers form a point-to-point-like adjacency across the transit area (see Area 0.0.0.4).

Figure 98: OSPF areas



compatible-rfc1583

Syntax

[no] **compatible-rfc1583**

Context

config>service>vpn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables OSPF summary and external route calculations in compliance with RFC1583 and earlier RFCs.

RFC1583 and earlier RFCs use a different method to calculate summary and external route costs. To avoid routing loops, all routers in an OSPF domain should perform the same calculation method.

Although it would be favorable to require all routers to run a more current compliance level, this command allows the router to use obsolete methods of calculation.

The **no** form of this command enables the post-RFC1583 method of summary and external route calculation.

Default

compatible-rfc1583

export

Syntax

export *policy-name* [*policy-name*...]

no export

Context

config>service>vprn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command associates export route policies to determine which routes are exported from the route table to OSPF. Export policies are only in effect if OSPF is configured as an ASBR.

If no export policy is specified, non-OSPF routes are not exported from the routing table manager to OSPF.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered overrides the previous command. A maximum of five policy names can be specified.

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

Default

no export

Parameters

policy-name

Specifies the export route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.

The specified names must already be defined.

external-db-overflow

Syntax

external-db-overflow *limit interval*

no external-db-overflow

Context

config>service>vprn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables limits on the number of non-default AS-external-LSA entries that can be stored in the LSDB and specifies a wait timer before processing these after the limit is exceeded.

The *limit* value specifies the maximum number of non-default AS-external-LSA entries that can be stored in the link-state database (LSDB). Placing a limit on the non-default AS-external-LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reach or exceed the *limit*, the table is in an overflow state. When in an overflow state, the router does not originate any new AS-external-LSAs. In fact, it withdraws all the self-originated non-default external LSAs.

The *interval* specifies the amount of time to wait after an overflow state before regenerating and processing non-default AS-external-LSAs. The waiting period acts like a dampening period preventing the router from continuously running Shortest Path First (SPF) calculations caused by the excessive number of non-default AS-external LSAs.

The **external-db-overflow** must be set identically on all routers attached to any regular OSPF area. OSPF stub areas and not-so-stubby areas (NSSAs) are excluded.

The **no** form of this command disables limiting the number of non-default AS-external-LSA entries.

Default

no external-db-overflow

Parameters

limit

Specifies 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

interval

Specifies 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

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the preference for OSPF external routes.

A route can be learned by the router from different protocols in which case the costs are not comparable; when this occurs the preference is used to decide which route is used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of what route to use is determined by the configuration of the **config>router>ecmp** command.

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

Default

external-preference 150

Parameters

preference
Specifies the preference for external routes expressed as a decimal integer.

Route Type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes

Route Type	Preference	Configurable
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes
IS-IS level 2 external	165	Yes
BGP	170	Yes

Note:

1. Preference for OSPF internal routes is configured with the **preference** command.

Values 1 to 255

ignore-dn-bit

Syntax

[no] **ignore-dn-bit**

Context

config>service>vpn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies whether to ignore the DN bit for OSPF LSA packets for this instance of OSPF on the router. When enabled, the DN bit for OSPF LSA packets are ignored. When disabled, the DN bit is not ignored for OSPF LSA packets.

import

Syntax

import *policy-name* [*policy-name...*(up to 5 max)]

no import

Context

config>service>vpn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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.

This configuration can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific level is used.

When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be specified. The first policy that matches is applied.

When multiple import commands are issued, the last command entered overrides the previous command.

The **no** form of this command removes the policy association. To remove the association of all policies, use `no import` without any arguments.

Default

`no import`

Parameters

policy-name

Specifies the route policy name. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

overload

Syntax

overload [*timeout seconds*]

no overload

Context

`config>service>vpn>ospf`

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command changes the overload state of the local router so that it appears to be overloaded. When overload is enabled, the router can participate in OSPF routing, but is not used for transit traffic. Traffic destined to directly attached interfaces continue to reach the router.

To put the IGP in an overload state enter a timeout value. The IGP enters the overload state until the timeout timer expires or a **no overload** command is executed.

If the **overload** command is encountered during the execution of an [overload-on-boot](#) command, this command takes precedence. This could occur as a result of a saved configuration file where both parameters are saved. When the file is saved by the system the **overload-on-boot** command is saved after the **overload** command.

The **no** form of this command returns to the default. When the **no overload** command is executed, the overload state is terminated regardless the reason the protocol entered overload state.

Default

no overload

Parameters

timeout seconds

Specifies the number of seconds to reset overloading.

Values 60 to 1800

Default 60

overload-include-stub

Syntax

[no] overload-include-stub

Context

config>service>vpn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures whether the OSPF stub networks should be advertised with a maximum metric value when the system goes into overload state for any reason. When enabled, the system uses the maximum metric value. When this command is enabled and the router is in overload, all stub interfaces, including loopback and system interfaces, are advertised at the maximum metric.

Default

no overload-include-stub

overload-on-boot

Syntax

overload-on-boot [*timeout seconds*]

no overload

Context

config>service>vpn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the IGP upon bootup in the overload state until one of the following events occur:

- The timeout timer expires.
- A manual override of the current overload state is entered with the **no overload** command.

When the router is in an overload state, the router is used only if there is no other router to reach the destination.

The **no overload** command does not affect the **overload-on-boot** function.

The **no** form of this command removes the overload-on-boot functionality from the configuration.

Default

no overload-on-boot

Parameters

timeout seconds

Specifies the number of seconds to reset overloading.

Values 60 to 1800

Default 60

preference

Syntax

preference *preference*

no preference

Context

```
config>service>vpn>ospf
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the preference for OSPF internal routes.

A route can be learned by the router from different protocols in which case the costs are not comparable, when this occurs the preference is used to decide to which route is used.

Different protocols should not be configured with the same preference, if this occurs the tiebreaker is per the default preference table as defined in the following table. If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

If multiple routes are learned with an identical preference using the same protocol and the costs (metrics) are equal, the decision of what route to use is determined by the configuration of the **config router ecmp** command.

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

Default

```
preference 10
```

Parameters

preference

Specifies the preference for internal routes expressed as a decimal integer. The following table lists the defaults for different route types.

Table 109: Route type preference defaults

Route type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes ¹⁷
IS-IS level 1 internal	15	Yes
IS-IS level 2 internal	18	Yes
RIP	100	Yes
OSPF external	150	Yes
IS-IS level 1 external	160	Yes

¹⁷ Preference for OSPF internal routes is configured with the **preference** command.

Route type	Preference	Configurable
IS-IS level 2 external	165	Yes
BGP	170	Yes

Values 1 to 255

reference-bandwidth

Syntax

reference-bandwidth *reference-bandwidth*

no reference-bandwidth

Context

config>service>vpn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command configures the reference bandwidth in kilobits per second (Kbps) that provides the reference for the default costing of interfaces based on their underlying link speed.

The default interface cost is calculated as follows:

cost = reference-bandwidth / bandwidth

The default *reference-bandwidth* is 100,000,000 Kbps or 100 Gbps, so the default auto-cost metrics for various link speeds are as follows:

- 10 Mbs link default cost of 10000
- 100 Mbs link default cost of 1000
- 1 Gbps link default cost of 100
- 10 Gbps link default cost of 10

The **reference-bandwidth** command assigns a default cost to the interface based on the interface speed. To override this default cost on a particular interface, use the **metric** *metric* command in the **config>router>ospf>area>interface** *ip-int-name* context.

The **no** form of this command reverts the reference-bandwidth to the default value.

Default

reference-bandwidth 100000000

Parameters

reference-bandwidth

Specifies the reference bandwidth in kilobits per second expressed as a decimal integer.

Values 1 to 1000000000

super-backbone

Syntax

[no] super-backbone

Context

config>service>vpn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies whether CE-PE functionality is required or not. The OSPF super backbone indicates the type of the LSA generated as a result of routes redistributed into OSPF. When enabled, the redistributed routes are injected as summary, external or NSSA LSAs. When disabled, the redistributed routes are injected as either external or NSSA LSAs only.

Default

no super-backbone

suppress-dn-bit

Syntax

[no] suppress-dn-bit

Context

config>service>vpn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies whether to suppress the setting of the DN bit for OSPF LSA packets generated by this instance of OSPF on the router. When enabled, the DN bit for OSPF LSA packets generated by this

instance of the OSPF router is not set. When disabled, this instance of the OSPF router follows the normal procedure to determine whether to set the DN bit.

Default

no suppress-dn-bit

timers

Syntax

timers

Context

config>service>vpn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

Commands in this context configure of OSPF timers. Timers control the delay between receipt of a link state advertisement (LSA) requiring a Dijkstra (Shortest Path First (SPF)) calculation and the minimum time between successive SPF calculations.

Changing the timers affect CPU utilization and network reconvergence times. Lower values reduce convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase reconvergence time.

spf-wait

Syntax

spf-wait *max-spf-wait* [*spf-initial-wait* [*spf-second-wait*]]

no spf-wait

Context

config>service>vpn>ospf>timers

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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) occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, the next SPF runs after 2000 milliseconds, and the next SPF runs after 4000 milliseconds, and so on, until it reaches the **spf-wait** value. The SPF interval stays 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 drops back to *spf-initial-wait*.

The timer must be entered in increments of 100 milliseconds. Values entered that do not match this requirement are rejected.

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

Default

no spf-wait

Parameters

max-spf-wait

Specifies the maximum interval in milliseconds between two consecutive SPF calculations.

Values 1 to 120000

Default 1000

spf-initial-wait

Specifies the initial SPF calculation delay in milliseconds after a topology change.

Values 10 to 100000

Default 1000

spf-second-wait

Specifies the hold time in milliseconds between the first and second SPF calculation.

Values 10 to 100000

Default 1000

vpn-domain

Syntax

vpn-domain [*type* {0005 | 0105 | 0205 | 8005}] *id id*

no vpn-domain

Context

config>service>vpn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID. This applies to VPRN instances of OSPF only. An attempt to modify the value of this object results in an inconsistent value error when it is not a VPRN instance. The parameters are mandatory and can be entered in either order.

Default

no vpn-domain

Parameters

id

Specifies the OSPF VPN domain in the "xxxx.xxxx.xxxx" format. This is exchanged using BGP in the extended community attribute associated with a prefix. This object applies to VPRN instances of OSPF only.

type

Specifies the type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID.

Values 0005, 0105, 0205, 8005

vpn-tag

Syntax

vpn-tag *vpn-tag*

no vpn-tag

Context

config>service>vprn>ospf

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command specifies the route tag for an OSPF VPN on a PE router. This field is set in the tag field of the OSPF external LSAs generated by the PE. This is mainly used to prevent routing loops. This applies to VPRN instances of OSPF only. An attempt to modify the value of this object results in an inconsistent value error when it is not a VPRN instance.

Default

vpn-tag 0

lsa-arrival

Syntax

lsa-arrival *lsa-arrival-time*

no lsa-arrival

Context

config>service>vpn>ospf>timers

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This parameter defines the minimum delay that must pass between receipt of the same Link State Advertisements (LSAs) arriving from neighbors. It is recommended that the neighbors configured **lsa-generate lsa-second-wait** interval is equal or greater than the **lsa-arrival** timer configured here.

Use the **no** form of this command to return to the default.

Default

no lsa-arrival

Parameters

lsa-arrival-time

Specifies the timer in milliseconds. Values entered that do not match this requirement are rejected.

Values 0 to 600000

lsa-generate

Syntax

lsa-generate *max-lsa-wait* [*lsa-initial-wait* [*lsa-second-wait*]]

no lsa-generate-interval

Context

config>service>vpn>ospf>timers

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This parameter customizes the throttling of OSPF LSA-generation. Timers that determine when to generate the first, second, and subsequent LSAs can be controlled with this command. Subsequent LSAs are generated at increasing intervals of the *lsa-second-wait* timer until a maximum value is reached. Configuring the **lsa-arrival** interval to equal or less than the **lsa-second-wait** interval configured in the **lsa-generate** command is recommended.

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

Default

no lsa-generate

Parameters

max-lsa-wait

Specifies the maximum interval, in milliseconds, between two consecutive occurrences of an LSA being generated.

The timer must be entered as either 1 or in millisecond increments. Values entered that do not match this requirement are rejected.

Values 1 to 600000

8.4.2.2 Show commands

egress-label

Syntax

egress-label *start-label* [*end-label*]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command display services using the range of egress labels.

If only the mandatory *start-label* parameter is specified, only services using the specified label are displayed.

If both *start-label* and *end-label* parameters are specified, the services using the range of labels X where *start-label* <= X <= *end-label* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters

start-label

Specifies the starting egress label value for which to display services using the label range. If only *egress-label/1* is specified, services only using *egress-label/1* are displayed.

Values 0 | 2048 to 131071

end-label

Specifies the ending egress label value for which to display services using the label range.

Values 2049 to 131071

Default The *egress-label/1* value.

Output

The following output is an example of service egress label information, and [Table 110: Output fields: egress label](#) describes the output fields.

Sample output

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

Table 110: Output fields: egress label

Label	Description
Svc Id	The ID that identifies a service.
Sdp Id	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 egress label range.

ingress-label

Syntax

ingress-label *start-label* [*end-label*]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 the range of labels X where *start-label* <= X <= *end-label* are displayed.

Use the **show router vprn-service-id ldp bindings** command to display dynamic labels.

Parameters

start-label

Specifies 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

Specifies the ending ingress label value for which to display services using the label range.

Values 2048 to 131071

Default The *start-label* value.

Output

The following output is an example of service ingress label information, and [Table 111: Output fields: ingress label](#) describes the output fields.

Sample output

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

Table 111: Output fields: ingress label

Label	Description
Svc ID	The service identifier.
SDP Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
I.Lbl	The ingress label used by the far-end device to send packets to this device in this service by the SDP.
E.Lbl	The egress label used by this device to send packets to the far-end device in this service by the SDP.
Number of Bindings Found	The number of SDP bindings within the label range specified.

sap-using

Syntax

sap-using [**sap** *sap-id*]
sap-using interface [*ip-address* | *ip-int-name*]
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**ingress** | **egress**] **qos-policy** *qos-policy-id*

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode.

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

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

interface

Specifies matching SAPs with the specified IP interface. This parameter can be used on 7210 SAS platforms operating in access-uplink mode only if the specified interface is configured within an IES context.

ip-address

Specifies the IP address of the interface for which to display matching SAPs. This parameter can be used on 7210 SAS platforms operating in access-uplink mode only if the specified interface is configured within an IES context.

Values a.b.c.d

ip-int-name

Specifies the IP interface name for which to display matching SAPs. This parameter can be used on 7210 SAS platforms operating in access-uplink mode only if the specified interface is configured within an IES context.

ingress

Specifies matching an ingress policy.

egress

Specifies matching an egress policy.

qos-policy qos-policy-id

Specifies the ingress or egress QoS Policy ID for which to display matching SAPs.

Values 1 to 65535

filter filter-id

Specifies the ingress or egress filter policy ID for which to display matching SAPs.

Values 1 to 65535

Output

The following output is an example of service SAP information, and [Table 112: Output fields: SAP-using](#) describes the output fields.

Sample output

```
*A:ALA-12# show service sap-using sap 1/1
=====
Service Access Points
=====
PortId          SvcId      SapMTU  I.QoS  I.Mac/IP  E.QoS  E.Mac/IP  A.Pol  Adm  Opr
-----
1/1/7:0         1          1518    10     8         10     none     none   Up   Up
1/1/11:0        100         1514    1     none      1     none     none   Down Down
1/1/7:300       300         1518    10     none      10     none     1000   Up   Up
-----
Number of SAPs : 3
-----
*A:ALA-12#
```

Table 112: Output fields: SAP-using

Label	Description
Port ID	The ID of the access port where the SAP is defined.
Svc ID	The service identifier.
SapMTU	The SAP MTU value.
I.QoS	The SAP ingress QoS policy number specified on the ingress SAP.
I.MAC/IP	The MAC or IP filter policy ID applied to the ingress SAP.
E.QoS	The SAP egress QoS policy number specified on the egress SAP.
E.Mac/IP	The MAC or IP filter policy ID applied to the egress SAP
A.Pol	The accounting policy ID assigned to the SAP.
Adm	The desired state of the SAP.

Label	Description
Opr	The actual state of the SAP.

sdp

Syntax

sdp [*sdp-id* | *far-end ip-address*] [**detail** | **keep-alive-history**]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays SDP information.

If no optional parameters are specified, a summary SDP output for all SDPs is displayed.

Parameters

sdp-id

Specifies the SDP ID for which to display information.

Values 1 to 17407

Default All SDPs.

far-end ip-address

Displays only SDPs matching with the specified far-end IP address.

Default SDPs with any far-end IP address.

detail

Displays detailed SDP information.

Default SDP summary output.

keep-alive-history

Displays the last fifty SDP keepalive events for the SDP.

Default SDP summary output.

Output

The following output is an example of SDP information, and [Table 113: Output fields: SDP](#) describes the output fields.

Sample output

```

*A:ALA-12# show service sdp
=====
Services: Service Destination Points
=====
SdpId      Adm MTU    Opr MTU    IP address      Adm  Opr      Deliver Signal
-----
10         4462      4462      10.20.1.3       Up   Dn NotReady MPLS    TLDP
40         4462      1534      10.20.1.20      Up   Up        MPLS    TLDP
-----
Number of SDPs : 5
=====
*A:ALA-12#

*A:ALA-12# show service sdp 8
=====
Service Destination Point (Sdp Id : 8)
=====
SdpId      Adm MTU    Opr MTU    IP address      Adm  Opr      Deliver Signal
-----
8          4462      4462      10.10.10.104    Up   Dn NotReady MPLS    TLDP
=====
Service Destination Point (Sdp Id : 8) Details
-----
Sdp Id 8  -(10.10.10.104)
-----
Description          : MPLS-10.10.10.104
SDP Id               : 8
Admin Path MTU       : 0
Far End              : 10.10.10.104
Admin State          : Up
Flags                : SignalingSessDown TransportTunnDown
Signaling             : TLDP
Last Status Change   : 02/01/2007 09:11:39
Last Mgmt Change     : 02/01/2007 09:11:46
VLAN VC Etype        : 0x8100
Adv. MTU Over.       : No

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

Associated LSP LIST :
Lsp Name              : to-104
Admin State           : Up
Oper State            : Down
Time Since Last Tran* : 01d07h36m
=====
* indicates that the corresponding row element may have been truncated.
*A:ALA-12#

```

Table 113: Output fields: SDP

Label	Description
SDP Id	The SDP identifier.
Adm MTU	Specifies the largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.

Label	Description
Opr MTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
IP address	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Adm Admin State	Specifies the state of the SDP.
Opr Oper State	Specifies the operating state of the SDP.
Flags	Specifies all the conditions that affect the operating status of this SDP.
Signal Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
Last Status Change	Specifies the time of the most recent operating status change to this SDP.
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SDP.
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Deliver Delivered	Specifies the type of delivery used by the SDP: MPLS.
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hello Timeout	Specifies the number of seconds to wait for an SDP echo response message before declaring a timeout.
Unmatched Replies	Specifies the number of SDP unmatched message replies.

Label	Description
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
TX Hello Msgs	Specifies the number of SDP echo request messages transmitted since the keepalive was administratively enabled or the counter was cleared.
Rx Hello Msgs	Specifies the number of SDP echo request messages received since the keepalive was administratively enabled or the counter was cleared.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field.

sdp-using

Syntax

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

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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

Specifies the virtual circuit identifier.

Values 1 to 4294967295

far-end ip-address

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

Default Services with any far-end IP address.

Output

The following output is an example of service SDP information, and [Table 114: Output Fields: SDP-using](#) describes the output fields.

Sample output

```
*A:ALA-1# show service sdp-using 300
=====
Service Destination Point (Sdp Id : 300)
=====
SvcId      SdpId      Type Far End      Opr State I.Label  E.Label
-----
1          300:1      Mesh 10.0.0.13     Up        131071   131071
2          300:2      Spok 10.0.0.13     Up        131070   131070
100        300:100    Mesh 10.0.0.13     Up        131069   131069
101        300:101    Mesh 10.0.0.13     Up        131068   131068
102        300:102    Mesh 10.0.0.13     Up        131067   131067
-----
Number of SDPs : 5
-----
*A:ALA-1#

