



7210 Service Access System

Release 23.9.R1

7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Services Guide

3HE 19291 AAAB TQZZA
Edition: 01
October 2023

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.

© 2023 Nokia.

Table of contents

List of tables	20
List of figures	25
1 Getting started	29
1.1 About this guide.....	29
1.1.1 Document structure and content.....	29
1.2 7210 SAS modes of operation.....	30
1.3 7210 SAS port modes.....	32
1.4 7210 SAS services configuration process.....	34
1.5 Conventions.....	35
1.5.1 Precautionary and information messages.....	35
1.5.2 Options or substeps in procedures and sequential workflows.....	36
2 Services overview	37
2.1 Introduction.....	37
2.1.1 Service types on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	37
2.1.2 Service types on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	38
2.1.3 Service policies on 7210 SAS-D and 7210 SAS-Dxp.....	38
2.2 Service policies on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP + 8C.....	39
2.3 Nokia service model on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	39
2.4 Nokia service model on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	40
2.5 Service entities on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	40
2.6 Service entities on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	40
2.6.1 Customers.....	41
2.7 SAPs on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	41
2.8 SAPs on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	42
2.9 SAP encapsulation types and identifiers.....	43
2.9.1 Ethernet encapsulations supported on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.....	43
2.9.2 Services and SAP encapsulations.....	44
2.9.3 Default SAPs on a dot1q port.....	44
2.9.4 Default SAPs on a QinQ port.....	45

2.9.4.1	Configuration notes for default QinQ SAPs for transit service in a ring deployment.....	47
2.9.5	SAP configuration considerations applicable for SAPs configured on access-uplink ports.....	47
2.9.5.1	Configuration guidelines for 7210 SAS-D and 7210 SAS-Dxp.....	49
2.9.5.2	Configuration guidelines for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.....	51
2.10	VLAN range for SAPs (or dot1q range SAPs) in an Epipe or VPLS service.....	52
2.10.1	Epipe processing behavior for dot1q range SAPs in access-uplink mode on the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	52
2.10.1.1	Epipe VLAN range SAPs feature support and restrictions.....	53
2.10.1.2	Epipe processing behavior for dot1q range SAPs on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	54
2.10.2	VLAN range SAPs in a VPLS service on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	55
2.10.2.1	VPLS processing behavior for dot1q range SAPs on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	56
2.10.2.2	VPLS VLAN range SAP feature support and restrictions.....	56
2.10.3	Epipe emulation using dot1q VLAN range SAP in VPLS with G.8032.....	58
2.10.3.1	Epipe emulation using dot1q VLAN range SAP in VPLS with G.8032 — configuration guidelines and restrictions.....	59
2.11	SDPs on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	60
2.11.1	SDP binding.....	61
2.11.2	Spoke and mesh SDPs.....	61
2.11.3	SDP using BGP route tunnel.....	61
2.11.4	SDP keepalives.....	62
2.11.5	SDP administrative groups.....	62
2.11.6	Mixed-LSP mode of operation.....	63
2.12	G.8032/Ethernet ring protection switching.....	65
2.12.1	Overview of G.8032 operation.....	65
2.12.2	G.8032 Ethernet ring sub-rings.....	69
2.12.3	G.8032 virtual and non-virtual channel.....	71
2.12.3.1	G.8032 Ethernet ring sub-ring using non-virtual link.....	72
2.12.4	G.8032 OAM considerations.....	75
2.12.5	G.8032 QoS considerations.....	75
2.12.6	G.8032 service and solution combinations.....	76
2.12.7	G.8032 configuration guidelines.....	76
2.13	Layer 2 control processing.....	76

2.14	Service creation process overview.....	78
2.14.1	Deploying and provisioning services.....	79
2.14.1.1	Phase 1: core network construction.....	79
2.14.1.2	Phase 2: service administration.....	79
2.14.1.3	Phase 3: service provisioning.....	80
2.14.2	Configuration notes.....	80
2.14.2.1	General.....	80
2.15	Configuring global service entities with CLI.....	80
2.15.1	Service model entities.....	80
2.16	Basic configuration.....	80
2.17	Common configuration tasks.....	81
2.17.1	Configuring customer accounts.....	81
2.17.1.1	Customer information.....	81
2.18	Service creation process overview on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C....	81
2.18.1	Deploying and provisioning services on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	82
2.18.1.1	Phase 1: core network construction.....	82
2.18.1.2	Phase 2: service administration.....	82
2.18.1.3	Phase 3: service provisioning.....	83
2.18.2	Configuration notes.....	83
2.18.2.1	General.....	83
2.18.3	Configuring global service entities with CLI.....	83
2.18.3.1	Service model entities.....	83
2.18.4	Basic configuration.....	84
2.18.5	Common configuration tasks.....	84
2.18.5.1	Configuring customer accounts.....	84
2.18.5.2	Customer information.....	85
2.18.5.3	Configuring an SDP.....	85
2.18.5.4	SDP configuration tasks.....	85
2.18.6	Configuring an SDP.....	86
2.18.7	Configuring a mixed-LSP SDP.....	87
2.19	ETH-CFM.....	87
2.19.1	Common actionable failures.....	90
2.19.2	MEP and MIP support.....	91
2.19.3	Configuring ETH-CFM parameters.....	91
2.20	Service management tasks.....	93