A:ALA-48# show service sdp-using
=====
SDP Using
=====
SvcId      SdpId      Type Far End      Opr State I.Label  E.Label
-----
3          2:3        Spok 10.20.1.2     Up        n/a      n/a
103        3:103      Spok 10.20.1.3     Up        131067   131068
103        4:103      Spok 10.20.1.2     Up        131065   131069
105        3:105      Spok 10.20.1.3     Up        131066   131067
-----
Number of SDPs : 4
-----
A:ALA-48
```

Table 114: Output Fields: SDP-using

Label	Description
Svc ID	The service identifier.
Sdp ID	The SDP identifier.
Type	Type of SDP: spoke or mesh.
Far End	The far end address of the SDP.
Oper State	The operational state of the service.

Label	Description
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.

service-using

Syntax

service-using [**epipe**] [**ies**] [**vpls**] [**vprn**][**sdp** *sdp-id*] [*customer customer-id*]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the services matching certain usage properties.

If no optional parameters are specified, all services defined on the system are displayed.

Parameters

epipe

Displays matching Epipe services.

ies

Displays matching IES instances.

vpls

Displays matching VPLS instances.

vprn

Displays matching VPRN services. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

sdp *sdp-id*

Displays only services bound to the specified SDP ID.

Values 1 to 17407

Default Services bound to any SDP ID.

customer *customer-id*

Displays services only associated with the specified customer ID.

Values 1 to 2147483647

Default Services associated with an customer.

Output

The following output is an example of service information, and [Table 115: Output Fields: Service-using](#) describes the output fields.

Sample output

```
*A:ALA-12# show service service-using customer 10
=====
Services
=====
ServiceId   Type      Adm   Opr      CustomerId  Last Mgmt Change
-----
1           VPLS      Up    Up        10          09/05/2006 13:24:15
100        IES       Up    Up        10          09/05/2006 13:24:15
300        Epipe     Up    Up        10          09/05/2006 13:24:15
900        VPRN      Up    Up        2           11/04/2006 04:55:12
-----
Matching Services : 4
=====
*A:ALA-12#

*A:ALA-12# show service service-using epipe
=====
Services [epipe]
=====
ServiceId   Type      Adm   Opr      CustomerId  Last Mgmt Change
-----
6           Epipe     Up    Up        6           06/22/2006 23:05:58
7           Epipe     Up    Up        6           06/22/2006 23:05:58
8           Epipe     Up    Up        3           06/22/2006 23:05:58
103        Epipe     Up    Up        6           06/22/2006 23:05:58
-----
Matching Services : 4
=====
*A:ALA-12#

A:del14# show service service-using
=====
Services
=====
ServiceId   Type      Adm   Opr      CustomerId  Last Mgmt Change
-----
1           uVPLS     Up    Up        1           10/26/2006 15:44:57
2           Epipe     Up    Down      1           10/26/2006 15:44:57
10          mVPLS     Down  Down      1           10/26/2006 15:44:57
11          mVPLS     Down  Down      1           10/26/2006 15:44:57
100        mVPLS     Up    Up        1           10/26/2006 15:44:57
101        mVPLS     Up    Up        1           10/26/2006 15:44:57
102        mVPLS     Up    Up        1           10/26/2006 15:44:57
999        uVPLS     Down  Down      1           10/26/2006 16:14:33
-----
Matching Services : 8
-----
A:del14#
```

Table 115: Output Fields: Service-using

Label	Description
Service Id	The service identifier.
Type	Specifies the service type configured for the service ID.
Adm	The desired state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

id

Syntax

id service-id {all | arp | base | fdb | labels | mfib | sap | sdp | split-horizon-group | stp}

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information for a particular *service-id*.

Parameters

service-id

Specifies the unique service identification number that identifies the service in the service domain.

all

Display detailed information about the service.

arp

Display ARP entries for the service.

base

Display basic service information.

fdb

Display FDB entries.

- interface**
Display service interfaces.
- labels**
Display labels being used by this service.
- sap**
Display SAPs associated to the service.
- sdp**
Display SDPs associated with the service. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.
- split-horizon-group**
Display split horizon group information.
- stp**
Display STP information.

all

Syntax

all

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays detailed information for all aspects of the service.

Output

The following output is an example of detailed service information, and [Table 116: Output fields: service ID All](#) describes the output fields.

Sample output

```
*A:7210SAS>show>service>id# all

=====
Service Detailed Information
=====
Service Id      : 1                Vpn Id      : 0
Service Type    : Epipe
Description     : (Not Specified)
Customer Id     : 1
Last Status Change: 02/12/2002 23:51:07
Last Mgmt Change  : 02/12/2002 23:50:18
Admin State     : Up              Oper State   : Up
```

```

SAP Count      : 2
Uplink Type:   : L2
SAP Type:      : Any
Customer vlan: : n/a
-----
Service Access Points
-----
SAP 1/1/9:600.*
-----
Service Id      : 1
SAP             : 1/1/9:600.*
QinQ Dot1p     : Default
Description     : (Not Specified)
Admin State     : Up
Flags           : None
Last Status Change : 02/12/2002 23:51:06
Last Mgmt Change  : 02/12/2002 23:50:18
Dot1Q Ethertype : 0x8100
QinQ Ethertype  : 0x8100

Admin MTU       : 9212
Ingr IP Fltr-Id : n/a
Ingr Mac Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a
tod-suite       : None
Endpoint        : N/A

Oper MTU        : 9212
Egr IP Fltr-Id  : n/a
Egr Mac Fltr-Id : n/a
Egr IPv6 Fltr-Id : n/a

Acct. Pol       : None
Collect Stats   : Disabled
-----
QoS
-----
Ingress qos-policy : n/a
-----
Aggregate Policer
-----
rate            : n/a
burst           : n/a
-----
Ingress QoS Classifier Usage
-----
Classifiers Allocated: 2
Classifiers Used      : 1
Meters Allocated      : 1
Meters Used           : 1
-----
Sap Statistics
-----
Ingress Stats:      Packets      Octets
                   0             0
Egress Stats:      26941105      18014193523
Extra-Tag Drop Stats: n/a         n/a
-----
SAP 1/1/12:90
-----
Service Id      : 1
SAP             : 1/1/12:90
QinQ Dot1p     : (Not Specified)
Description     : (Not Specified)
Admin State     : Up
Flags           : None
Last Status Change : 02/12/2002 23:51:07
Last Mgmt Change  : 02/13/2002 00:05:46
Dot1Q Ethertype : 0x8100
Loopback Mode    : Internal
QinQ Ethertype  : 0x8100
No-svc-port used : 1/1/25
Encap           : q-tag
Oper State      : Up

```



```

Loopback Src Addr : 00:00:01:00:02:00
Loopback Dst Addr : 00:00:01:00:03:00

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
Endpoint           : N/A

Acct. Pol          : None                      Collect Stats      : Disabled

-----
QoS
-----
Ingress qos-policy : 1
-----
Aggregate Policer
-----
rate              : n/a                      burst             : n/a
-----
Ingress QoS Classifier Usage
-----
Classifiers Allocated: 2                      Meters Allocated  : 1
Classifiers Used     : 1                      Meters Used       : 1
-----
Sap Statistics
-----
Ingress Stats:      Packets      Octets
                   26940595      18013850572
Egress Stats:       0            0
Ingress Drop Stats: 0            0

Extra-Tag Drop Stats: n/a          n/a
-----
Sap per Meter stats (in/out counter mode)
-----
                   Packets      Octets

Ingress Meter 1
For. InProf        : 8          4265
For. OutProf       : 26941156   18014224039
-----
Service Endpoints
-----
No Endpoints found.
=====
*A:7210SAS>show>service>id#

```

Table 116: Output fields: service ID All

Label	Description
Service Detailed Information	
Service Id	The service identifier.
VPN Id	The number which identifies the VPN.

Label	Description
Customer Id	The customer identifier.
Last Status Change	The date and time of the most recent change in the administrative or operating status of the service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Admin State	The current administrative state.
Oper State	The current operational state.
Route Dist.	Displays the route distribution number.
AS Number	Displays the autonomous system number.
Router Id	Displays the router ID for this service.
Auto Bind	Specifies the automatic binding type for the SDP assigned to this service.
Vrf Target	Specifies the VRF target applied to this service.
Vrf Import	Specifies the VRF import policy applied to this service.
Vrf Export	Specifies the VRF export policy applied to this service.
Description	Generic information about the service.
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service.
Split Horizon Group	Name of the split horizon group for this service.
Description	Description of the split horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.

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.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the keepalive protocol.
Oper State	The current status of the keepalive protocol.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field.

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.
Number of SDPs	The total number SDPs applied to this service ID.
Service Access Points	
Service Id	The service identifier.
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Admin State	The desired state of the SAP.
Oper State	The operating state of the SAP.
Last Changed	The date and time of the last change.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The SAP ingress QoS policy ID.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
Spoke SDPs	
Managed by Service	Specifies the service-id of the management VPLS managing this spoke-SDP.
Managed by Spoke	Specifies the sap-id inside the management VPLS managing this spoke-SDP.
Prune state	Specifies the STP state inherited from the management VPLS.
Peer Pw Bits	Indicates the bits set by the LDP peer when there is a fault on its side of the pseudowire. LAC failures occur on the SAP that has been configured on the pipe service, PSN bits are set by SDP-binding failures on the pipe service. The pwNotForwarding bit is set when none of the preceding failures apply, such as an MTU mismatch failure. This value is only applicable if the peer is using the pseudowire status signalling method to indicate faults. pwNotForwarding — Pseudowire not forwarding

Label	Description
	lacIngressFault Local — Attachment circuit RX fault lacEgresssFault Local — Attachment circuit TX fault psnIngressFault Local — PSN-facing PW RX fault psnEgressFault Local — PSN-facing PW TX fault pwFwdingStandby — Pseudowire in standby mode
Max IPv4 Routes	Maximum IPv4 routes configured for use with the service.
Last Changed	The date and time of the most recent management-initiated change.
Dot1Q Ethertype	The Dot1q ethertype in use by the SAP.
Ingr IP Fltr-Id	The policy ID of the IP filter applied at ingress.
Ingr Mac Fltr-Id	The policy ID of the MAC filter applied at ingress.
Egr IP Fltr-Id	The policy ID of the IP filter applied at egress.
Egr Mac Fltr-Id	The policy ID of the MAC filter applied at egress.
tod-suite	The TOD suite applied for use by this SAP.
rate	Specifies the SAP aggregate rate configured for the aggregate policer/meter used by this SAP.
burst	Specifies the burst to be used with SAP aggregate policer/meter used by this SAP.
Classifiers Allocated	Number of SAP ingress QoS resources allocated for use by this SAP.
Classifiers Used	Number of SAP ingress QoS resources in use by this SAP.
Meters Allocated	Number of SAP ingress meter resources allocated for use by this SAP. This is set to half the number of classifiers allocated to this SAP.
Meters Used	Number of SAP ingress meters in use.
Ingress Stats	The number of received packets/octets for this SAP.
Egress Stats	The number of packets/octets forwarded out of this SAP.
Ingress Drop Stats	Number of packets/octets dropped by the system.
Extra-Tag Drop Stats	Number of packets received with the count of VLAN tags exceeding the count of VLAN tags implied by the SAP encapsulation.
Ingress Meter 1	The index of the ingress QoS meter of this SAP.

Label	Description
For. InProf	Number of in-profile packets/octetets received on this SAP.
For. OutProf	Number of out-of-profile packets/octetets received on this SAP.
If Name	IP interface name assigned by user.
Protocols	Protocols enabled for use on this interface.
Oper (v4/v6)	Operational status of this interface for IPv4 and IPv6.
IP Addr/mask	IPv4 address and Mask assigned to this interface.
Address Type	Whether the address is a primary or secondary address.
Broadcast Address	Type of broadcast address used. It can be host-ones or all-ones.
If Index	The interface Index assigned by the system. It is used with SNMP IfTable.
Virt. If Index	The interface index assigned by the system. It is used with SNMP.
Last Oper Chg	Timestamp associated with the last operational change.
Global If Index	This is the system wide Interface index allotted by the system.
If Type	Network — The IP interface is a network/core IP interface. Service — The IP interface is a service IP interface.
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled.
Arp Timeout	Specifies the timeout for an ARP entry learned on the interface.
IP Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the port to the far-end router, without requiring the packet to be fragmented.
LdpSyncTimer	Specifies the value used for IGP-LDP synchronization.
Redirects	Specifies the rate for ICMP redirect messages.
Unreachables	Specifies the rate for ICMP unreachable messages.
TTL Expired	Specifies the rate for ICMP TTL messages.
MAC Address	Specifies the 48-bit IEEE 802.3 MAC address.

arp

Syntax

arp [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*] [**sdp** *sdp-id:vc-id*] [**summary**]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Displays the ARP table for the IES instance.

Parameters

ip-address

Displays only ARP entries in the ARP table with the specified IP address.

Default All IP addresses.

mac ieee-address

Displays only ARP entries in the ARP table with the specified 48-bit MAC address. The MAC address can be expressed in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

Default All MAC addresses.

sap sap-id

Displays SAP information for the specified SAP ID. See [Common CLI command descriptions](#) for command syntax.

sdp-id:vc-id

Displays SDP information for the specified SDP ID and VC ID. This parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

port id

Specifies matching service ARP entries associated with the specified IP interface.

ip-address

Specifies the IP address of the interface for which to display matching ARP entries.

Values a.b.c.d

ip-int-name

Specifies the IP interface name for which to display matching ARPs.

Output

The following output is an example of ARP information, and [Table 117: Output fields: ARP](#) describes the output fields.

Sample output

```
*A:ALA-12# show service id 2 arp
=====
ARP Table
=====
IP Address      MAC Address      Type   Age      Interface      Port
-----
10.11.1.1       00:03:fa:00:08:22 Other   00:00:00   ies-100-10.11.1 1/1/11:0
=====
*A:ALA-12#
```

Table 117: Output fields: ARP

Label	Description
Service ID	The service ID number.
MAC	The specified MAC address
Source-Identifier	The location the MAC is defined.
Type	Static FDB entries created by management.
	Learned Dynamic entries created by the learning process.
	OAM Entries created by the OAM process.
Age	The time elapsed since the service was enabled.
Interface	The interface applied to the service.
Port	The port where the SAP is applied.

base

Syntax

base

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays basic information about the service ID including service type, description, SAPs and SDPs.



Note:

SDP information is not displayed for 7210 SAS platforms operating in access-uplink mode.

Output

The following outputs are examples of basic service information, and the associated tables describes the output fields.

- [Sample output, Table 118: Output fields: base](#)
- [Sample output — BGP PIC, Table 119: Output Fields: ID base BGP PIC](#)

Sample output

```
*A:ALA-12# show service id 1 base
=====
Service Basic Information
=====
Service Id       : 1                Vpn Id       : 0
Service Type     : VPRN
Customer Id      : 1
Last Status Change: 02/01/2007 09:11:39
Last Mgmt Change  : 02/01/2007 09:11:46
Admin State      : Up               Oper State     : Down
Route Dist.      : 10001:1
AS Number        : 10000            Router Id      : 10.10.10.103
ECMP              : Enabled          ECMP Max Routes : 8
Max Routes       : No Limit          Auto Bind      : LDP
Vrf Target       : target:10001:1
Vrf Import       : vrfImpPolCust1
Vrf Export       : vrfExpPolCust1
SAP Count        : 1                SDP Bind Count : 18
-----
Service Access & Destination Points
-----
Identifier                Type      AdmMTU  OprMTU  Adm    Opr
-----
sap:1/1/7:0               q-tag    1518    1518    Up      Up
sdp:10:1 M(10.20.1.3)     TLDP     4462    4462    Up      TLDP Down
sdp:20:1 M(10.20.1.4)     TLDP     4462    4462    Up      TLDP Down
sdp:30:1 M(10.20.1.5)     TLDP     4462    4462    Up      TLDP Down
sdp:40:1 M(10.20.1.20)    TLDP     1534    4462    Up      Up
sdp:200:1 M(10.20.1.30)   TLDP     1514    4462    Up      Up
sdp:300:1 M(10.20.1.31)   TLDP     4462    4462    Up      TLDP Down
sdp:500:1 M(10.20.1.50)   TLDP     4462    4462    Up      TLDP Down
=====
*A:ALA-12#
```

Table 118: Output fields: base

Label	Description
Service Id	The service identifier
Vpn Id	Specifies the VPN ID assigned to the service.
Service Type	Specifies the type of service
Description	Generic information about the service
Customer Id	The customer identifier
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer
Adm	The desired state of the service
Oper	The operating state of the service
Mtu	The largest frame size (in octets) that the service can handle.
Def. Mesh VC Id	This object is only valid in services that accept mesh SDP bindings It is used to validate the VC ID portion of each mesh SDP binding defined in the service
SAP Count	The number of SAPs defined on the service
SDP Bind Count	The number of SDPs bound to the service
Identifier	Specifies the service access (SAP) and destination (SDP) points
Type	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP
AdmMTU	Specifies the desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented
OprMTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented
Opr	The operating state of the SDP

Sample output — BGP PIC

```
*A:7210SAS>show>service>id# show service id 1 base
```

```
=====
Service Basic Information
```

```

=====
Service Id      : 1                      Vpn Id      : 0
Service Type    : VPRN
Name            : (Not Specified)
Description     : Default Description For VPRN ID 1
Customer Id     : 1
Last Status Change: 01/08/2000 22:57:35
Last Mgmt Change : 01/08/2000 22:57:35
Admin State     : Up                    Oper State    : Up

Route Dist.     : 100:1                  VPRN Type     : regular
AS Number       : 100                    Router Id     : 1.1.1.1
ECMP            : Enabled                 ECMP Max Routes : 1
Max IPv4 Routes : No Limit                Auto Bind     : MPLS
Max IPv6 Routes : No Limit
Ignore NH Metric : Disabled
Hash Label      : Disabled
Vrf Target      : target:200:1
Vrf Import      : None
Vrf Export      : None
MVPN Vrf Target : None
MVPN Vrf Import : None
MVPN Vrf Export : None
Label mode      : vrf
BGP VPN Backup  : ipv4 ipv6

SAP Count       : 1                      SDP Bind Count : 3

-----
Service Access & Destination Points
-----
Identifier      Type      AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/2:1     q-tag    9212   9212   Up   Up
sdp:1002:1 S(2.2.2.2) Spok     0     9186   Up   Up
sdp:1003:1 S(3.3.3.3) Spok     0     9186   Up   Up
sdp:1004:1 S(4.4.4.4) Spok     0     9186   Up   Up
=====
*A:7210SAS>show>service>id#

```

Table 119: Output Fields: ID base BGP PIC

Label	Description
Service Id	The service identifier
Service Type	The type of service: VPRN
Name	The service name
Description	Generic information about the service
Customer Id	The customer identifier
Last Status Change	The date and time of the most recent status change to this service
Last Mgmt Change	The date and time of the most recent management-initiated change to this service

Label	Description
Admin State	The desired state of the service
Oper State	The operating state of the service
Route Dist.	The largest frame size (in octets) that the service can handle
VPRN Type	Only valid in services that accept mesh SDP bindings. It validates the VC ID portion of each mesh SDP binding defined in the service.
AS Number	The autonomous system number
Router ID	The router ID for this service
ECMP	Displays equal cost multipath information
ECMP Max Routes	The maximum number of routes that can be received from the neighbors in the group or for the specific neighbor
Max IPv4 Routes	The maximum number of routes that can be used for path sharing
Auto Bind	The automatic binding type for the SDP assigned to this service
Max IPv6 Routes	Not applicable
Vrf Target	The route target in the VRF applied to this service
Vrf Import	The VRF import policy applied to this service
Vrf Export	The VRF export policy applied to this service
MVPN Vrf Target	The route target in the MVPN 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
SAP Count	The number of SAPs defined on the service
SDP Bind Count	The number of SDPs bound to the service
Service Access and Destination Points	
Identifier	The service access (SAP) and destination (SDP) points
Type	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP
AdmMTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented

Label	Description
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

statistics

Syntax

statistics [*sap sap-id*] (network and access-uplink mode)

statistics [*sdp sdp-id:vc-id*] (network mode)

statistics [*interface interface-name*] (access-uplink mode for IES)

Context

show>service>id>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays DHCP statistics information.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for the command syntax.

sdp-id

Specifies the SDP identifier.

Values 1 to 17407

vc-id

Specifies the virtual circuit ID on the SDP ID for which to display information.

Values 1 to 4294967295

interface interface-name

Displays information for the specified IP interface.

Output

The following output is an example of DHCP statistics, and [Table 120: Output Fields: DHCP statistics](#) describes the output fields.

Sample output

```
A:sim1# show service id 11 dhcp statistics
=====
DHCP Global Statistics, service 11
=====
Rx Packets                : 32
Tx Packets                : 12
Rx Malformed Packets     : 0
Rx Untrusted Packets     : 0
Client Packets Discarded  : 0
Client Packets Relayed    : 11
Client Packets Snooped    : 21
Server Packets Discarded  : 0
Server Packets Relayed    : 0
Server Packets Snooped    : 0
=====
A:sim1#
```

Table 120: Output Fields: DHCP statistics

Label	Description
Received Packets	The number of packets received from the DHCP clients.
Transmitted Packets	The number of packets transmitted to the DHCP clients.
Received Malformed Packets	The number of corrupted/invalid packets received from the DHCP clients.
Received Untrusted Packets	The number of untrusted packets received from the DHCP clients. In this case, a frame is dropped due to the client sending a DHCP packet with Option 82 filled in before "trust" is set under the DHCP interface command.
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded.
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded.
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped.
Server Packets Discarded	The number of packets received from the DHCP server that were discarded.
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded.
Server Packets Snooped	The number of packets received from the DHCP server that were snooped.

interface

Syntax

interface [*ip-address* | *ip-int-name*] [**detail**]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode with IES

Description

This command displays information for the IP interfaces associated with the service.

If no optional parameters are specified, a summary of all IP interfaces associated to the service are displayed.

Parameters

ip-address

The IP address of the interface for which to display information.

Values a.b.c.d

ip-int-name

The IP interface name for which to display information.

detail

Displays detailed IP interface information.

Default IP interface summary output.

Output

The following output is an example of service interface information, and [Table 121: Output Fields: Interface](#) describes the output fields.

Sample output

```
*A:ALA-12# show service id 321 interface
=====
Interface Table
=====
Interface-Name          Type IP-Address      Adm   Opr   Type
-----
test                    Pri  10.11.1.1/24      Up    Up    IES
-----
Interfaces : 1
=====
*A:ALA-12#
```

```

A:ALA-49# show service id 88 interface detail
=====
Interface Table
=====
Interface
-----
If Name      : Sector A
Admin State  : Up                               Oper State   : Down
Protocols    : None

IP Addr/mask : Not Assigned
-----
Details
-----
Description  :
If Index     : 26                               Virt. If Index : 26
SAP Id       : 71/1/1.2.2
TOS Marking  : Untrusted                       If Type       : IES
SNTP B.Cast  : False                           IES ID        : 88
MAC Address  : Not configured.                 Arp Timeout   : 14400
IP MTU       : 1500                            ICMP Mask Reply : True
Arp Populate : Disabled
Cflowd      : None

Proxy ARP Details
Proxy ARP    : Enabled                         Local Proxy ARP : Disabled
Policies     : ProxyARP

DHCP Details
Admin State  : Up                               Lease Populate : 0
Action       : Keep                            Trusted        : Disabled
ICMP Details
Redirects    : Number - 100                     Time (seconds) - 10
Unreachables : Number - 100                     Time (seconds) - 10
TTL Expired  : Number - 100                     Time (seconds) - 10
-----
Interface
-----
If Name      : test
Admin State  : Up                               Oper State   : Down
Protocols    : None
IP Addr/mask : Not Assigned
-----
Details
-----
Description  :
If Index     : 27                               Virt. If Index : 27
SAP Id       : 101/1/2:0
TOS Marking  : Untrusted                       If Type       : IES
SNTP B.Cast  : False                           IES ID        : 88
MAC Address  : Not configured.                 Arp Timeout   : 14400
Arp Populate : Disabled

Proxy ARP Details
Proxy ARP    : Disabled                         Local Proxy ARP : Disabled

ICMP Details
Redirects    : Number - 100                     Time (seconds) - 10
Unreachables : Number - 100                     Time (seconds) - 10
TTL Expired  : Number - 100                     Time (seconds) - 10
-----
Interfaces : 2
=====

```


A:ALA-49#

Table 121: Output Fields: Interface

Label	Description
Interface-Name	The name used to refer to the interface.
Type	Specifies the interface type.
IP-Address	Specifies the IP address/IP subnet/broadcast address of the interface.
Adm	The desired state of the interface.
Opr	The operating state of the interface.
Interface	
If Name	The name used to refer to the interface.
Admin State	The desired state of the interface.
Oper State	The operating state of the interface.
IP Addr/mask	Specifies the IP address/IP subnet/broadcast address of the interface.
Details	
If Index	The index corresponding to this interface. The primary index is 1. For example, all interfaces are defined in the Base virtual router context.
If Type	Specifies the interface type.
Port Id	Specifies the SAP port ID.
SNTP B.Cast	Specifies whether SNTP broadcast client mode is enabled or disabled.
Arp Timeout	Specifies the timeout for an ARP entry learned on the interface.
MAC Address	Specifies the 48-bit IEEE 802.3 MAC address.
ICMP Mask Reply	Specifies whether ICMP mask reply is enabled or disabled.
ICMP Details	
Redirects	Specifies the rate for ICMP redirect messages.
Unreachables	Specifies the rate for ICMP unreachable messages.
TTL Expired	Specifies the rate for ICMP TTL messages.

sap

Syntax

sap sap-id [detail]]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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

Specifies the ID that displays SAPs for the service. See [Common CLI command descriptions](#) for command syntax.
- detail

Displays detailed information for the SAP.

Output

The following output is an example of service SAP information, and [Table 122: Output Fields: service ID SAP](#) describes the output fields.

Sample output

```
*A:ALA-12# show service id 321 sap 1/1/4:0
=====
Service Access Points(SAP)
=====
Service Id      : 321
SAP             : 1/1/4:0
Dot1Q Ethertype : 0x8100
Admin State     : Up
Flags           : PortOperDown
                  SapIngressQoSMismatch
Last Status Change : 02/03/2007 12:58:37
Last Mgmt Change  : 02/03/2007 12:59:10
Admin MTU        : 1518
Ingress qos-policy : 100
Ingress Filter-Id : n/a
Multi Svc Site   : None
Acct. Pol        : None
Encap            : q-tag
QinQ Ethertype   : 0x8100
Oper State       : Down
Oper MTU         : 1518
Egress qos-policy : 1
Egress Filter-Id : n/a
Collect Stats    : Disabled
=====
*A:ALA-12#
```

```

*A:ALA-12# show service id 321 sap 1/1/4:0 detail
=====
Service Access Points(SAP)
=====
Service Id      : 321
SAP             : 1/1/4:0          Encap           : q-tag
Dot1Q Ethertype : 0x8100          QinQ Ethertype  : 0x8100

Admin State     : Up               Oper State      : Down
Flags           : PortOperDown
                 SapIngressQoSMismatch
Last Status Change : 02/03/2007 12:58:37
Last Mgmt Change  : 02/03/2007 12:59:10
Admin MTU        : 1518           Oper MTU        : 1518
Ingress qos-policy : 100          Egress qos-policy : 1
Ingress Filter-Id : n/a           Egress Filter-Id  : n/a
Multi Svc Site   : None
Acct. Pol        : None           Collect Stats    : Disabled
-----
Sap Statistics
-----
Packets          Octets
Forwarding Engine Stats
Dropped          : 0              0
Off. HiPrio      : 0              0
Off. LowPrio     : 0              0
Off. Uncolor     : 0              0
Queueing Stats(Egress QoS Policy 1)
Dro. InProf      : 0              0
Dro. OutProf     : 0              0
For. InProf      : 0              0
For. OutProf     : 0              0
-----

=====
*A:ALA-12#

*A:dut-a>config>log# /show service id 100 sap 1/1/22:100 sap-stats
=====
Service Access Points(SAP)
=====
Service Id      : 100
SAP             : 1/1/22:100       Encap           : q-tag
Description     : (Not Specified)
Admin State     : Up               Oper State      : Up
Flags           : None
Last Status Change : 02/17/2016 10:24:49
Last Mgmt Change  : 02/17/2016 10:24:46
-----
Ingress QoS Classifier Usage
-----
Classifiers Allocated: 2           Meters Allocated : 1
Classifiers Used      : 1           Meters Used       : 1
-----
Sap Statistics
-----
Packets          Octets
Ingress Stats:    0              0
Egress Stats:    76990984        116872316748
Ingress Drop Stats: 0              0

```

```

Extra-Tag Drop Stats:  n/a                      n/a
-----
Sap per Meter stats (in/out counter mode)
-----
                        Packets                      Octets
Ingress Meter 1
For. InProf           : 0                      0
For. OutProf          : 0                      0
-----
Egr sap agg-meter stats
-----
                        Packets                      Octets
Drop                  : 385943060              73232696583
Forward               : 74671326              14168884298
=====
*A:dut-a>

```

Table 122: Output Fields: service ID SAP

Label	Description
Service Id	The service identifier.
SAP	The SAP and qtag.
Encap	The encapsulation type of the SAP.
Ethertype	Specifies an Ethernet type II Ethertype value.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, PortMTUTooSmall, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, RxProtSrcMac, ParentIfAdminDown, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, ServiceMTUTooSmall, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
Last Status Change	Specifies the time of the most recent operating status change to this SAP
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SAP.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.

Label	Description
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Egress qos-policy	The egress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Table-based	Indicates the use of table-based resource classification: Enabled (table-based) or Disabled (CAM-based)
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether collect stats is enabled.
Dropped	The number of packets and octets dropped due to SAP state, ingress MAC or IP filter, same segment discard, bad checksum, and so on.
Off. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Dro. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, and so on.
Dro. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, and so on.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the ingress Qchip.
For. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, and so on.
Dro. InProf	The number of in-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, and so on.

Label	Description
Dro. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, and so on.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the egress Qchip.
For. OutProf	The number of out-of-profile packets and octets (rate above CIR) forwarded by the egress Qchip.
Ingress TD Profile	The profile ID applied to the ingress SAP.
Egress TD Profile	The profile ID applied to the egress SAP.
Alarm Cell Handling	The indication that OAM cells are being processed.
AAL-5 Encap	The AAL-5 encapsulation type.

sdp

Syntax

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

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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

Default All SDPs.

far-end ip-addr

Displays only SDPs matching with the specified far-end IP address.

Default SDPs with any far-end IP address.

detail

Displays detailed SDP information.

Output

The following output is an example of service SDP information, and [Table 123: Output Fields: service ID SDP](#) describes the output fields.

Sample output

```
A:Dut-A# show service id 1 sdp detail
=====
Services: Service Destination Points Details
=====
Sdp Id 1:1  -(10.20.1.2)
-----
Description      : Default sdp description
SDP Id           : 1:1                               Type           : Spoke
VC Type          : Ether                               VC Tag          : n/a
Admin Path MTU   : 0                                  Oper Path MTU   : 9186
Far End          : 10.20.1.2                           Delivery        : MPLS

Admin State      : Up                                Oper State      : Up
Acct. Pol       : None                               Collect Stats   : Disabled
Ingress Label   : 2048                               Egress Label    : 2048
Ing mac Fltr    : n/a                                Egr mac Fltr    : n/a
Ing ip Fltr     : n/a                                Egr ip Fltr     : n/a
Ing ipv6 Fltr   : n/a                                Egr ipv6 Fltr   : n/a
Admin ControlWord : Not Preferred                     Oper ControlWord : False
Last Status Change : 05/31/2007 00:45:43              Signaling       : None
Last Mgmt Change  : 05/31/2007 00:45:43

Class Fwding State : Up
Flags              : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None
Max Nbr of MAC Addr : No Limit                       Total MAC Addr  : 0
Learned MAC Addr   : 0                               Static MAC Addr  : 0

MAC Learning      : Enabled                           Discard Unkwn Srce: Disabled
MAC Aging         : Enabled
L2PT Termination  : Disabled                           BPDU Translation : Disabled
MAC Pinning       : Disabled

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

Statistics         :
I. Fwd. Pkts.      : 0                                I. Dro. Pkts.    : 0
I. Fwd. Octs.      : 0                                I. Dro. Octs.    : 0
E. Fwd. Pkts.      : 0                                E. Fwd. Octets   : 0
MCAC Policy Name   :
MCAC Max Unconst BW: no limit                           MCAC Max Mand BW : no limit
MCAC In use Mand BW: 0                                MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                                MCAC Avail Opnl BW: unlimited

Associated LSP LIST :
Lsp Name           : A_B_1
Admin State         : Up                                Oper State        : Up
Time Since Last Tr*: 00h26m35s
```

```

Lsp Name      : A_B_2
Admin State   : Up
Time Since Last Tr*: 00h26m35s
Oper State    : Up

Lsp Name      : A_B_3
Admin State   : Up
Time Since Last Tr*: 00h26m34s
Oper State    : Up

Lsp Name      : A_B_4
Admin State   : Up
Time Since Last Tr*: 00h26m34s
Oper State    : Up

Lsp Name      : A_B_5
Admin State   : Up
Time Since Last Tr*: 00h26m34s
Oper State    : Up

Lsp Name      : A_B_6
Admin State   : Up
Time Since Last Tr*: 00h26m34s
Oper State    : Up

Lsp Name      : A_B_7
Admin State   : Up
Time Since Last Tr*: 00h26m34s
Oper State    : Up

Lsp Name      : A_B_8
Admin State   : Up
Time Since Last Tr*: 00h26m35s
Oper State    : Up

Lsp Name      : A_B_9
Admin State   : Up
Time Since Last Tr*: 00h26m34s
Oper State    : Up

Lsp Name      : A_B_10
Admin State   : Up
Time Since Last Tr*: 00h26m34s
Oper State    : Up
-----
Class-based forwarding :
-----
Class forwarding      : enabled
Default LSP           : A_B_10
Multicast LSP         : A_B_9
=====
FC Mapping Table
=====
FC Name      LSP Name
-----
af           A_B_3
be           A_B_1
ef           A_B_6
h1           A_B_7
h2           A_B_5
l1           A_B_4
l2           A_B_2
nc           A_B_8
=====
Stp Service Destination Point specifics
-----
Mac Move      : Blockable
Stp Admin State : Up
Core Connectivity : Down
Port Role     : N/A
Port Number   : 2049
Port Path Cost : 10
Admin Edge    : Disabled
Stp Oper State : Down
Port State    : Forwarding
Port Priority  : 128
Auto Edge     : Enabled
Oper Edge     : N/A

```



```

Link Type      : Pt-pt      BPDU Encap      : Dot1d
Root Guard    : Disabled   Active Protocol : N/A
Last BPDU from : N/A
Designated Bridge : N/A      Designated Port Id: 0

Fwd Transitions : 0          Bad BPDUs rcvd  : 0
Cfg BPDUs rcvd  : 0          Cfg BPDUs tx   : 0
TCN BPDUs rcvd  : 0          TCN BPDUs tx   : 0
RST BPDUs rcvd  : 0          RST BPDUs tx   : 0
-----
Number of SDPs : 1
-----
* indicates that the corresponding row element may have been truncated.
-----
A:Dut-A#

```

Table 123: Output Fields: service ID SDP

Label	Description
Sdp Id	The SDP identifier.
Type	Indicates whether the SDP is a spoke or a mesh.
Split Horizon Group	Name of the split horizon group that the SDP belongs to.
VC Type	Displays the VC type: ether or vlan.
VC Tag	Displays the explicit dot1Q value used when encapsulating to the SDP far end.
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP.
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.)
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP.
Oper State	The operational state of this SDP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP.

Label	Description
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the keepalive process.
Oper State	he operational state of the keepalive process.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts.	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field.

aggregate

Syntax

aggregate [active]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays aggregated routes.

Parameters

active

This keyword filters out inactive aggregates.

Output

The following output is an example of aggregate route information, and [Table 124: Output Fields: Aggregate](#) describes the output fields.

Sample output

```
*A:ALA-12# show router 3 aggregate
=====
Aggregates (Service: 3)
=====
Prefix                Summary  AS Set   Aggr AS   Aggr IP-Address  State
-----
No. of Aggregates: 0
-----
*A:ALA-12#
```

Table 124: Output Fields: Aggregate

Label	Description
Prefix	Displays the destination address of the aggregate route in dotted decimal notation.
Summary	Specifies whether the aggregate or more specific components are advertised.
AS Set	Displays an aggregate where the path advertised for the route consists of all elements contained in all paths that are being summarized.
Aggr AS	Displays the aggregator path attribute to the aggregate route.
Aggr IP-Address	The IP address of the aggregated route.
State	The operational state of the aggregated route.
No. of Aggregates	The total number of aggregated routes.

arp

Syntax

arp [*ip-address* | *ip-int-name* | **mac** *ieee-mac-addr*]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode with IES

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

Only displays ARP entries associated with the specified IP address.

ip-int-name

Only displays ARP entries associated with the specified IP interface name.

macieeee-mac-addr

Only displays ARP entries associated with the specified MAC address.

Output

The following output is an example of router ARP information, and [Table 125: Output fields: ARP](#) describes the output fields.

Sample output

```
*A:ALA-12# show router 3 arp
=====
ARP Table (Service: 3)
=====
IP Address      MAC Address      Expiry      Type      Interface
-----
10.10.10.103    04:67:ff:00:00:01 00h00m00s 0th      system
10.10.4.3       00:00:00:00:00:00 00h00m00s 0th      ALA-1-2
10.10.5.3       00:00:00:00:00:00 00h00m00s 0th      ALA-1-3
10.10.7.3       00:00:00:00:00:00 00h00m00s 0th      ALA-1-5
10.10.0.16      00:00:00:00:00:00 00h00m00s 0th      bozo
10.10.3.3       00:00:00:00:00:00 00h00m00s 0th      gizmo
10.10.2.3       00:00:00:00:00:00 00h00m00s 0th      hobo
10.10.1.17      00:00:00:00:00:00 00h00m00s 0th      int-cflowd
10.0.0.92       00:00:00:00:00:00 04h00m00s Dyn      to-104
10.0.0.103      04:67:01:01:00:01 00h00m00s 0th[I]   to-104
10.0.0.104      04:68:01:01:00:01 03h59m49s Dyn[I]   to-104
10.10.36.2      00:00:00:00:00:00 00h00m00s 0th      tuesday
192.168.2.98    00:03:47:c8:b4:86 00h14m37s Dyn[I]   management
192.168.2.103   00:03:47:dc:98:1d 00h00m00s 0th[I]   management
-----
No. of ARP Entries: 14
=====
*A:ALA-12#

*A:ALA-12# show router 3 arp 10.10.0.3
=====
ARP Table
=====
IP Address      MAC Address      Expiry      Type      Interface
-----
10.10.0.3       04:5d:ff:00:00:00 00:00:00 0th      system
```

```
=====
*A:ALA-12#

*A:ALA-12# show router 3 arp to-ser1
=====
ARP Table
=====
IP Address      MAC Address      Expiry    Type    Interface
-----
10.10.13.1      04:5b:01:01:00:02 03:53:09  Dyn    to-ser1
=====
*A:ALA-12#
```

Table 125: Output fields: ARP

Label	Description
IP Address	The IP address of the ARP entry.
MAC Address	The MAC address of the ARP entry.
Expiry	The age of the ARP entry.
Type	Dyn The ARP entry is a dynamic ARP entry.
	Inv The ARP entry is an inactive static ARP entry (invalid).
	Oth The ARP entry is a local or system ARP entry.
	Sta The ARP entry is an active static ARP entry.
Interface	The IP interface name associated with the ARP entry.
No. of ARP Entries	The number of ARP entries displayed in the list.

damping

Syntax