2.20.1	Modifying customer accounts.....	93
2.20.2	Deleting customers.....	94
2.21	Global services command reference.....	94
2.21.1	Command hierarchies.....	94
2.21.1.1	Customer commands.....	94
2.21.1.2	Pseudowire (PW) commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	95
2.21.1.3	Pseudowire (PW) routing commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	95
2.21.1.4	SDP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	96
2.21.1.5	SAP commands for 7210 SAS-D and 7210 SAS-Dxp.....	96
2.21.1.6	SAP commands for 7210 SAS-K 2F1C2T.....	97
2.21.1.7	SAP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	97
2.21.1.8	ETH-CFM configuration commands.....	97
2.21.1.9	Show commands.....	98
2.21.1.10	Tools perform commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP + 8C.....	99
2.21.2	Command descriptions.....	99
2.21.2.1	Global service configuration commands.....	99
2.21.2.2	Show, clear, debug commands.....	138
3	VLL services.....	180
3.1	Epipe.....	180
3.2	Epipe service overview.....	180
3.3	Epipe oper state decoupling on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.....	181
3.4	Dynamic Multi-Segment Pseudowire routing.....	182
3.4.1	MS-PW routing — pseudowire routing.....	186
3.4.1.1	MS-PW routing — static routing.....	187
3.4.1.2	MS-PW routing — explicit paths.....	187
3.4.2	MS-PW routing — configuring VLLs using dynamic MS-PWs.....	188
3.4.2.1	MS-PW routing — active/passive T-PE selection.....	188
3.4.2.2	MS-PW routing — automatic endpoint configuration.....	188
3.4.2.3	MS-PW routing — selecting a path for an MS-PW.....	189
3.4.2.4	MS-PW routing — pseudowire templates.....	189
3.4.3	MS-PW routing — pseudowire redundancy.....	189
3.4.4	MS-PW routing — VCCV OAM for dynamic MS-PWs.....	190

3.4.5	MS-PW routing — VCCV-ping on dynamic MS-PWs.....	191
3.4.6	MS-PW routing — VCCV-trace on dynamic MS-PWs.....	191
3.4.7	MS-PW routing — example dynamic MS-PW configuration.....	191
3.5	Master-slave operation.....	194
3.5.1	PW redundancy — operation of master-slave pseudowire redundancy with existing scenarios.....	196
3.5.1.1	PW redundancy — VLL resilience.....	196
3.5.2	VLL resilience for a switched pseudowire path.....	197
3.5.2.1	PW redundancy — VLL resilience for a switched PW path.....	199
3.5.3	Access node resilience using MC-LAG and pseudowire redundancy.....	200
3.5.4	Epipe pseudowire switching on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C....	202
3.5.4.1	Pseudowire switching with protection.....	203
3.5.4.2	Pseudowire switching behavior.....	204
3.5.4.3	Pseudowire redundancy.....	206
3.5.5	VLL service resilience with pseudowire on 7210 SAS-K 3SFP+ 8C.....	206
3.5.5.1	VLL resilience with two destination PE nodes.....	206
3.5.5.2	Operation of master-slave pseudowire redundancy with existing scenarios.....	209
3.5.5.3	VLL resilience for a switched pseudowire path.....	210
3.5.5.4	Pseudowire redundancy service models.....	211
3.5.5.5	Processing endpoint SAP active/standby status bits.....	213
3.5.5.6	Processing and merging.....	213
3.5.6	MPLS entropy label and hash label.....	215
3.6	VLL service considerations.....	215
3.6.1	SDPs.....	215
3.6.2	SAP encapsulations.....	215
3.6.3	QoS policies.....	215
3.6.4	Filter policies.....	216
3.6.5	MAC resources.....	216
3.7	Configuring a VLL service with CLI.....	216
3.7.1	Basic configurations.....	216
3.7.2	Common configuration tasks.....	216
3.7.2.1	Creating an Epipe service for 7210 SAS-D.....	216
3.7.2.2	Creating an Epipe service for 7210 SAS-Dxp.....	217
3.7.2.3	Creating an Epipe service with range SAPs.....	217
3.7.2.4	Creating an Epipe service for 7210 SAS-K 2F1C2T.....	218
3.7.2.5	Creating an Epipe service for 7210 SAS-K 2F1C2T with range SAPs.....	218