```
damping [ip-prefix/mask | ip-address] [detail]
damping [damp-type] [detail]
```

Context

```
show>router>bgp
```

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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.

When the keyword **detail** is included, more detailed information displays.

When only the command is entered (without any parameters included except **detail**), all dampened routes are listed.

When a parameter is specified, the matching routes are listed.

When a **decayed**, **history**, or **suppressed** keyword is specified, only those types of dampened routes are listed.

Parameters

ip-prefix/mask

Displays damping information for the specified IP prefix and mask length.

ip-address

Displays damping entry for the best match route for the specified IP address.

damp-type

Displays damping type for the specified IP address.

decayed

Displays damping entries that are decayed but are not suppressed.

history

Displays damping entries that are withdrawn but have history.

suppressed

Displays damping entries suppressed because of route damping.

detail

Displays detailed information.

Output

The following output is an example of BGP damping, and [Table 126: Output fields: damping](#) describes the output fields.

Sample output

```
*A:ALA-12# show router 3 bgp damping
=====
BGP Router ID : 10.0.0.14      AS : 65206   Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes
=====
```

```

Flag Network          From          Reuse          AS-Path
-----
ud*i  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
                        2914 7459
si    10.155.40.0/21     10.0.28.1      00h28m24s      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:ALA-12#

*A:ALA-12# show router 3 bgp damping detail
=====
BGP Router ID : 10.0.0.14      AS : 65206      Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * -
valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes
=====
Network : 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 : 10.32.27.203
Local Pref       : none
Age              : 00h22m09s          Last update   : 02d00h58m
FOM Present      : 738                FOM Last upd. : 2039
Number of Flaps  : 2                  Flags         : ud*i
Path             : 60203 65001 19855 3356 1239 22406
Applied Policy   : default-damping-profile
-----
Network : 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 : 10.32.27.203
Local Pref       : none
Age              : 00h00m38s          Last update   : 02d01h20m

```

```

FOM Present      : 2011          FOM Last upd.    : 2023
Number of Flaps  : 2             Flags            : ud*i
Path             : 60203 65001 19855 3356 3561 5551 1889
Applied Policy   : default-damping-profile
-----
Network : 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 : 10.32.27.203
Local Pref   : none
Age          : 00h00m38s         Last update   : 02d01h20m
FOM Present  : 2011             FOM Last upd. : 2023
Number of Flaps : 2             Flags          : ud*i
Path         : 60203 65001 19855 1299 702 1889
Applied Policy : default-damping-profile
-----
Network : 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 : 10.32.27.203
Local Pref   : none
Age          : 00h00m07s         Last update   : 02d01h20m
FOM Present  : 1018             FOM Last upd. : 1024
Number of Flaps : 1             Flags          : ud*i
Path         : 60203 65001 19855 1299 702 1889
Applied Policy : default-damping-profile
-----
*A:ALA-12#

*A:ALA-12# show router 3 bgp damping 10.203.192.0/18 detail
=====
BGP Router ID : 10.0.0.14      AS : 65206   Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes 10.203.192.0/18
=====
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 : 10.32.27.203
Local Pref   : none
Age          : 00h00m42s         Last update   : 02d01h20m
FOM Present  : 2003             FOM Last upd. : 2025
Number of Flaps : 2             Flags          : ud*i
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Paths : 1
=====
*A:ALA-12#

*A:ALA-12# show router 3 bgp damping suppressed detail
=====
BGP Router ID : 10.0.0.14      AS : 65206   Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best

```



```

=====
BGP Damped Routes (Suppressed)
=====
Network : 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 : 10.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update  : 02d01h20m
FOM Present  : 2936              FOM Last upd. : 3001
Number of Flaps : 3              Flags        : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 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 : 10.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update  : 02d01h20m
FOM Present  : 2936              FOM Last upd. : 3001
Number of Flaps : 3              Flags        : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 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 : 10.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update  : 02d01h20m
FOM Present  : 2936              FOM Last upd. : 3001
Number of Flaps : 3              Flags        : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 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 : 10.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update  : 02d01h20m
FOM Present  : 2936              FOM Last upd. : 3001
Number of Flaps : 3              Flags        : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
*A:ALA-12#

```

Table 126: Output fields: damping

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.

Label	Description
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	Route IP prefix and mask length for the route.
Flags	Legend: Status codes: u- used, s-suppressed, h-history, d-decayed, *-valid. If a * is not present, the status is invalid. Origin codes: i-IGP, e-EGP, ?-incomplete, >-best
Network	The IP prefix and mask length for the route.
From	The originator ID path attribute value.
Reuse time	The time when a suppressed route can be used again.
AS Path	The BGP AS path for the route.
Peer	The router ID of the advertising router.
NextHop	BGP nexthop for the route.
Peer AS	The autonomous system number of the advertising router.
Peer Router-Id	The router ID of the advertising router.
Local Pref	BGP local preference path attribute for the route.
Age	The time elapsed since the service was enabled.
Last update	The time when BGP was updated last in second/minute/hour (SS:MM:HH) format.
FOM Present	The current Figure of Merit (FOM) value.
Number of Flaps	The number of flaps in the neighbor connection.
Reuse time	The time when the route can be reused.
Path	The BGP AS path for the route.
Applied Policy	The applied route policy name.

group

Syntax

group [*name*] [*detail*]

Context

show>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 operational state. Other valid states are:

Up - BGP global process is configured and running.

Down - BGP global process is administratively shutdown and not running.

Disabled - BGP global process is operationally disabled. The process must be restarted by the operator.

Parameters

name

Displays information for the BGP group specified.

detail

Displays detailed information.

Output

The following output is an example of BGP group information, and [Table 127: Output fields: group](#) describes the output fields.

Sample output

```
*A:ALA-12# show router 3 bgp group
=====
BGP Groups
=====
Group           : To_AS_40000
-----
Description      : Not Available
Group Type       : No Type           State           : Up
Peer AS          : 40000              Local AS        : 65206
Local Address    : n/a                Loop Detect     : Ignore
Export Policy    : direct2bgp
Hold Time        : 90                  Keep Alive      : 30
Cluster Id       : None                Client Reflect  : Enabled
NLRI             : Unicast              Preference     : 170

List of Peers
- 10.0.0.1       : To_Jukebox
- 10.0.0.12      : Not Available
- 10.0.0.13      : Not Available
- 10.0.0.14      : To_ALA-1
```

```

- 10.0.0.15      : To_H-215
Total Peers      : 5                      Established      : 2
=====
*A:ALA-12#

```

Table 127: Output fields: group

Label	Description
Group	BGP group name
Group Type	No Type Peer type not configured.
	External Peer type configured as external BGP peers.
	Internal Peer type configured as internal BGP peers.
State	Disabled The BGP peer group has been operationally disabled.
	Down The BGP peer group is operationally inactive.
	Up The BGP peer group is operationally active.
Peer AS	The configured or inherited peer AS for the specified peer group.
Local AS	The configured or inherited local AS for the specified peer group.
Local Address	The configured or inherited local address for originating peering for the specified peer group.
Loop Detect	The configured or inherited loop detect setting for the specified peer group.
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.

Label	Description
MED Out	The configured or inherited MED value assigned to advertised routes without a MED attribute.
Min Route Advt.	The minimum amount of time that must pass between route updates for the same IP prefix.
Min AS Originate	The minimum amount of time that must pass between updates for a route originated by the local router.
Multihop	The maximum number of router hops a BGP connection can traverse.
Multipath	The configured or inherited multipath value, determining the maximum number of ECMP routes BGP can advertise to the RTM.
Prefix Limit	No Limit No route limit assigned to the BGP peer group.
	1 - 4294967295 The maximum number of routes BGP can learn from a peer.
Passive	Disabled BGP attempts to establish BGP connections with neighbors in the specified peer group.
	Enabled BGP does not actively attempt to establish BGP connections with neighbors in the specified peer group.
Next Hop Self	Disabled BGP is not configured to send only its own IP address as the BGP nexthop in route updates to neighbors in the peer group.
	Enabled BGP sends only its own IP address as the BGP nexthop in route updates to neighbors in the specified peer group.
Aggregator ID 0	Disabled BGP is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group.
	Enabled BGP is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group.

Label	Description
Remove Private	Disabled BGP does 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.
Export Policy	The configured export policies for the peer group.
Import Policy	The configured import policies for the peer group.
Hold Time	The configured hold time setting.
Keep Alive	The configured keepalive setting.
Cluster Id	None No cluster ID has been configured.
Client Reflect	Disabled The BGP route reflector does not reflect routes to this neighbor.
	Enabled The BGP route reflector is configured to reflect routes to this neighbor.
NLRI	The type of NLRI information that the specified peer group can accept.
	Unicast IPv4 unicast routing information can be carried.
Preference	The configured route preference value for the peer group.
List of Peers	A list of BGP peers configured under the peer group.
Total Peers	The total number of peers configured under the peer group.
Established	The total number of peers that are in an established state.

neighbor

Syntax

neighbor [*ip-address* [[**family** *family*] *filter1*]]

neighbor [*as-number* [[**family** *family*] *filter2*]]

Context

show>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

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 AS number, information regarding only that specific peer or peers with the same AS display.

When either **received-routes** or **advertised-routes** is specified, the routes received from or sent to the specified peer is listed (see second output example). Note: This information is not available by SNMP.

When either **history** or **suppressed** is specified, the routes learned from those peers that either have a history or are suppressed (respectively) are listed.

The 'State' field displays the BGP peer protocol state. In addition to the standard protocol states, this field can also display the 'Disabled' operational state which indicates the peer is operationally disabled and must be restarted by the operator.

Parameters

ip-addr

Displays the BGP neighbor with the specified IP address.

family family

Specifies the type of routing information to be distributed by the BGP instance.

Values ipv4 | vpn-ipv4 | ipv6 | vpn-ipv6 | l2-vpn | ms-pw

filter1

Specifies route criteria.

Values received-routes, advertised-routes, history, suppressed, detail

filter2

Specifies route criteria.

Values history, suppressed, detail

Output

The following outputs are examples of BGP neighbor information, and the associated tables describe the output fields.

- [Sample output, Table 128: Output fields: neighbor](#)
- [Sample output for Received routes, Table 129: Output fields: neighbor received routes](#)
- [Sample output — Add-path, Table 130: Output fields: show neighbor add-path](#)

Sample output

```
*A:ALA-12# show router 3 bgp neighbor
=====
BGP Neighbor
=====
-----
Peer : 10.0.0.15          Group : To_AS_40000
-----
Peer AS      : 65205
Peer Address : 10.0.0.15      Peer Port    : 0
Local AS     : 65206
Local Address : 10.0.0.16     Local Port    : 0
Peer Type    : External
State        : Active        Last State    : Connect
Last Event   : openFail
Last Error   : Hold Timer Expire
Hold Time    : 90
Active Hold Time : 0          Keep Alive    : 30
Cluster Id   : None          Active Keep Alive: 0
Preference   : 170
Recd. Prefixes : 0           Num of Flaps   : 0
Recd. Paths   : 0           Active Prefixes : 0
Input Queue   : 0           Suppressed Paths : 0
i/p Messages  : 0           Output Queue   : 0
i/p Octets    : 0           o/p Messages   : 0
i/p Updates   : 0           o/p Octets     : 0
Export Policy : direct2bgp   o/p Updates    : 0
=====
*A:ALA-12#

*A:ALA-12# show router 3 bgp neighbor detail
=====
BGP Neighbor (detail)
=====
-----
Peer : 10.0.0.15          Group : To_AS_40000
-----
Peer AS      : 65205
Peer Address : 10.0.0.15      Peer Port    : 0
Local AS     : 65206
Local Address : 10.0.0.16     Local Port    : 0
Peer Type    : External
State        : Active        Last State    : Connect
Last Event   : openFail
Last Error   : Hold Timer Expire
Connect Retry : 20           Local Pref.   : 100
Min Route Advt. : 30         Min AS Orig.  : 15
Multipath     : 1            Multihop      : 5
Damping       : Disabled     Loop Detect    : Ignore
MED Out       : No MED Out   Authentication : None
Next Hop Self  : Disabled    AggregatorID Zero: Disabled
Remove Private : Disabled    Passive       : Disabled
```



```

Prefix Limit      : No Limit
Hold Time         : 90
Active Hold Time  : 0
Cluster Id        : None
Preference        : 170
Recd. Prefixes    : 0
Recd. Paths       : 0
Input Queue       : 0
i/p Messages      : 0
i/p Octets        : 0
i/p Updates       : 0
Export Policy     : direct2bgp
Keep Alive        : 30
Active Keep Alive : 0
Client Reflect    : Enabled
Num of Flaps      : 0
Active Prefixes   : 0
Suppressed Paths  : 0
Output Queue      : 0
o/p Messages      : 0
o/p Octets        : 0
o/p Updates       : 0

```

```

=====
*A:ALA-12#

```

Table 128: Output fields: neighbor

Label	Description
Peer	The IP address of the configured BGP peer.
Group	The BGP peer group to which this peer is assigned.
Peer AS	The configured or inherited peer AS for the peer group.
Peer Address	The configured address for the BGP peer.
Peer Port	The TCP port number used on the far-end system.
Local AS	The configured or inherited local AS for the peer group.
Local Address	The configured or inherited local address for originating peering for the peer group.
Local Port	The TCP port number used on the local system.
Peer Type	External Peer type configured as external BGP peers. Internal Peer type configured as internal BGP peers.
State	Idle The BGP peer is not accepting connections. Active BGP is listening for and accepting TCP connections from this peer. Connect BGP is attempting to establish a TCP connection from this peer. Open Sent BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer.

Label	Description
	Open Confirm BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION. Established BGP has successfully established a peering and is exchanging routing information.
Last State	Idle The BGP peer is not accepting connections. Active BGP is listening for and accepting TCP connections from this peer. Connect BGP is attempting to establish a TCP connection with this peer. Connect BGP is attempting to establish a TCP connections from this peer. Open Sent BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer. Open Confirm BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION. Open Confirm BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION.
Last Event	start BGP has initialized the BGP neighbor. stop BGP has disabled the BGP neighbor. open BGP transport connection opened. close BGP transport connection closed. openFail BGP transport connection failed to open. error BGP transport connection error.

Label	Description
	connectRetry Connect retry timer expired. holdTime Hold time timer expired. keepAlive Keepalive timer expired. recvOpen Receive an OPEN message. revKeepalive Receive an KEEPALIVE message. recvUpdate Receive an UPDATE message. recvNotify Receive an NOTIFICATION message. None No events have occurred.
Last Error	Displays the last BGP error and sub-code to occur on the BGP neighbor.
Connect Retry	The configured or inherited connect retry timer value.
Local Pref.	The configured or inherited local preference value.
Min Route Advt.	The minimum amount of time that must pass between route updates for the same IP prefix.
Min AS Originate	The minimum amount of time that must pass between updates for a route originated by the local router.
Multihop	The maximum number of router hops a BGP connection can traverse.
Multipath	The configured or inherited multipath value, determining the maximum number of ECMP routes BGP can advertise to the RTM.
Damping	Disabled BGP neighbor is configured not to dampen route flaps. Enabled BGP neighbor is configured to dampen route flaps.
Loop Detect	Ignore

Label	Description
	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 assigned to advertised routes without a MED attribute.
Authentication	None No authentication is configured. MD5 MD5 authentication is configured.
Next Hop Self	Disabled BGP is not configured to send only its own IP address as the BGP nexthop in route updates to the specified neighbor. Enabled BGP sends 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 does not remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor. Enabled BGP does remove all private AS numbers from the AS path attribute, in updates sent to the specified neighbor.
Passive	Disabled BGP actively attempts to establish a BGP connection with the specified neighbor. Enabled BGP does not actively attempt to establish a BGP connection with the specified neighbor.

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.
Hold Time	The configured hold time setting.
Keep Alive	The configured keepalive setting.
Active Hold Time	The negotiated hold time, if the BGP neighbor is in an established state.
Active Keep Alive	The negotiated keepalive time, if the BGP neighbor is in an established state.
Cluster Id	The configured route reflector cluster ID. None No cluster ID has been configured
Client Reflect	Disabled The BGP route reflector is configured not to reflect routes to this neighbor. Enabled The BGP route reflector is configured to reflect routes to this neighbor.
Preference	The configured route preference value for the peer group.
Num of Flaps	The number of flaps in the neighbor connection.
Recd. Prefixes	The number of routes received from the BGP neighbor.
Active Prefixes	The number of routes received from the BGP neighbor and active in the forwarding table.
Recd. Paths	The number of unique sets of path attributes received from the BGP neighbor.
Suppressed Paths	The number of unique sets of path attributes received from the BGP neighbor and suppressed due to route damping.
Input Queue	The number of BGP messages to be processed.
Output Queue	The number of BGP messages to be transmitted.
i/p Messages	Total number of packets received from the BGP neighbor.
o/p Messages	Total number of packets sent to the BGP neighbor.

Label	Description
i/p Octets	Total number of octets received from the BGP neighbor.
o/p Octets	Total number of octets sent to the BGP neighbor.
i/p Updates	Total number of BGP updates received from the BGP neighbor.
o/p Updates	Total number of BGP updates sent to the BGP neighbor.
Export Policy	The configured export policies for the peer group.
Import Policy	The configured import policies for the peer group.

Sample output for Received routes

```
*A:ALA-12# show router 3 bgp neighbor 10.0.0.16 received-routes
=====
BGP Router ID : 10.0.0.16          AS : 65206    Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Neighbor
=====
Flag  Network          Nexthop          LocalPref  MED      As-Path
-----
?    10.0.0.16/32        10.0.0.16        100        none     No As-Path
?    10.0.6.0/24         10.0.0.16        100        none     No As-Path
?    10.0.8.0/24         10.0.0.16        100        none     No As-Path
?    10.0.12.0/24        10.0.0.16        100        none     No As-Path
?    10.0.13.0/24        10.0.0.16        100        none     No As-Path
?    10.0.204.0/24       10.0.0.16        100        none     No As-Path
=====
*A:ALA-12#
```

Table 129: Output fields: neighbor received routes

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, it is the same value as the AS
Flag	u - used s - suppressed h - history d - decayed * - valid i - igp

Label	Description
	? - incomplete > - best
Network	Route IP prefix and mask length for the route
Next Hop	BGP nexthop for the route
LocalPref	BGP local preference path attribute for the route
MED	BGP Multi-Exit Discriminator (MED) path attribute for the route
AS Path	The BGP AS path for the route

Sample output — Add-path

```
*A:7210SAS# show router bgp neighbor 2.2.2.2

=====
BGP Neighbor
=====
-----
Peer   : 10.2.2.2
Group  : toPE
-----
Peer AS      : 100          Peer Port      : 50854
Peer Address : 10.2.2.2     Local Port    : 179
Local AS     : 100
Local Address : 10.1.1.1
Peer Type    : Internal
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  : 0
IPv4 Recd. Prefixes : 0    IPv4 Active Prefixes : 0
IPv4 Suppressed Pfxs : 0    VPN-IPv4 Suppr. Pfxs : 0
VPN-IPv4 Recd. Pfxs : 0    VPN-IPv4 Active Pfxs : 0
Mc IPv4 Recd. Pfxs : 0     Mc IPv4 Active Pfxs : 0
Mc IPv4 Suppr. Pfxs : 0    IPv6 Suppressed Pfxs : 0
IPv6 Recd. Prefixes : 0    IPv6 Active Prefixes : 0
VPN-IPv6 Recd. Pfxs : 0    VPN-IPv6 Active Pfxs : 0
VPN-IPv6 Suppr. Pfxs : 0    L2-VPN Suppr. Pfxs : 0
L2-VPN Recd. Pfxs : 0     L2-VPN Active Pfxs : 0
MVPN-IPv4 Suppr. Pfxs : 0    MVPN-IPv4 Recd. Pfxs : 0
MVPN-IPv4 Active Pfxs : 0    MDT-SAFI Suppr. Pfxs : 0
MDT-SAFI Recd. Pfxs : 0     MDT-SAFI Active Pfxs : 0
FLOW-IPv4-SAFI Suppr*: 0    FLOW-IPv4-SAFI Recd.*: 0
FLOW-IPv4-SAFI Activ*: 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 Recd. Pf*: 0    Mc Vpn Ipv4 Active P*: 0
Backup Vpn IPv4 Pfxs : 0    Backup Vpn IPv6 Pfxs : 0
Input Queue    : 0         Output Queue    : 0
```

```

i/p Messages      : 9042          o/p Messages      : 65
i/p Octets        : 111           o/p Octets        : 278
i/p Updates       : 0             o/p Updates       : 0
TTL Security      : Disabled      Min TTL Value     : n/a
Graceful Restart  : Disabled      Stale Routes Time : n/a
Advertise Inactive : Disabled      Peer Tracking     : Disabled
Advertise Label   : ipv4 ipv6
Auth key chain    : n/a
Disable Cap Nego  : Disabled      Bfd Enabled       : Enabled
Flowspec Validate : Disabled      Default Route Tgt : Disabled
L2 VPN Cisco Interop : Disabled
Local Capability  : RtRefresh MPBGP 4byte ASN
Remote Capability : RtRefresh MPBGP 4byte ASN
Local AddPath Capabi*: Send - VPN-IPv4 (1) VPN-IPv6 (4)
                  : Receive - VPN-IPv6
Remote AddPath Capab*: Send - VPN-IPv6
                  : Receive - VPN-IPv4 VPN-IPv6
Import Policy     : None Specified / Inherited
Export Policy     : P1

```

```

-----
Neighbors : 1
=====

```

* indicates that the corresponding row element may have been truncated.

*A:7210SAS#

Table 130: Output fields: show neighbor add-path

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
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

Label	Description
	Connect - BGP is attempting to establish a TCP connection with this peer Open Sent - BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer Open Confirm - BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION Established - BGP has successfully established a peering session and is exchanging routing information
Last State	Idle - the BGP peer is not accepting connections Active - BGP is listening for and accepting TCP connections from this peer Connect - BGP is attempting to establish a TCP connections with this peer Open Sent - BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer Open Confirm - BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION
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

Label	Description
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 is configured
Preference	The configured route preference value for the peer group
Num of Flaps	The number of route flaps in the neighbor connection
Recd. Prefixes	The number of routes received from the BGP neighbor
Recd. Paths	The number of unique sets of path attributes received from the BGP neighbor
IPv4 Recd. Prefixes	The number of unique sets of IPv4 path attributes received from the BGP neighbor
IPv4 Active Prefixes	The number of IPv4 routes received from the BGP neighbor and active in the forwarding table
IPv4 Suppressed Pfxs	The number of unique sets of IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
VPN-IPv4 Suppr. Pfxs	The number of unique sets of VPN-IPv4 path attributes received from the BGP neighbor and suppressed due to route damping
VPN-IPv4 Recd. Pfxs	The number of unique sets of VPN-IPv4 path attributes received from the BGP neighbor
VPN-IPv4 Active Pfxs	The number of VPN-IPv4 routes received from the BGP neighbor and active in the forwarding table
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

Label	Description
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
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
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
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)
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

Label	Description
Bfd Enabled	Enabled - BFD is enabled Disabled - BFD is disabled
Local Capability	The capability of the local BGP speaker; for example, route refresh, MP-BGP, ORF
Remote Capability	The capability of the remote BGP peer; for example, route refresh, MP-BGP, ORF
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

paths

Syntax

paths

Context

show>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays a summary of BGP path attributes.

Output

The following output is an example of BGP path information, and [Table 131: Output fields: paths](#) describes the output fields.

Sample output

```
*A:ALA-12# show router 3 bgp paths
=====
BGP Router ID : 10.0.0.14      AS : 65206   Local AS : 65206
=====
BGP Paths
-----
Path: 60203 65001 19855 3356 15412
```

```

-----
Origin      : IGP                Next Hop      : 10.0.28.1
MED         : 60203             Local Preference : none
Refs        : 4                 ASes          : 5
Segments    : 1
Flags       : EBGP-learned
Aggregator  : 15412 62.216.140.1
-----
Path: 60203 65001 19855 3356 1      1236 1236 1236 1236
-----
Origin      : IGP                Next Hop      : 10.0.28.1
MED         : 60203             Local Preference : none
Refs        : 2                 ASes          : 9
Segments    : 1
Flags       : EBGP-learned
-----
*A:ALA-12#

```

Table 131: Output fields: paths

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.
Path	The AS path attribute.
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.
Next Hop	The advertised BGP nexthop.
MED	The Multi-Exit Discriminator value.
Local Preference	The local preference value.
Refs	The number of routes using a specified set of path attributes.
ASes	The number of autonomous system numbers in the AS path attribute.
Segments	The number of segments in the AS path attribute.
Flags	eBGP-learned

Label	Description
	Path attributes learned by an eBGP peering.
	iBGP-Learned Path attributes learned by an iBGP peering.
Aggregator	The route aggregator ID.
Community	The BGP community attribute list.
Originator ID	The originator ID path attribute value.
Cluster List	The route reflector cluster list.

routes

Syntax

routes [**family** *family*] [*prefix* [**detail** | **longer**]]

routes [**family** *family*] [*prefix* [**hunt** | **brief**]]

routes [**family** *family*] [**community** *comm-id*]

routes [**family** *family*] [**aspath-regex** *reg-ex1*]

routes [**family** *family*] [*ipv6-prefix*[/*prefix-length*] [**detail** | **longer**] | [**hunt** [**brief**]]]

Context

show>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays BGP route information.

When this command is issued without any parameters, the entire BGP routing table is displayed.

When this command is issued with an IP prefix/mask or IP address, the best match for the parameter is displayed.

Parameters

family *family*

Specifies the type of routing information to be distributed by the BGP instance.

Values **ipv4** — Displays only those BGP peers that have the IPv4 family enable and not those capable of exchanging IP-VPN routes. **vpn-ipv4** — Displays the BGP peers that are IP-VPN capable. **ipv6** — Displays

the BGP peers that are IPv6 capable. **mcast-ipv4** — Displays the BGP peers that are mcast-ipv4 capable.

prefix

Specifies the type of routing information to display.

Values	<i>rd</i> [<i>rd:</i>] <i>ip-address</i> [/ <i>mask</i>]	
	rd	{ip-address:number1 as-number1:number2 as-number2:number3}
	number1	1 to 65535
	as-number1	1 to 65535
	number2	0 to 4294967295
	as-number2	1 to 4294967295
	number3	0 to 65535
	ip-address	a.b.c.d
	mask	0 to 32

filter

Specifies route criteria.

Values	hunt Displays entries for the specified route in the RIB-In, RIB-Out, and RTM.
	longer Displays the specified route and subsets of the route.
	detail Display the longer, more detailed version of the output.

aspath-regex "reg-exp"

Displays all routes with an AS path matching the specified regular expression *reg-exp*.

community comm.-id

Displays all routes with the specified BGP community.

Values	<i>[as-number1:comm-val1 ext-comm well-known-comm]</i>	
	ext-comm	type:{ip-address:comm-val1 as-number1:comm-val2 as- number2:comm-val1}
	as-number1	0..65535
	comm-val1	0..65535
	type	keywords: target, origin

```

ip-address      a.b.c.d
comm-val2      0 to 4294967295
as-number2      0 to 4294967295

well-known-comm no-export, no-export-subconfed, no-
advertise

```

Output

The following output is an example of BGP route information, and [Table 132: Output fields: routes](#) describes the output fields.

Sample output

```

*A:ALA-12>config>router>bgp# show router 3 bgp routes family ipv4
=====
BGP Router ID : 10.10.10.103      AS : 200      Local AS : 200
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Routes
=====
Flag  Network                Nexthop      LocalPref  MED
     VPN Label              As-Path
-----
No Matching Entries Found
=====
*A:ALA-12>config>router>bgp#

A:SR-12# show router bgp routes 10.0.0.0/31 hunt
=====
BGP Router ID : 10.20.1.1      AS : 100Local AS : 100
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, > - best
=====
BGP Routes
=====
RIB In Entries
-----
Network      : 10.0.0.0/31
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:SR-12#

```

Table 132: Output fields: routes

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 it is the same as the system AS.
Network	The IP prefix and mask length.
Nexthop	The BGP nexthop.
From	The advertising BGP neighbor IP address.
Res. Nexthop	The resolved nexthop.
Local Pref.	The local preference value.
Flag	u used
	s suppressed
	h history
	d decayed
	* valid
	i igp
	e egp
	? incomplete
	>

Label	Description
	best
Aggregator AS	The aggregator AS value. none No aggregator AS attributes are present.
Aggregator	The aggregator attribute value. none no Aggregator attributes are present.
Atomic Aggr.	Atomic The atomic aggregator flag is set.
	Not Atomic The atomic aggregator flag is not set.
MED	The MED metric value. none No MED metric is present.
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.
AS-Path	The BGP AS path attribute.
VPRN Imported	Displays the VPRNs where a particular BGP-VPN received route has been imported and installed.

summary

Syntax

summary [all]

Context

show>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays a summary of BGP neighbor information.

If confederations are not configured, that portion of the output does not display.

The "State" field displays the global BGP operational state. The valid values are:

Up — BGP global process is configured and running.

Down — BGP global process is administratively shutdown and not running.

Disabled — BGP global process is operationally disabled. The process must be restarted by the operator.

For example, if a BGP peer is operationally disabled, the state in the summary table shows the state 'Disabled'.

Parameters

all

Displays BGP peers in all instances.

Output

The following output is an example of summary BGP information, and [Table 133: Output fields: BGP summary](#) describes the output fields.

Sample output

```
*A:ALA-12# show router 3 bgp summary
=====
BGP Router ID : 10.0.0.14      AS : 65206   Local AS : 65206
=====
BGP Admin State      : Up           BGP Oper State      : Up
Confederation AS     : 40000
Member Confederations : 65205 65206 65207 65208

Number of Peer Groups : 2           Number of Peers      : 7
Total BGP Active Routes : 86689     Total BGP Routes     : 116999
Total BGP Paths        : 35860     Total Path Memory    : 2749476
Total Suppressed Routes : 0           Total History Routes : 0
Total Decayed Routes   : 0
=====
BGP Summary
=====
Neighbor      AS PktRcvd PktSent InQ OutQ Up/Down State|Recv/Actv/Sent
-----
10.0.0.1      65206    5  21849  0   0 00h01m29s 32/0/86683
10.0.0.12     65206    0    0    0   0 00h01m29s Active
10.0.0.13     65206    5  10545  0   0 00h01m29s 6/0/86683
10.0.0.15     65205    0    0    0   0 00h01m29s Active
10.0.0.16     65206    5   9636  0   0 00h01m29s 6/0/86683
10.0.27.1     2        0    0    0   0 00h01m29s Active
10.0.28.1     60203  22512   15   0   0 00h01m29s 116955/86689/9
=====
*A:ALA-12#
```

Table 133: Output fields: BGP summary

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 it is the same as the system 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.
Confederation AS	The configured confederation AS.
Member Confederations	The configured members of the BGP confederation.
Number of Peer Groups	The total number of configured BGP peer groups.
Number of Peers	The total number of configured BGP peers.
Total BGP Active Routes	The total number of BGP routes used in the forwarding table.
Total BGP Routes	The total number of BGP routes learned from BGP peers.
Total BGP Paths	The total number of unique sets of BGP path attributes learned from BGP peers.
Total Path Memory	Total amount of memory used to store the path attributes.
Total Suppressed Routes	Total number of suppressed routes due to route damping.
Total History Routes	Total number of routes with history due to route damping.
Total Decayed Routes	Total number of decayed routes due to route damping.
Neighbor	BGP neighbor address.

Label	Description
AS (Neighbor)	BGP neighbor autonomous system number.
PktRcvd	Total number of packets received from the BGP neighbor.
PktSent	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	The BGP neighbor current state (if not established) or the number of received routes, active routes and sent routes (if established).

interface

Syntax

interface *[[ip-address | ip-int-name][detail]] | summary*

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays the router IP interface table sorted by interface index.

Parameters

ip-address

Displays the interface information associated with the specified IP address.

ip-int-name

Displays the interface information associated with the specified IP interface name.

detail

Displays detailed IP interface information.

summary

Displays summary IP interface information for the router.

Output

The following outputs are examples of router interface information, and the associated tables describe the output fields.

- [Sample output — Standard, Table 134: Output fields: interface](#)
- [Sample output — Detailed, Table 135: Output fields: interface detail](#)
- [Sample output — Summary, Table 136: Output fields: interface summary](#)

Sample output — Standard

```
*A:7210SAS>show>router interface il detail

=====
Interface Table (Router: Base)
=====

-----
Interface
-----
If Name       : il
Admin State   : Up                               Oper (v4/v6)   : Down/--
Protocols     : None

IP Addr/mask  : Not Assigned
-----
Details
-----
Description   : (Not Specified)
If Index      : 2                               Virt. If Index : 2
Last Oper Chg: 03/07/2001 01:47:29             Global If Index: 127
Port Id       : 1/1/1
TOS Marking   : Trusted                         If Type        : Network
Egress Filter : none                           Ingress Filter  : none
Egr IPv6 Flt  : none                           Ingr IPv6 Flt   : none
SNTP B.Cast   : False                          QoS Policy     : 2
Queue-group   : None
MAC Address   : 00:25:ba:0d:27:32               Arp Timeout    : 14400
IP Oper MTU   : 9198
LdpSyncTimer  : None                           Strip-Label    : Disabled
uRPF Chk      : disabled                       uRPF Chk Fail Pk*: 0
uRPF Fail By*: 0

ICMP Details
Redirects     : Number - 100                     Time (seconds) - 10
Unreachables  : Number - 100                     Time (seconds) - 10
TTL Expired   : Number - 100                     Time (seconds) - 10

=====
Meter Statistics
=====

-----
Packets      Octets
-----
Ingress Meter 1 (Unicast)
For. InProf   : 0                               0
For. OutProf  : 0                               0
Ingress Meter 9 (Multipoint)
For. InProf   : 0                               0
For. OutProf  : 0                               0
=====
=====
```

* indicates that the corresponding row element may have been truncated.
 *A:7210SAS>show>router#

Table 134: Output fields: interface

Label	Description
Interface-Name	The IP interface name.
Type	n/a No IP address has been assigned to the IP interface, so the IP address type is not applicable. Pri The IP address for the IP interface is the Primary address on the IP interface. Sec The IP address for the IP interface is a secondary address on the IP interface.
IP-Address	The IP address and subnet mask length of the IP interface. n/a — Indicates no IP address has been assigned to the IP interface.
Adm	Down The IP interface is administratively disabled. Up The IP interface is administratively enabled.
Opr	Down The IP interface is operationally disabled. Up The IP interface is operationally enabled.
Mode	Network The IP interface is a network/core IP interface. Service The IP interface is a service IP interface.

Sample output — Detailed

```
*A:ALA-12# show router 3 interface detail
```

```
=====
```

```
Interface Table
```

```
=====
```

```
Interface
```

```

-----
If Name      : to-ser1
Admin State  : Up
Oper State   : Up

IP Addr/mask : 10.10.13.3/24
IGP Inhibit  : Disabled
Address Type : Primary
Broadcast Address: Host-ones

IP Addr/mask : 10.200.0.1/16
IGP Inhibit  : Enabled
Address Type : Secondary
Broadcast Address: Host-ones
-----
Details
-----
If Index     : 2
Port Id      : 1/1/2
Egress Filter: none
QoS Policy   : 1
MAC Address  : 04:5d:01:01:00:02
If Type      : Network
Ingress Filter : 100
SNTP Broadcast : False
Arp Timeout  : 14400

ICMP Details
Redirects     : Disabled
Unreachables : Number - 100
TTL Expired  : Number - 100
Time (seconds) - 10
Time (seconds) - 10
=====
*A:ALA-12#

```

Table 135: Output fields: interface detail

Label	Description
If Name	The IP interface name.
Admin State	Down The IP interface is administratively disabled.
	Up The IP interface is administratively enabled.
Oper State	Down The IP interface is operationally disabled.
	Up The IP interface is operationally disabled.
IP Addr/mask	The IP address and subnet mask length of the IP interface. Not Assigned - indicates no IP address has been assigned to the IP interface
Address Type	Primary The IP address for the IP interface is the Primary address on the IP interface.
	Secondary

Label	Description
	The IP address for the IP interface is a Secondary address on the IP interface.
IGP Inhibit	Disabled The secondary IP address on the interface is recognized as a local interface by the IGP.
	Enabled The secondary IP address on the interface is not recognized as a local interface by the IGP.
Broadcast Address	All-ones The broadcast format on the IP interface is all ones.
	Host-ones The broadcast format on the IP interface is host ones.
If Index	The interface index of the IP router interface.
If Type	Network The IP interface is a network/core IP interface.
	Service The IP interface is a service IP interface.
Port Id	The port ID of the IP interface.
Egress Filter	The egress IP filter policy ID associated with the IP interface. none Indicates no egress filter policy is associated with the interface.
Ingress Filter	The ingress IP filter policy ID associated with the IP interface. none Indicates no ingress filter policy is associated with the interface.
QoS Policy	The QoS policy ID associated with the IP interface.
SNTP Broadcast	False Receipt of SNTP broadcasts on the IP interface is disabled.
	True Receipt of SNTP broadcasts on the IP interface is enabled.
MAC Address	The MAC address of the IP interface.

Label	Description
Arp Timeout	The ARP timeout for the interface, in seconds, which is the time an ARP entry is maintained in the ARP cache without being refreshed.
ICMP Mask Reply	False The IP interface does not reply to a received ICMP mask request.
	True The IP interface replies to a received ICMP mask request.
Redirects	Specifies the maximum number of ICMP redirect messages the IP interface issues in a specific period of time (Time (seconds)). Disabled Indicates the IP interface does not generate ICMP redirect messages.
Unreachables	Specifies the maximum number of ICMP destination unreachable messages the IP interface issues in a specific period of time. Disabled Indicates that the IP interface does not generate ICMP destination unreachable messages.
TTL Expired	The maximum number (Number) of ICMP TTL expired messages the IP interface issues in a specific period of time (Time (seconds)). Disabled Indicates that the IP interface does not generate ICMP TTL expired messages.

Sample output — Summary

```
*A:ALA-12# show router 3 interface summary
=====
Router Summary (Interfaces)
=====
Instance  Router Name                Interfaces  Admin-Up  Oper-Up
-----
1         Base                        7          7         5
=====
*A:ALA-12#
```

Table 136: Output fields: interface summary

Label	Description
Instance	The router instance number.

Label	Description
Router Name	The name of the router instance.
Interfaces	The number of IP interfaces in the router instance.

mvpn

Syntax

mvpn

Context

show>router

Platforms

7210 SAS-T, 7210 SAS-Mxp, and 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC)

Description

This command displays Multicast VPN related information. The router instance must be specified using the **show router** command.

Output

The following output is an example of multicast VPN information, and [Table 137: Output fields: MVPN](#) describes the output fields.

Sample output

```
*A:Dut-y# show router 10 mvpn

=====
MVPN 10 configuration data
=====
signaling : Bgp auto-discovery : Default
UMH Selection : Highest-IP intersite-shared : Enabled
vrf-import : N/A
vrf-export : N/A
vrf-target : unicast
C-Mcast Import RT : target:16.16.16.16:3

ipmsi : ldp
i-pmsi P2MP AdmSt : Up

spmsi : ldp
s-pmsi P2MP AdmSt : Up
max-p2mp-spmsi : 251
data-delay-interval: 3 seconds
enable-asm-mdt : N/A
data-threshold : 224.0.0.0/4 --> 1 kbps

=====
*A:Dut-y#
```

Table 137: Output fields: MVPN

Label	Description
signaling	Displays the signaling type.
UMH Selection	Displays the UMH selection method.
vrf-import	Displays the VRF import policy in use.
vrf-export	Displays the VRF export policy in use.
vrf-target	Displays the VRF target.
C-Mcast Import RT	Displays the c-multicast import router PE system address or loopback address. This address is common for all VPNs on the PE.
ipmsi	Displays the signaling protocol used to setup the I-PMSI tree transport tunnel.
i-pmsi P2MP AdmSt	Displays I-PMSI P2MP administrative state.
spmsi	Displays signaling protocol used to setup the S-PMSI tree transport tunnel.
s-pmsi P2MP AdmSt	Displays the S-PMSI P2MP administrative state.
max-p2mp-spmsi	Displays the maximum number of P2MP S-PMSIs.
data-delay-interval	Displays the interval, in seconds, before a PE router connected to the source switches traffic from the inclusive provider tunnel to the selective provider tunnel.
enable-asm-mdt	Displays whether ASM MDT is enabled.
data-threshold	Displays the data threshold.

mvpn-list

Syntax

mvpn-list [**type** *type*] [**auto-discovery** *auto-discovery*] [**signalling** *signalling*] [**group** *group*]

Context

show>router

Platforms

7210 SAS-T, 7210 SAS-Mxp, and 7210 SAS-Sx/S 1/10GE (standalone and standalone-VC)

Description

This command displays Multicast VPN list related information.

Parameters

- type**

Specifies the MVPN type.

Values pim, rsvp, ldp
- auto-discovery**

Specifies the auto-discovery mode.

Values none, default, mdt-safi
- signalling**

Specifies the signaling type.

Values bgp, pim
- group**

Specifies the group address.

Values grp-address

Output

The following output is an example of multicast VPN list information, and [Table 138: Output Fields: MVPN list](#) describes the output fields.

Sample output

```
*A:Dut-y# show router mvpn-list

=====
MVPN List
=====
VprnID Sig A-D iPmsi/sPmsi GroupAddr/Lsp-Template (S,G)/(*,G)
-----
10 Bgp Default Mldp/Mldp N/A 512/0
20 Bgp Default Mldp/Mldp N/A 512/0
30 Bgp Default None/None N/A 0/0
-----
Total PIM I-PMSI tunnels : 0
Total RSVP I-PMSI tunnels : 0
Total MLDP I-PMSI tunnels : 2
Total PIM TX S-PMSI tunnels : 0
Total RSVP TX S-PMSI tunnels : 0
Total MLDP TX S-PMSI tunnels : 502
Total PIM RX S-PMSI tunnels : 0
Total RSVP RX S-PMSI tunnels : 0
Total MLDP RX S-PMSI tunnels : 0
Total (S,G) : 1024
Total (*,G) : 0
Total Mvpns : 3
Sig = Signal Pim-a = pim-asm Pim-s = pim-ssm A-D = Auto-Discovery
=====
```

*A:Dut-y#

Table 138: Output Fields: MVPN list

Label	Description
Total PIM I-PMSI tunnels	Displays the total number of PIM I-PMSI tunnels.
Total RSVP I-PMSI tunnels	Displays the total number of RSVP I-PMSI tunnels.
Total MLDP I-PMSI tunnels	Displays the total number of MLDP I-PMSI tunnels.
Total PIM TX I-PMSI tunnels	Displays the total number of PIM I-PMSI transmit tunnels.
Total RSVP TX I-PMSI tunnels	Displays the total number of RSVP I-PMSI transmit tunnels.
Total MLDP TX I-PMSI tunnels	Displays the total number of MLDP I-PMSI transmit tunnels.
Total PIM RX I-PMSI tunnels	Displays the total number of PIM I-PMSI receive tunnels.
Total RSVP RX I-PMSI tunnels	Displays the total number of RSVP I-PMSI receive tunnels.
Total MLDP RX I-PMSI tunnels	Displays the total number of MLDP I-PMSI receive tunnels.
Total (S,G)	Displays the total number of (S,G) multicast groups.
Total (*,G)	Displays the total number of (*,G) multicast groups.
Total Mvpngs	Displays the total number of MVPNs.

route-table

Syntax

route-table [*ip-prefix* [*/mask*] [*longer*] | [*protocol protocol*] | [*summary*]]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays the active routes in the routing table.

If no command line arguments are specified, all routes are displayed, sorted by prefix.

Parameters

ip-prefix[/mask]

Displays routes only matching the specified *ip-prefix* and optional *mask*.

longer

Displays routes matching the *ip-prefix/mask* and routes with longer masks.

protocol *protocol*

Displays routes learned from the specified protocol.

Values bgp, isis, local, ospf, rip, static, aggregate

summary

Displays a route table summary information.

Output

The following output is an example of route table information, and [Table 139: Output Fields: Router Table](#) describes the output fields.

Sample output

```
*A:ALA-12# show router 3 route-table
=====
Route Table
=====
Dest Address      Next Hop      Type      Protocol      Age      Metric      Pref
-----
10.10.0.1/32      10.10.13.1    Remote    OSPF           65844    1001        10
10.10.0.2/32      10.10.13.1    Remote    OSPF           65844    2001        10
10.10.0.3/32      0.0.0.0       Local     Local          1329261  0           0
10.10.0.4/32      10.10.34.4    Remote    OSPF           3523     1001        10
10.10.0.5/32      10.10.35.5    Remote    OSPF           1084022  1001        10
10.10.12.0/24     10.10.13.1    Remote    OSPF           65844    2000        10
10.10.13.0/24     0.0.0.0       Local     Local          65859    0           0
10.10.15.0/24     10.10.13.1    Remote    OSPF           58836    2000        10
10.10.24.0/24     10.10.34.4    Remote    OSPF           3523     2000        10
10.10.25.0/24     10.10.35.5    Remote    OSPF           399059   2000        10
10.10.34.0/24     0.0.0.0       Local     Local          3543     0           0
10.10.35.0/24     0.0.0.0       Local     Local          1329259  0           0
10.10.45.0/24     10.10.34.4    Remote    OSPF           3523     2000        10
10.200.0.0/16     0.0.0.0       Local     Local          4513     0           0
192.168.0.0/20    0.0.0.0       Local     Local          1329264  0           0
192.168.254.0/24 0.0.0.0       Remote    Static         11       1           5
-----

*A:ALA-12#

*A:ALA-12# show router 3 route-table 10.10.0.4
=====
Route Table
=====
Dest Address      Next Hop      Type      Protocol      Age      Metric      Pref
-----
```

```

10.10.0.4/32      10.10.34.4      Remote OSPF      3523      1001      10
-----
*A:ALA-12#

*A:ALA-12# show router 3 route-table 10.10.0.4/32 longer
=====
Route Table
=====
Dest Address      Next Hop          Type    Protocol    Age      Metric    Pref
-----
10.10.0.4/32      10.10.34.4      Remote  OSPF        3523     1001     10
-----
No. of Routes: 1
=====
+ : indicates that the route matches on a longer prefix
*A:ALA-12#

*A:ALA-12# show router 3 route-table protocol ospf
=====
Route Table
=====
Dest Address      Next Hop          Type    Protocol    Age      Metric    Pref
-----
10.10.0.1/32      10.10.13.1      Remote  OSPF        65844    1001     10
10.10.0.2/32      10.10.13.1      Remote  OSPF        65844    2001     10
10.10.0.4/32      10.10.34.4      Remote  OSPF        3523     1001     10
10.10.0.5/32      10.10.35.5      Remote  OSPF        1084022  1001     10
10.10.12.0/24     10.10.13.1      Remote  OSPF        65844    2000     10
10.10.15.0/24     10.10.13.1      Remote  OSPF        58836    2000     10
10.10.24.0/24     10.10.34.4      Remote  OSPF        3523     2000     10
10.10.25.0/24     10.10.35.5      Remote  OSPF        399059   2000     10
10.10.45.0/24     10.10.34.4      Remote  OSPF        3523     2000     10
-----
*A:ALA-12#

*A:ALA-12# show router 3 route-table summary
=====
Route Table Summary
=====
Active              Available
-----
Static              1              1
Direct              6              6
BGP                  0              0
OSPF                 9              9
ISIS                 0              0
RIP                  0              0
Aggregate            0              0
-----
Total                15             15
=====
*A:ALA-12#

```

Table 139: Output Fields: Router Table

Label	Description
Dest Address	The route destination address and mask.

Label	Description
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.

static-arp

Syntax

static-arp [*ip-address* | *ip-int-name* | **mac** *ieee-mac-addr*]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

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

Displays static ARP entries associated with the specified IP address.

ip-int-name

Displays static ARP entries associated with the specified IP interface name.

mac ieee-mac-addr

Displays static ARP entries associated with the specified MAC address.

Output

The following output is an example of static ARP information, and [Table 140: Output Fields: Static ARP](#) describes the output fields.

Sample output

```
*A:ALA-12# show router 3 static-arp
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
10.200.0.253    00:00:5a:40:00:01 00:00:00 Sta  to-ser1
10.200.1.1      00:00:5a:01:00:33 00:00:00 Inv  to-ser1a
-----
No. of ARP Entries: 2
=====

*A:ALA-12#

*A:ALA-12# show router 3 static-arp 10.200.1.1
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
10.200.1.1      00:00:5a:01:00:33 00:00:00 Inv  to-ser1 a
=====

*A:ALA-12#

*A:ALA-12# show router 3 static-arp to-ser1
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
10.200.0.253    00:00:5a:40:00:01 00:00:00 Sta  to-ser1
=====

S*A:ALA-12#

*A:ALA-12# show router 3 static-arp mac 00:00:5a:40:00:01
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
10.200.0.253    00:00:5a:40:00:01 00:00:00 Sta  to-ser1
=====

*A:ALA-12#
```

Table 140: Output Fields: Static ARP

Label	Description
IP Address	The IP address of the static ARP entry.
MAC Address	The MAC address of the static ARP entry.
Age	The age of the ARP entry. Static ARPs always have 00:00:00 for the age.

Label	Description
Type	Inv The ARP entry is an inactive static ARP entry (invalid).
	Sta The ARP entry is an active static ARP entry.
Interface	The IP interface name associated with the ARP entry.
No. of ARP Entries	The number of ARP entries displayed in the list.

static-route

Syntax

static-route [*ip-prefix lmask*] | [**preference** *preference*] | [**next-hop** *ip-addr* | **tag** *tag*] [**detail**]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command displays the static entries in the routing table.

If no options are present, all static routes are displayed sorted by prefix.

Parameters

ip-prefix lmask

Displays static routes only matching the specified *ip-prefix* and *mask*.

preference preference

Displays static routes with the specified route preference.

Values 0 to 65535

next-hop ip-addr

Displays static routes with the specified next hop IP address.

detail

Displays detailed information about the static route.

tag

Displays the tag used to add a 32-bit integer tag to the static route. The tag is used in route policies to control distribution of the route into other protocols.

Values 1 to 4294967295

Output

The following output is an example of static route information, and [Table 141: Output Fields: Static Route](#) describes the output fields.

Sample output

```
*A:ALA-12# show router 3 static-route
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.250.0/24  5    1    ID   10.200.10.1    to-ser1        Y
192.168.252.0/24  5    1    NH   10.10.0.254    n/a            N
192.168.253.0/24  5    1    NH   to-ser1        n/a            N
192.168.253.0/24  5    1    NH   10.10.0.254    n/a            N
192.168.254.0/24  4    1    BH   black-hole     n/a            Y
=====

*A:ALA-12#

*A:ALA-12# show router 3 static-route 192.168.250.0/24
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.250.0/24  5    1    ID   10.200.10.1    to-ser1        Y
=====

*A:ALA-12#

*A:ALA-12# show router 3 static-route preference 4
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.254.0/24  4    1    BH   black-hole     n/a            Y
=====

*A:ALA-12#

*A:ALA-12# show router 3 static-route next-hop 10.10.0.254
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.253.0/24  5    1    NH   10.10.0.254    n/a            N
=====

*A:ALA-12#

*A:Dut-B# show router static-route
=====
Static Route Table (Router: Base) Family: IPv4
=====
Prefix                                Tag      Met      Pref Type Act
```

```

Next Hop                                Interface
-----
10.2.3.4/32                             0          1          5      NH      Y
10.11.25.6
ip-10.11.25.5_base_to_cpe_static
10.11.15.0/24                           0          1          5      NH      Y
10.11.25.6
ip-10.11.25.5_base_to_cpe_static
-----
No. of Static Routes: 2
=====

*A:Dut-B# show router static-route detail
=====
Static Route Table (Router: Base)  Family: IPv4
=====
Network      : 10.2.3.4/32
Nexthop      : 10.11.25.6
Type         : Nexthop
Interface    : ip-10.11.25.5_base_to_cpe_stat*
Metric       : 1
Admin State  : Up
BFD          : disabled
CPE-check    : enabled
Target       : 10.11.18.6
Interval     : 1
Log          : N
CPE Host Up Time : 0d 00:00:02
CPE Echo Req Tx : 3
CPE Up Trans  : 1
CPE TTL      : 2
Nexthop Type : IP
Active       : Y
Preference   : 5
Tag          : 0
State        : n/a
Drop Count   : 3

Network      : 10.11.15.0/24
Nexthop      : 10.11.25.6
Type         : Nexthop
Interface    : ip-10.11.25.5_base_to_cpe_stat*
Metric       : 1
Admin State  : Up
BFD          : disabled
CPE-check    : disabled
Nexthop Type : IP
Active       : Y
Preference   : 5
Tag          : 0

No. of Static Routes: 2
=====

```

Table 141: Output Fields: Static Route

Label	Description
IP Addr/mask	The static route destination address and mask.
Pref	The route preference value for the static route.
Metric	The route metric value for the static route.
Type	BH The static route is a blackhole route. The Nexthop for this type of route is black hole.

Label	Description
	ID The static route is an indirect route, where the nexthop 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. The Nexthop for this type of route is either the next hop IP address or an egress IP interface name.
Next Hop	The next hop for the static route destination.
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.
Active	N The static route is inactive; for example, the static route is disabled or the next hop IP interface is down.
	Y The static route is active.
No. of Routes:	The number of routes displayed in the list.

tunnel-table

Syntax

tunnel-table [ip-address[/mask] [protocol protocol | sdp sdp-id]

tunnel-table [summary]

Context

show>router

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command displays tunnel table information.

When the **auto-bind** command is used when configuring a VPRN service, it means the MP-BGP NH resolution is referring to core routing instance for IP reachability. For a VPRN service this object specifies the lookup to be used by the routing instance if no SDP to the destination exists.

Parameters

ip-address[/mask]

Displays the specified tunnel table destination IP address and mask.

protocol protocol

Displays LDP protocol information.

sdp sdp-id

Displays information pertaining to the specified SDP.

summary

Displays summary tunnel table information.

Output

The following output is an example of tunnel table information, and [Table 142: Output Fields: Tunnel Table](#) describes the output fields.

Sample output

```
*A:ALA-12>config>service# show router 3 tunnel-table summary
=====
Tunnel Table Summary (Router: Base)
=====
Active Available
-----
LDP      1      1
SDP      1      1
=====
*A:ALA-12>config>service#
```

Table 142: Output Fields: Tunnel Table

Label	Description
Destination	The route destination address and mask.
Owner	Specifies the tunnel owner.
Encap	Specifies the tunnel encapsulation type.
Tunnel ID	Specifies the tunnel (SDP) identifier.
Pref	Specifies 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.

8.4.2.3 Clear commands

arp-host

Syntax

arp-host

arp-host {*mac ieee-address* | **sap** *sap-id* | **ip-address** *ip-address*[*/mask*]}

arp-host [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id* | **no-inter-dest-id**]

arp-host statistics [**sap** *sap-id* | **interface** *interface-name*]

Context

clear>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears ARP host data.

forwarding-table

Syntax

forwarding-table [*slot-number*]

Context

clear>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears the route table on the specified IOM with the route table.

If the slot number is not specified, the command forces the route table to be recalculated.

Parameters

slot-number

Clears the specified IOM slot.

Values 1 - 10 (depending on chassis model)

Default all IOMs

interface

Syntax

interface [*ip-int-name* | *ip-addr*] [**icmp**]

Context

clear>router

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears IP interface statistics.

If no IP interface is specified either by IP interface name or IP address, the command performs the clear operation on all IP interfaces.

Parameters

ip-int-name* | *ip-addr

Specifies the IP interface name or IP interface address.

Default All IP interfaces.

icmp

Specifies to reset the ICMP statistics for the IP interfaces used for ICMP rate limit.

damping

Syntax

damping [[*ip-prefix/mask*] [**neighbor** *ip-address*]] | [**group name**]

Context

clear>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command clears or resets the route damping information for received routes.

Parameters

ip-prefix/mask

Clears damping information for entries that match the IP prefix and mask length.

neighbor ip-address

Clears damping information for entries received from the BGP neighbor.

group name

Clears damping information for entries received from any BGP neighbors in the peer group.

flap-statistics

Syntax

flap-statistics *[[ip-prefix/mask] [neighbor ip-addr]] | [group group-name] | [regex reg-exp] | [policy policy-name]*

Context

clear>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command clears route flap statistics.

Parameters

ip-prefix/mask

Clears route flap statistics for entries that match the specified IP prefix and mask length.

neighbor ip-addr

Clears route flap statistics for entries received from the specified BGP neighbor.

group group-name

Clears route flap statistics for entries received from any BGP neighbors in the specified peer group.

regex reg-exp

Clears route flap statistics for all entries which have the regular expression and the AS path that matches the regular expression.

policy policy-name

Clears route flap statistics for entries that match the specified route policy.

neighbor

Syntax

neighbor {*ip-addr* | **as** *as-number* | **external** | **all**} [**soft** | **soft-inbound** | **statistics**]

Context

clear>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command resets the specified BGP peer or peers. This can cause existing BGP connections to be shutdown and restarted.

Parameters

ip-addr

Resets the BGP neighbor with the specified IP address.

as *as-number*

Resets all BGP neighbors with the specified peer AS.

external

Resets all eBGP neighbors.

all

Resets all BGP neighbors.

soft

The specified BGP neighbors reevaluates all routes in the Local-RIB against the configured export policies.

soft-inbound

The specified BGP neighbors reevaluates all routes in the RIB-In against the configured import policies.

statistics

The BGP neighbor statistics.

protocol

Syntax

protocol

Context

clear>router>bgp

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command resets the entire BGP protocol. If the AS number was previously changed, the BGP AS number does not inherit the new value.

id

Syntax

id *service-id*

Context

clear>service

clear>service>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears commands for a specific service.

Parameters

service-id

The ID that uniquely identifies a service.

Values 1 to 2147483648

sap

Syntax

sap *sap-id* {all | counters | stp}

Context

clear>service>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears SAP statistics for a SAP.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

spoke-sdp

Syntax

spoke-sdp *sdp-id:vc-id* ingress-vc-label

Context

clear>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command clears and resets the spoke-SDP bindings for the service.

Parameters

sdp-id

Specifies the spoke-SDP ID to be reset.

Values 1 to 17407

vc-id

Specifies the virtual circuit ID on the SDP ID to be reset.

Values 1 to 4294967295

sdp

Syntax

sdp *sdp-id* keep-alive

Context

clear>service>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command clears keepalive statistics associated with the SDP ID.

Parameters

sdp-id

Specifies the SDP ID for which to clear keepalive statistics.

Values 1 to 17407

counters

Syntax

counters

Context

clear>service>statistics>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears all traffic queue counters associated with the service ID.

spoke-sdp

Syntax

spoke-sdp *sdp-id[:vc-id]* {all | **counters** | **stp**}

Context

clear>service>statistics>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command clears statistics for the spoke-SDP bound to the service.

Parameters

sdp-id

Specifies the spoke-SDP ID for which to clear statistics.

Values 1 to 17407

vc-id

Specifies 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.

stp

Clears all STP statistics associated with the SDP.

```
stp
```

Syntax

```
stp
```

Context

```
clear>service>statistics>id
```

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command clears all spanning tree statistics for the service ID.

8.4.2.4 Debug commands

```
id
```

Syntax

```
[no] id service-id
```

Context

debug>service

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command debugs commands for a specific service.

The **no** form of this command disables debugging.

Parameters

service-id

Specifies the ID that uniquely identifies a service.

sap

Syntax

[no] **sap** *sap-id*

Context

debug>service>id

debug>service>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for a specific SAP.

The **no** form of this command disables debugging.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

sdp

Syntax

[no] **sdp** *sdp-id:vc-id*

Context

debug>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, except those operating in access-uplink mode

Description

This command enables STP debugging for a specific SDP.

The **no** form of this command disables debugging.

event-type

Syntax

[no] event-type {config-change | svc-oper-status-change | sap-oper-status-change | sdpbind-oper-status-change}

Context

debug>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables debugging for a particular event type.

The **no** form of this command disables debugging.



Note:

The **sdpbind-oper-status-change** parameter is not supported on 7210 SAS platforms operating in access-uplink mode.

event-type

Syntax

[no] event-type {config-change | oper-status-change}

Context

debug>service>id>sap

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables debugging for a particular event type.

The **no** form of this command disables debugging.

```
stp
```

Syntax

[no] stp

Context

debug>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables the context for debugging STP.

The **no** form of this command disables debugging.

```
all-events
```

Syntax

all-events

Context

debug>service>id>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for all events.

The **no** form of this command disables debugging.

```
bpdu
```

Syntax

[no] bpdu

Context

debug>service>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for received and transmitted BPDUs.

The **no** form of this command disables debugging.

core-connectivity**Syntax**

[no] core-connectivity

Context

debug>service>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for core connectivity.

The **no** form of this command disables debugging.

exception**Syntax**

[no] exception

Context

debug>service>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for exceptions.

The **no** form of this command disables debugging.

fsm-state-changes

Syntax

[no] fsm-state-changes

Context

debug>service>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for FSM state changes.

The **no** form of this command disables debugging.

fsm-timers

Syntax

[no] fsm-timers

Context

debug>service>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for FSM timer changes.

The **no** form of this command disables debugging.

port-role

Syntax

[no] port-role

Context

debug>service>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for changes in port roles.

The **no** form of this command disables debugging.

port-state

Syntax

[no] port-state

Context

debug>service>stp

Platforms

Supported on all 7210 SAS platforms as described in this document, including those operating in access-uplink mode

Description

This command enables STP debugging for port states.

The **no** form of this command disables debugging.

9 Common CLI command descriptions

This chapter provides information about Command Line Interface (CLI) syntax and command usage for common service commands.

9.1 Command descriptions

9.1.1 SAP syntax

sap

Syntax

[no] **sap** *sap-id*

Context

Platforms

7210 SAS-T, 7210 SAS-Sx/S 1/10GE, 7210 SAS-Sx 10/100GE, and 7210 SAS-Mxp

Description

This command specifies the physical port identifier portion of the SAP definition.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition.
The *sap-id* can be configured in one of the formats listed in the following table.

Table 143: SAP-ID formats

Type	Syntax	Example
port-id	<i>slot/mda/port[.channel]</i>	1/1/5
null	<i>[port-id lag-id]</i>	<i>port-id:</i> 1/1/3 <i>lag-id:</i> lag-3
dot1q	<i>[port-id lag-id]:qtag1</i>	<i>port-id:</i> qtag1: 1/1/3:100 <i>lag-id:</i> qtag1:lag-3:102

Type	Syntax	Example
		<i>cp.conn-prof-id: 1/2/1:cp.2</i>
qinq	<i>[port-id lag-id]:qtag1.qtag2</i>	<i>port-id:qtag1.qtag2: 1/1/3:100.10</i> <i>lag-id:qtag1.qtag2: lag-10:</i>

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 144: Encapsulation types

Port type	Encap-type	Allowed values	Comments
Ethernet	Null	0	The SAP is identified by the port.
Ethernet	Dot1q	0 to 4094	The SAP is identified by the 802.1Q tag on the port. ¹⁸
Ethernet	QinQ	qtag1: 0 to 4094 qtag2: 0 to 4094	The SAP is identified by two 802.1Q tags on the port. ^{19 20}

¹⁸ A 0 qtag1 value also accepts untagged packets on the dot1q port.

¹⁹ A 0 qtag1 value is allowed with some 7210 SAS platforms. See [SAP configuration notes for 7210 SAS platforms in access-uplink operating mode](#) for information about platforms and frame processing.

²⁰ A 0 qtag2 value is not allowed on 7210 SAS platforms as described in this document. See [SAP configuration notes for 7210 SAS platforms in access-uplink operating mode](#) for more information about allowed SAPs and their processing behavior.

10 Appendix: Port-based split horizon

This chapter provides Port-Based Split Horizon configuration information.

10.1 Overview



Note:

Port-based split horizon is supported on all 7210 SAS platforms as described in this document, except the 7210 SAS-Mxp.

The port-based split horizon feature can be used to disable local switching on the 7210 SAS. A loop-free topology can be achieved using split horizon on 7210 SAS switches.


Traffic arriving on an access or a network port within a split horizon group will not be copied to other access and a network ports in the same split horizon group, but will be copied to an access or network ports in other split horizon groups.

Since split horizon is a per port feature in 7210 SAS, all SAPs associated with the port becomes part of split horizon group configured on that port.

10.1.1 Topology

The following figure shows an example of split horizon groups used to prevent communication between two access SAPs and between two network ports.

Figure 99: Split horizon group example



Using 7210 SAS-1 as an example:

1. Split horizon group "access" is created to prevent any communication between the SAP part of port 1/1/4 and port 1/1/5 (configured as access port) within the same VPLS.
2. Split horizon group "network" is created to prevent any communication between port 1/1/1 and port 1/1/3 (configured as a network port) within the same VPLS.
3. VPLS 100 is created on 7210 SAS-1 with spoke-SDPs on network port 1/1/1 and 1/1/3, and SAPs on 1/1/4 and 1/1/5 as part of this VPLS. CE1, CE2 and CE3 are the customer sites.

4. With this configuration, any communication between ports 1/1/4 and 1/1/5 gets blocked, similarly communication between ports 1/1/1 and 1/1/3 gets blocked but any traffic received on ports (for example, spoke-SDPs on these ports) that belong to split horizon group "network" will be switched to ports (for example, SAPs on these ports) that belong to split horizon group "access" and the other way around based on the FDB entries for VPLS 100.

10.2 Configuration guidelines

The following configuration guidelines must be followed to configure a split horizon group:



Note:

When configuring split horizon on a port, it must be configured before creating any SAPs associated with the port.

1. Create a split horizon group in the **config** prompt. The group name must be unique across the system.

```
7210-SAS1>config#info
#-----
echo "Split-horizon-group Configuration"
#-----
split-horizon-group access create
description "Block access between access Ports"
split-horizon-group network create
description "Block access between network Ports"
exit
#-----
7210-SAS1>config#
```

2. Configure ports 1/1/4 and 1/1/5 as access ports and associate these ports with split horizon group "access".

```
7210-SAS1>config#info
#-----
echo "Port Configuration"
#-----
port 1/1/4
split-horizon-group access
ethernet
mode access
access
exit
exit
no shutdown
exit
port 1/1/5
split-horizon-group access
ethernet
mode access
access
exit
exit
no shutdown
exit
#-----
7210-SAS1>config#
```

3. Configure ports 1/1/1 and 1/1/3 as network ports and associate these ports with split horizon group "network". The default Ethernet encapsulation for network port is null.

```
7210-SAS1>config# info
#-----
echo "Port Configuration"
#-----
    port 1/1/1
        split-horizon-group network
        ethernet
        exit
        no shutdown
    exit
    port 1/1/3
        split-horizon-group network
        ethernet
        exit
        no shutdown
    exit
#-----
7210-SAS1>config#
```

4. Create a VPLS instance 100.

```
#-----
echo "Service Configuration"
#-----
    service
        customer 2 create
        exit
        vpls 100 customer 2 create
            stp
                shutdown
            exit
            sap 1/1/4 create
            exit
            sap 1/1/5 create
            exit
            spoke-sdp 1:1 create
            exit
            spoke-sdp 2:1 create
            exit
        no shutdown
    exit
    ...
#-----
```

10.2.1 Verification

The following output verifies the split horizon configuration on a 7210 SAS.

```
7210-SAS1# show split-horizon-group
=====
Port: Split Horizon Group
=====
Name                               Description
-----
access                             Block access between access Ports
network                             Block access between network Ports
```

```
No. of Split Horizon Groups: 2
=====
7210-SAS1#
```

The following shows the command usage to verify the port association with split horizon groups:

```
7210-SAS1# show split-horizon-group access
=====
Port: Split Horizon Group
=====
Name                               Description
-----
access                             Block access between access Ports
-----
Associations
-----
Port1/1/4                          10/100/Gig Ethernet SFP
Port1/1/5                          10/100/Gig Ethernet SFP

Ports Associated : 2
=====
7210-SAS1#

7210-SAS1# show split-horizon-group network
=====
Port: Split Horizon Group
=====
Name                               Description
-----
network                             Block access between network Ports
-----
Associations
-----
Port1/1/1                          10/100/Gig Ethernet SFP
Port1/1/3                          10/100/Gig Ethernet SFP

Ports Associated : 2
=====
7210-SAS1#
```

11 Appendix: DHCP management

This chapter provides information about using DHCP, including theory, supported features and configuration process overview.


11.1 DHCP principles

In a Triple Play network, client devices (such as a routed home gateway, a session initiation protocol (SIP) phone or a set-top box) use Dynamic Host Configuration Protocol (DHCP) to dynamically obtain their IP address and other network configuration information. The 7210 auto-init procedure also uses DHCP to dynamically obtain the BOF file used for first-time booting of the system (along with IP address required to retrieve the BOF file, the configuration file and the Timos software image from the network). DHCP is defined and shaped by several RFCs and drafts in the IETF DHC working group including the following:

- RFC 2131, *Dynamic Host Configuration Protocol*
- RFC 3046, *DHCP Relay Agent Information Option*

The DHCP operation is shown in the following figure.

Figure 100: IP address assignment with DHCP



1. During boot-up, the client device sends a DHCP discover message to get an IP address from the DHCP Server. The message contains:

- destination MAC address - broadcast
- source MAC address - MAC of client device
- client hardware address - MAC of client device

If this message passes through a DSLAM or other access node (possibly a 7210 SAS device), typically the Relay information option (Option 82) field is added, indicating shelf, slot, port, VPI, VCI and other fields, to identify the subscriber.

DHCP relay is enabled on the first IP interface in the upstream direction. Depending on the scenario, the DSLAM, BSA or the BSR will relay the discover message as a unicast packet toward the configured DHCP server. DHCP relay is configured to insert the giaddr to indicate to the DHCP server in which subnet an address should be allocated.

2. The DHCP server will lookup the client MAC address and Option 82 information in its database. If the client is recognized and authorized to access the network, an IP address will be assigned and a DHCP offer message returned. The BSA or BSR will relay this back to the client device.
3. It is possible that the discover reached more than one DHCP server, and therefore that more than one offer was returned. The client selects one of the offered IP addresses and confirms it needs to use this in a DHCP request message, sent as unicast to the DHCP server that offered it.
4. The DHCP server confirms that the IP address is still available, updates its database to indicate it is now in use, and replies with a DHCP ACK message back to the client. The ACK also contains the Lease Time of the IP address.

11.1.1 DHCP features

11.1.1.1 Using Option 82 field

Option 82, or the relay information option is specified in RFC 3046, *DHCP Relay Agent Information Option*, allows the router to append some information to the DHCP request that identifies where the original DHCP request arrives from.

There are two sub-options under Option 82:

- Agent Circuit ID Sub-option (RFC 3046, section 3.1): This sub-option specifies data which must be unique to the box that is relaying the circuit.
- Remote ID Sub-option (RFC 3046 section 3.2): This sub-option identifies the host at the other end of the circuit. This value must be globally unique.

Both sub-options are supported by the 7210 SAS and can be used separately or together.

Inserting Option 82 information is supported independently of DHCP relay.

When the circuit ID sub-option field is inserted by the 7210 SAS, it can take following values:

- *sap-id* - the SAP index (only under a IES or VPRN service)
- *ifindex* - the index of the IP interface (only under a IES or VPRN service)
- *ascii-tuple* - an ASCII-encoded concatenated tuple, consisting of [system-name | service-id | interface-name] (for VPRN or IES) or [system-name | service-id | sap-id] (for VPLS)

- *vlan-ascii-tuple* - an ASCII-encoded concatenated tuple, consisting of the ascii-tuple followed by Dot1p bits and Dot1q tags

Note that for VPRN the ifindex is unique only within a VRF. The DHCP relay function automatically prepends the VRF ID to the ifindex before relaying a DHCP Request.

When a DHCP packet is received with Option 82 information already present, the system can do one of three things. The available actions are:

- **Replace**

On ingress the existing information-option is replaced with the information-option parameter configured on the 7210 SAS. On egress (toward the customer) the information-option is stripped (per the RFC).

- **Drop**

The DHCP packet is dropped and a counter is incremented.

- **Keep**

The existing information is kept on the packet and the router does not add any more information. On egress the information option is not stripped and is sent on to the downstream node.

In accordance with the RFC, the default behavior is to keep the existing information; except if the giaddr of the packet received is identical to a local IP address on the router, then the packet is dropped and an error incremented regardless of the configured action.

The maximum packet size for a DHCP relay packet is 1500 bytes. If adding the Option 82 information would cause the packet to exceed this size, the DHCP relay request will be forwarded without the Option 82 information. This packet size limitation exists to ensure that there will be no fragmentation on the end Ethernet segment where the DHCP server attaches.

In the downstream direction, the inserted Option 82 information should not be passed back toward the client (as per RFC 3046, DHCP Relay Agent Information Option). To enable downstream stripping of the option 82 field, DHCP snooping should be enabled on the SDP or SAP connected to the DHCP server.

11.1.1.2 Trusted and untrusted

There is a case where the relay agent could receive a request where the downstream node added Option 82 information without also adding a giaddr (giaddr of 0). In this case the default behavior is for the router to drop the DHCP request. This behavior is in line with the RFC.

The 7210 SAS supports a command `trusted`, which allows the router to forward the DHCP request even if it receives one with a giaddr of 0 and Option 82 information attached. This could occur with older access equipment. In this case the relay agent would modify the request's giaddr to be equal to the ingress interface. This only makes sense when the action in the information option is `keep`, and the service is IES or VPRN. In the case where the Option 82 information gets replaced by the relay agent, either through explicit configuration or the VPLS DHCP Relay case, the original Option 82 information is lost, and the reason for enabling the trusted option is lost.

11.1.2 Common configuration guidelines

11.1.2.1 Configuration guidelines for DHCP relay and snooping

The following configuration guidelines must be followed to configure DHCP relay and snooping.

- 7210 SAS devices does not support the ARP populate based on the DHCP lease, assigned to the DHCP client
- 7210 SAS devices does not maintain the DHCP lease assigned to the client
- 7210 SAS devices do not perform IP spoofing checks and MAC spoofing checks based on the DHCP parameters assigned to the client
- MAC learning must be enabled in the VPLS service, for DHCP snooping.
- DHCP snooping is not supported for B-SAPs in B-VPLS services and I-SAPs in I-VPLS services.
- Ingress ACLs cannot be used to drop DHCP control packet.
- DHCP packets received over a SDP cannot be identified and option-82 inserted by the node cannot be removed by the node, in the downstream direction. If this behavior is not needed user should not enable DHCP snooping in the VPLS service, if the DHCP server is reachable over the SDP (either spoke-SDP or mesh SDP).

11.1.2.2 Configuring Option 82 handling

Option 82, or "Relay Information Option" is a field in DHCP messages used to identify the subscriber. The Option 82 field can already be filled in when a DHCP message is received at the router, or it can be empty. MAC learning must be enabled in the VPLS service, for DHCP snooping. If the field is empty, the router should add identifying information (circuit ID, remote ID or both). If the field is not empty, the router can decide to replace it.

Example: Partial BSA configuration — adding Option 82 to a VPLS

The following is a sample partial BSA configuration output with Option 82 adding on a VPLS service. Note that snooping must be enabled explicitly on a SAP.

```
*A:7210SAS>config>service#
-----

vpls 2 customer 1 create
    shutdown
    stp
        shutdown
    exit
sap 1/1/12:100 create
    dhcp
        option //Configuration example to add option 82
            action replace
            circuit-id
            no remote-id
        exit
        no shutdown
    exit
    exit
    no shutdown
    exit

-----

*A:7210SAS>config>service#
```

Example: Partial BSA configuration — removing Option 82 from a VPLS

The following is a sample partial BSA configuration output to remove the Option 82 on a VPLS service.

```
vpls 2 customer 1 create
    stp
        shutdown
    exit
    sap 1/1/14:100 create      //Configuration example to remove option 82
        dhcp
            snoop
            no shutdown
        exit
    exit
```


12 Standards and protocol support



Note:

The information provided in this chapter is subject to change without notice and may not apply to all platforms.

Nokia assumes no responsibility for inaccuracies.

The following conventions are used in this section:

- T(A,N) indicates 7210 SAS-T in both Access-uplink mode and Network mode. Similarly, T(N) indicates 7210 SAS-T in network mode only.
- K5 indicates 7210 SAS-K 2F2T1C.
- K12 indicates 7210 SAS-K 2F4T6C.
- K30 indicates 7210 SAS-K 3SFP+ 8C.
- Sx/S-1/10GE indicates all variants of 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE platforms.
- Sx-10/100GE indicates only the variants of 7210 SAS-Sx 10/100GE. Sx-10/100GE QSFP28 indicates the 7210 SAS-Sx 10/100GE 64 SFP+ 4QSFP28 variant.
- Sx/S-1/10GE-VC indicates 7210 SAS-Sx 1/10GE and 7210 SAS-S 1/10GE platforms in standalone VC mode.
- R6 indicates 7210 SAS-R6.
- R12 indicates 7210 SAS-R12.
- D indicates 7210 SAS-D and 7210 SAS-D ETR. If a line item applies only to 7210 SAS-D ETR, then it is indicated as D-ETR.
- Dxp indicates 7210 SAS-Dxp-12p (2SFP+ 4F6T) 7210 SAS-Dxp-12p ETR (2SFP+ 4F6T), 7210 SAS-Dxp 16p (2SFP+ 4F10T), and 7210 SAS-Dxp-24p (2SFP+ 6F16T). If a line item applies only to a particular variant, the variant name will be called out explicitly against that item.
- This standards list is not applicable to platforms in the satellite mode of operation, as most of the features are supported on 7x50 SR platforms. For this reason, the host platforms standards compliance must be consulted for the satellite mode of operation.

12.1 BGP

draft-ietf-idr-add-paths-04, Advertisement of Multiple Paths in BGP is supported on T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-ietf-sidr-origin-validation-signaling-04, BGP Prefix Origin Validation State Extended Community is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1772, Application of the Border Gateway Protocol in the Internet is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1997, BGP Communities Attribute is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2385, Protection of BGP Sessions via the TCP MD5 Signature Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2439, BGP Route Flap Damping is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2545, Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2858, Multiprotocol Extensions for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2918, Route Refresh Capability for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3107, Carrying Label Information in BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3392, Capabilities Advertisement with BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4271, A Border Gateway Protocol 4 (BGP-4) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4360, BGP Extended Communities Attribute is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4456, BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

Only R6 and R12 supports RR server functionality. Rest of the platforms support only client function.

RFC 4659, BGP/MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4684, Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4760, Multiprotocol Extensions for BGP-4 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4798, Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4893, BGP Support for Four-octet AS Number Space is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5004, Avoid BGP Best Path Transitions from One External to Another is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5291, Outbound Route Filtering Capability for BGP-4 is supported on K12, K30, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5668, 4-Octet AS Specific BGP Extended Community is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6811, Prefix Origin Validation is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

12.2 Certificate Management

RFC 7468, Textual Encodings of PKIX, PKCS, and CMS Structures is supported on Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

Only PKIX and PKCS is supported.

12.3 Ethernet

IEEE 802.1AB, Station and Media Access Control Connectivity Discovery is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

TIA-1057, LLDP for Media endpoint devices is supported on Dxp, Sx/S-1/10GE, and Sx/S-1/10GE-VC

IEEE 802.1ad, Provider Bridges is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1ag, Connectivity Fault Management is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1ah, Provider Backbone Bridges is supported on T(N)

IEEE 802.1ax, Link Aggregation is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1D, MAC Bridges is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1p, Traffic Class Expediting is supported on D, E, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1Q, Virtual LANs is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1s, Multiple Spanning Trees is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1w, Rapid Reconfiguration of Spanning Tree is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.1X, Port Based Network Access Control is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ab, 1000BASE-T is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ac, VLAN Tag is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ad, Link Aggregation is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ae, 10 Gb/s Ethernet is supported on Dxp, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ah, Ethernet in the First Mile is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3ba, 40 Gb/s and 100 Gb/s Ethernet is supported on R6, R12, and Sx-10/100GE

IEEE 802.3i, Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3u, Fast Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3z, Gigabit Ethernet is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE 802.3af, Power Over Ethernet (PoE) is supported on Dxp, T-ETR, Mxp-ETR, Sx/S-1/10GE, and Sx/S-1/10GE-VC

**Note:**

Sx/S-1/10GE only on PoE variant and Sx-1/10GE fiber variant with two fixed copper ports. Only on Dxp-16p and Dxp-24p.

IEEE 802.3at, Power Over Ethernet (PoE+) is supported on Dxp, T-ETR, Mxp-ETR, Sx/S-1/10GE, and Sx/S-1/10GE-VC

**Note:**

Sx/S-1/10GE only on PoE variant and Sx-1/10GE fiber variant with two fixed copper ports. Only on Dxp-16p and Dxp-24p.

IEEE 802.3bt, Power Over Ethernet (PoE++/HPoE) is supported on Dxp

**Note:**

Only on Dxp-16p and Dxp-24p.

ITU-T G.8032, Ethernet Ring Protection Switching is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ITU-T Y.1731, OAM functions and mechanisms for Ethernet based networks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

12.4 EVPN

draft-snr-bess-evpn-proxy-arp-nd-00, Proxy-ARP/ND function in EVPN networks is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12

**Note:**

Sx/S-1/10GE standalone mode only.

RFC 7432, BGP MPLS-Based Ethernet VPN is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12

**Note:**

Sx/S-1/10GE standalone mode only.

draft-ietf-bess-evpn-vpws-14, Virtual Private Wire Service support in Ethernet VPN is supported on Mxp

12.5 Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, Operational management of Loop Free Alternates is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

With Segment Routing.

RFC 5286, Basic Specification for IP Fast Reroute: Loop-Free Alternates is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

With Segment Routing.

RFC 7490, Remote Loop-Free Alternate (LFA) Fast Reroute (FRR) is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

With Segment Routing.

12.6 Internet Protocol (IP) — General

draft-grant-tacacs-02, The TACACS+ Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-vrrp-unified-spec-02, Virtual Router Redundancy Protocol Version 3 for IPv4 and IPv6 is supported on Mxp

RFC 768, User Datagram Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 793, Transmission Control Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 854, TELNET Protocol Specifications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 951, Bootstrap Protocol (BOOTP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1034, Domain Names - Concepts and Facilities is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1035, Domain Names - Implementation and Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1350, The TFTP Protocol (revision 2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1534, Interoperation between DHCP and BOOTP is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1542, Clarifications and Extensions for the Bootstrap Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2131, Dynamic Host Configuration Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2132, DHCP Options and BOOTP Vendor Extensions is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2347, TFTP Option Extension is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2348, TFTP Blocksize Option is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2349, TFTP Timeout Interval and Transfer Size Options is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2428, FTP Extensions for IPv6 and NATs is supported on D, Dxp, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2865, Remote Authentication Dial In User Service (RADIUS) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2866, RADIUS Accounting is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3046, DHCP Relay Agent Information Option (Option 82) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3596, DNS Extensions to Support IP version 6 is supported on D, Dxp, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3768, Virtual Router Redundancy Protocol (VRRP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4250, The Secure Shell (SSH) Protocol Assigned Numbers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4251, The Secure Shell (SSH) Protocol Architecture is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4252, The Secure Shell (SSH) Authentication Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

All 7210 platforms support password and publickey based user authentication. 7210 SAS-D support only password based authentication.

RFC 4253, The Secure Shell (SSH) Transport Layer Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4254, The Secure Shell (SSH) Connection Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5246, The Transport Layer Security (TLS) Protocol Version 1.2 – TLS client, RSA public key is supported on Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

Only TLS client.

RFC 5425, Transport Layer Security (TLS) Transport Mapping for Syslog – RFC 3164 with TLS is supported on Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

Only TLS client.

RFC 5656, Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer – ECDSA is supported on Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7301, Transport Layer Security (TLS) Application Layer Protocol Negotiation Extension is supported on Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

Only TLS client.

RFC 8446, The Transport Layer Security (TLS) Protocol Version 1.3 is supported on Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

Only TLS client.

RFC 4632, Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5880, Bidirectional Forwarding Detection (BFD) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 5881, Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 5883, Bidirectional Forwarding Detection (BFD) for Multihop Paths is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

IPv4 only on all platforms listed. IPv4 and IPv6 only on Mxp.

RFC 6528, Defending against Sequence Number Attacks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7130, Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces is supported on T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

12.7 IP — Multicast

RFC 1112, Host Extensions for IP Multicasting is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

IGMP v1, v2, v3 is supported.

RFC 2236, Internet Group Management Protocol, Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3306, Unicast-Prefix-based IPv6 Multicast Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3376, Internet Group Management Protocol, Version 3 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3446, Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

MSDP supported only on Sx/S-1/10GE standalone.

RFC 3618, Multicast Source Discovery Protocol (MSDP) is supported on Sx/S-1/10GE

**Note:**

Only in standalone mode.

RFC 4601, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4604, Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

MLD not supported.

RFC 4607, Source-Specific Multicast for IP is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4608, Source-Specific Protocol Independent Multicast in 232/8 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4610, Anycast-RP Using Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5059, Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5384, The Protocol Independent Multicast (PIM) Join Attribute Format is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6513, Multicast in MPLS/BGP IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6514, BGP Encodings and Procedures for Multicast in MPLS/IP VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6515, IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6625, Wildcards in Multicast VPN Auto-Discover Routes is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 6826, Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7246, Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

RFC 7385, IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:
Only IPv4.

12.8 IP — Version 4

RFC 791, Internet Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 792, Internet Control Message Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 826, An Ethernet Address Resolution Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1519, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1812, Requirements for IPv4 Routers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1981, Path MTU Discovery for IP version 6 is supported on K12, K30, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2401, Security Architecture for Internet Protocol is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

Supported only for OSPFv3 authentication. Not supported for services.

RFC 2460, Internet Protocol, Version 6 (IPv6) Specification is supported on D, Dxp, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

D, Dxp, and T(A) for Management only.

12.9 IP — Version 6

RFC 2464, Transmission of IPv6 Packets over Ethernet Networks is supported on D, Dxp, K5 K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

D, Dxp, K5, and T(A) for Management only.

RFC 3021, Using 31-Bit Prefixes on IPv4 Point-to-Point Links is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3122, Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification is supported on K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3587, IPv6 Global Unicast Address Format is supported on D, Dxp, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

D, Dxp, K5, and T(A) for Management only.

RFC 4007, IPv6 Scoped Address Architecture is supported on D, Dxp, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

D, Dxp, K5, and T(A) for Management only.

RFC 4193, Unique Local IPv6 Unicast Addresses is supported on D, Dxp, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

D, Dxp, K5, and T(A) for Management only.

RFC 4291, Internet Protocol Version 6 (IPv6) Addressing Architecture is supported on D, Dxp, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

D, Dxp, K5, and T(A) for Management only.

RFC 4443, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification is supported on D, Dxp, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

D, Dxp, K5, and T(A) for Management only.

RFC 4861, Neighbor Discovery for IP version 6 (IPv6) is supported on D, Dxp, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

D, Dxp, K5, and T(A) for Management only.

RFC 4862, IPv6 Stateless Address Autoconfiguration (Router Only) is supported on D, Dxp, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

D, Dxp, K5, and T(A) for Management only.

RFC 5095, Deprecation of Type 0 Routing Headers in IPv6 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5952, A Recommendation for IPv6 Address Text Representation is supported on D, Dxp, K5, K12, K30, T(A, N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

D, Dxp, K5, and T(A) for Management only.

RFC 6106, IPv6 Router Advertisement Options for DNS Configuration is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6164, Using 127-Bit IPv6 Prefixes on Inter-Router Links is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

12.10 IPsec

RFC 2401, Security Architecture for the Internet Protocol is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

Only for use with OSPFv3 authentication. Not supported for services.

RFC 2406, IP Encapsulating Security Payload (ESP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**

Only for use with OSPFv3 authentication. Not supported for services.

12.11 IS-IS

draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-segment-routing-extensions-04, IS-IS Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

draft-kaplan-isis-ext-eth-02, Extended Ethernet Frame Size Support is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ISO/IEC 10589:2002, Second Edition, Nov. 2002, Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1195, Use of OSI IS-IS for Routing in TCP/IP and Dual Environments is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3359, Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3719, Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3787, Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4971, Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5120, M-ISIS: Multi Topology (MT) Routing in IS-IS is supported on K12, T(N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5130, A Policy Control Mechanism in IS-IS Using Administrative Tags is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5301, Dynamic Hostname Exchange Mechanism for IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5302, Domain-wide Prefix Distribution with Two-Level IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5303, Three-Way Handshake for IS-IS Point-to-Point Adjacencies is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5304, IS-IS Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5305, IS-IS Extensions for Traffic Engineering TE is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5306, Restart Signaling for IS-IS (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5308, Routing IPv6 with IS-IS is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5310, IS-IS Generic Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6232, Purge Originator Identification TLV for IS-IS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6233, IS-IS Registry Extension for Purges is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-mi-02, IS-IS Multi-Instance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12



Note:

K12, K30 support only a single instance and can operate in multi-instance deployment as it supports the processing of TLVs for multi-instance support.

draft-ietf-isis-segment-routing-extensions-04, IS-IS Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

12.12 Management

draft-ietf-snmpv3-update-mib-05, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-idr-bgp4-mib-05, Definitions of Managed Objects for the Fourth Version of Border Gateway Protocol (BGP-4) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-isis-wg-mib-06, Management Information Base for Intermediate System to Intermediate System (IS-IS) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mboned-msdp-mib-01, Multicast Source Discovery protocol MIB is supported on Sx/S-1/10GE



Note:

Only in standalone mode.

draft-ietf-mpls-ldp-mib-07, Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mpls-lsr-mib-06, Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-mpls-te-mib-04, Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-mib-update-08, OSPF Version 2 Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaaddressfamilynumbers-mib, IANA-ADDRESS-FAMILY-NUMBERS-MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaiftype-mib, IANAifType-MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

ianaiprouteprotocol-mib, IANA-RTPROTO-MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8021-CFM-MIB, IEEE P802.1ag(TM) CFM MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8021-PAE-MIB, IEEE 802.1X MIB is supported on D, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

IEEE8023-LAG-MIB, IEEE 802.3ad MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

LLDP-MIB, IEEE P802.1AB(TM) LLDP MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1157, A Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1215, A Convention for Defining Traps for use with the SNMP is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 1724, RIP Version 2 MIB Extension is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2021, Remote Network Monitoring Management Information Base Version 2 using SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2138, Remote Authentication Dial In User Service (RADIUS) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2206, RSVP Management Information Base using SMIv2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2213, Integrated Services Management Information Base using SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2571, An Architecture for Describing SNMP Management Frameworks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2572, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2573, SNMP Applications is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2575, View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP) is supported on D, E, Dxp, K5, K12, K30, T(A,N), X, Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2578, Structure of Management Information Version 2 (SMIv2) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2579, Textual Conventions for SMIv2 is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2787, Definitions of Managed Objects for the Virtual Router Redundancy Protocol is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2819, Remote Network Monitoring Management Information Base is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2856, Textual Conventions for Additional High Capacity Data Types is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2863, The Interfaces Group MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2864, The Inverted Stack Table Extension to the Interfaces Group MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2933, Internet Group Management Protocol MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3014, Notification Log MIB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3164, The BSD syslog Protocol is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3165, Definitions of Managed Objects for the Delegation of Management Scripts is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3231, Definitions of Managed Objects for Scheduling Management Operations is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3273, Remote Network Monitoring Management Information Base for High Capacity Networks is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3414, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3416, Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3417, Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3419, Textual Conventions for Transport Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3584, Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3593, Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals is supported on Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3635, Definitions of Managed Objects for the Ethernet-like Interface Types is supported on D, E, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3826, The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3877, Alarm Management Information Base (MIB) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4001, Textual Conventions for Internet Network Addresses is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4022, Management Information Base for the Transmission Control Protocol (TCP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4113, Management Information Base for the User Datagram Protocol (UDP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4220, Traffic Engineering Link Management Information Base is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4292, IP Forwarding Table MIB is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4293, Management Information Base for the Internet Protocol (IP) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5101, Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information is supported on Mxp, Sx/S-1/10GE, and R6

RFC 6241, Network Configuration Protocol (NETCONF) is supported on K5, K12, T(N), Mxp, Sx/S-1/10GE, R6, and R12

RFC 6242, Using the NETCONF Protocol over Secure Shell (SSH) is supported on K5, K12, T(N), Mxp, Sx/S-1/10GE, R6, and R12

12.13 MPLS — General

RFC 3031, Multiprotocol Label Switching Architecture is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3032, MPLS Label Stack Encoding is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3443, Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4182, Removing a Restriction on the use of MPLS Explicit NULL is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5332, MPLS Multicast Encapsulations is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12

12.14 MPLS — GMPLS

draft-ietf-ccamp-rsvp-te-srlg-collect-04, RSVP-TE Extensions for Collecting SRLG Information is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

12.15 MPLS — LDP

draft-ietf-mpls-ldp-ipv6-15, Updates to LDP for IPv6 is supported on Mxp

draft-pdutta-mpls-ldp-adj-capability-00, LDP Adjacency Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-pdutta-mpls-ldp-v2-00, LDP Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-pdutta-mpls-ldp-hello-reduce-04, Targeted LDP Hello Reduction is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3037, LDP Applicability is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3478, Graceful Restart Mechanism for Label Distribution Protocol (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5036, LDP Specification is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5283, LDP Extension for Inter-Area Label Switched Paths (LSPs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5443, LDP IGP Synchronization is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5561, LDP Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6388, Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12



Note:

P2MP LSPs only.

12.16 MPLS — MPLS-TP

RFC 5586, MPLS Generic Associated Channel is supported on T(N), R6, and R12

RFC 5921, A Framework for MPLS in Transport Networks is supported on T(N), R6, and R12

RFC 5960, MPLS Transport Profile Data Plane Architecture is supported on T(N), R6, and R12

RFC 6370, MPLS Transport Profile (MPLS-TP) Identifiers is supported on T(N), R6, and R12

RFC 6378, MPLS Transport Profile (MPLS-TP) Linear Protection is supported on T(N), R6, and R12

RFC 6426, MPLS On-Demand Connectivity and Route Tracing is supported on T(N), R6, and R12

RFC 6428, Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile is supported on T(N), R6, and R12

RFC 6478, Pseudowire Status for Static Pseudowires is supported on T(N), R6, and R12

RFC 7213, MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing is supported on T(N), R6, and R12

12.17 MPLS — OAM

RFC 6424, Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6425, Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping is supported on T(N), Mxp, R6, and R12

12.18 MPLS — RSVP-TE

RFC 2702, Requirements for Traffic Engineering over MPLS is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2747, RSVP Cryptographic Authentication is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2961, RSVP Refresh Overhead Reduction Extensions is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3097, RSVP Cryptographic Authentication -- Updated Message Type Value is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3209, RSVP-TE: Extensions to RSVP for LSP Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3477, Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4090, Fast Reroute Extensions to RSVP-TE for LSP Tunnels is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4561, Definition of a Record Route Object (RRO) Node-Id Sub-Object is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4875, Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs) is supported on T(N), Mxp, Sx/S-1/10GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4950, ICMP Extensions for Multiprotocol Label Switching is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5817, Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

12.19 OSPF

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 1765, OSPF Database Overflow is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2328, OSPF Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3101, The OSPF Not-So-Stubby Area (NSSA) Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3509, Alternative Implementations of OSPF Area Border Routers is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3623, Graceful OSPF Restart Graceful OSPF Restart (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3630, Traffic Engineering (TE) Extensions to OSPF Version 2 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4222, Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4552, Authentication/Confidentiality for OSPFv3 is supported on K12, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4576, Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4577, OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 4970, Extensions to OSPF for Advertising Optional Router Capabilities is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5185, OSPF Multi-Area Adjacency is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5187, OSPFv3 Graceful Restart (Helper Mode) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

RFC 5243, OSPF Database Exchange Summary List Optimization is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5250, The OSPF Opaque LSA Option is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5309, Point-to-Point Operation over LAN in Link State Routing Protocols is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5340, OSPF for IPv6 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5838, Support of Address Families in OSPFv3 is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6987, OSPF Stub Router Advertisement is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

draft-ietf-ospf-prefix-link-attr-06, OSPFv2 Prefix/Link Attribute Advertisement is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12

draft-ietf-ospf-segment-routing-extensions-04, OSPF Extensions for Segment Routing is supported on K12, K30, Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

12.20 Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, OAM Procedures for VPWS Interworking is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3916, Requirements for Pseudo- Wire Emulation Edge-to-Edge (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3985, Pseudo Wire Emulation Edge-to-Edge (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4385, Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4446, IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4447, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5659, An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6073, Segmented Pseudowire is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6310, Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6391, Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network is supported on K12, K30, Mxp, R6, and R12

RFC 6718, Pseudowire Redundancy is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6870, Pseudowire Preferential Forwarding Status bit is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7023, MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 7267, Dynamic Placement of Multi-Segment Pseudowires is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

12.21 Quality of Service

RFC 2430, A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE) is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 2598, An Expedited Forwarding PHB is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3140, Per Hop Behavior Identification Codes is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 3260, New Terminology and Clarifications for Diffserv is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

12.22 RIP

RFC 1058, Routing Information Protocol is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2082, RIP-2 MD5 Authentication is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

RFC 2453, RIP Version 2 is supported on Mxp and Sx/S-1/10GE



Note:

Only in standalone mode.

12.23 Timing

GR-1244-CORE, Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

GR-253-CORE, SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

IEEE 1588-2008, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx-1/10GE, Sx-10/100GE, R6, and R12



Note:

For Dxp, IEEE default profile is supported only includes the Dxp-12p ETR, Dxp-16p, Dxp-24p. Sx-10/100GE does not support IEEE default profile.

ITU-T G.781, Synchronization layer functions, issued 09/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

**Note:**

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.813, Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

**Note:**

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8261, Timing and synchronization aspects in packet networks, issued 04/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

**Note:**

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8262, Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

**Note:**

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8264, Distribution of timing information through packet networks, issued 10/2008 is supported on D-ETR, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, R6, and R12

**Note:**

For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, and Dxp-24p.

ITU-T G.8265.1, Precision time protocol telecom profile for frequency synchronization, issued 10/2010 is supported on D-ETR, K5, K12, K30, T(A,N), Mxp, Sx-1/10GE, R6, and R12

ITU-T G.8275.1, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014 is supported on Dxp, K12, K30, T(A,N), Mxp, Sx-1/10GE, Sx-10/100GE, R6, and R12

**Note:**

For 7210 SAS-Sx 10/100GE, the support only includes the Sx 10/100GE QSFP28 variant. For Dxp, the support only includes the Dxp-12p ETR, Dxp-16p, Dxp-24p.

IEC/IEEE 61850-9-3-2016, Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation is supported on Dxp-16p and Dxp-24p

IEEE C37.238-2017 - IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications is supported on Dxp-16p and Dxp-24p

RFC 5905, Network Time Protocol Version 4: Protocol and Algorithms Specification is supported on D, Dxp, K5, K12, K30, T(A,N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

12.24 VPLS

RFC 4761, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

**Note:**


On 7210 platforms, only BGP-AD is supported with TLDP signalling for PW. No BGP signalling is supported for PW establishment.

RFC 4762, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 5501, Requirements for Multicast Support in Virtual Private LAN Services is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

RFC 6074, Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)