3.7.2.6	Configuring default QinQ SAPs for Epipe transit traffic in a ring scenario in access-uplink mode.....	227
3.7.3	Service management tasks.....	229
3.7.3.1	Modifying Epipe service parameters.....	229
3.7.3.2	Disabling an Epipe service.....	229
3.7.3.3	Re-enabling an Epipe service.....	229
3.7.3.4	Deleting an Epipe service.....	230
3.8	VLL services command reference.....	230
3.8.1	Command hierarchy.....	230
3.8.1.1	Epipe service configuration commands.....	230
3.8.1.2	Connection profile commands.....	236
3.8.1.3	Show commands.....	236
3.8.1.4	Clear commands.....	236
3.8.2	Command descriptions.....	236
3.8.2.1	VLL service configuration commands.....	236
4	Ethernet Virtual Private Networks.....	299
4.1	EVPN applications.....	299
4.1.1	EVPN for MPLS tunnels in E-LAN services.....	299
4.2	EVPN for MPLS tunnels.....	300
4.2.1	BGP-EVPN control plane for MPLS tunnels.....	300
4.2.1.1	EVPN route type 3 — inclusive multicast Ethernet tag route.....	301
4.2.1.2	EVPN route type 2 — MAC/IP advertisement route.....	302
4.2.1.3	EVPN route type 1 — Ethernet Auto-Discovery route.....	303
4.2.1.4	EVPN route type 4 — ES route.....	304
4.2.1.5	BGP tunnel encapsulation extended community.....	304
4.2.2	EVPN for MPLS tunnels in VPLS services.....	304
4.2.2.1	EVPN and VPLS integration.....	307
4.2.2.2	Auto-derived RD in services with multiple BGP families.....	310
4.2.3	EVPN multi-homing in VPLS services.....	310
4.2.4	EVPN all-active multi-homing.....	311
4.2.4.1	All-active multi-homing procedures.....	311
4.2.4.2	All-active multi-homing service model.....	313
4.2.4.3	ES discovery and DF election procedures.....	314
4.2.4.4	Aliasing.....	319
4.2.4.5	Network failures and convergence for all-active multi-homing.....	321

4.2.4.6	Logical failures on ESs and black holes.....	322
4.2.4.7	Transient issues due to MAC route delays.....	323
4.2.5	EVPN single-active multi-homing.....	323
4.2.5.1	Single-active multi-homing service model.....	324
4.2.5.2	ES and DF election procedures.....	326
4.2.5.3	Backup PE function.....	327
4.2.5.4	Network failures and convergence for single-active multi-homing.....	328
4.3	General EVPN topics.....	329
4.3.1	ARP and ND snooping and proxy support.....	329
4.3.1.1	Proxy-ARP/ND periodic refresh, unsolicited refresh, and confirm-messages.....	333
4.3.1.2	Proxy-ND and the Router flag in neighbor advertisement messages.....	333
4.3.1.3	Procedure to add the R flag to a specified entry.....	334
4.3.2	BGP-EVPN MAC mobility.....	334
4.3.3	BGP-EVPN MAC duplication.....	335
4.3.4	Conditional static MAC and protection.....	336
4.3.5	BGP and EVPN route selection for EVPN routes.....	337
4.3.6	EVPN interaction with other features.....	338
4.3.6.1	EVPN-MPLS with existing VPLS features.....	338
4.3.6.2	EVPN with G.8032 in an access ring.....	339
4.3.7	Routing policies for BGP EVPN routes.....	344
4.4	Configuring an EVPN service with CLI.....	345
4.4.1	EVPN-MPLS configuration examples.....	345
4.4.1.1	EVPN all-active multi-homing example.....	345
4.4.1.2	EVPN single-active multi-homing example.....	347
4.5	EVPN command reference.....	348
4.5.1	Command hierarchies.....	348
4.5.1.1	EVPN configuration commands.....	348
4.5.1.2	EVPN show commands.....	350
4.5.1.3	EVPN clear commands.....	350
4.5.1.4	EVPN tools commands.....	350
4.5.2	Command descriptions.....	351
4.5.2.1	EVPN configuration commands.....	351
4.5.2.2	EVPN show commands.....	388
4.5.2.3	EVPN clear commands.....	399
4.5.2.4	Tools commands.....	400

5	Virtual Private LAN Service.....	403
5.1	VPLS service overview.....	403
5.1.1	VPLS packet walk-through for a local service.....	403
5.1.2	VPLS packet walk-through for a distributed service.....	405
5.2	VPLS features.....	408
5.2.1	VPLS enhancements on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	408
5.2.2	VPLS enhancements on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	408
5.2.3	VPLS over MPLS - supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C..	409
5.2.4	VPLS over SAPs.....	409
5.2.5	VPLS MAC learning and packet forwarding.....	409
5.2.5.1	IGMP snooping on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T....	410
5.2.5.2	IGMP snooping on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	411
5.2.5.3	Multicast VLAN Registration (MVR) support.....	412
5.2.6	DHCPv4 snooping.....	413
5.2.7	L2 forwarding table management.....	413
5.2.7.1	FIB size on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	414
5.2.7.2	FIB size on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	414
5.2.7.3	FIB size alarms.....	414
5.2.7.4	Local aging timers on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T..	415
5.2.7.5	Local and remote aging timers on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP + 8C.....	415
5.2.7.6	Disable MAC aging.....	415
5.2.7.7	Disable MAC learning.....	415
5.2.7.8	Unknown MAC discard.....	415
5.2.7.9	VPLS and rate limiting.....	416
5.2.7.10	MAC move.....	416
5.2.7.11	Split horizon SAP groups.....	416
5.2.8	VPLS and Spanning Tree Protocol.....	417
5.2.8.1	Spanning tree operating modes.....	417
5.2.8.2	Multiple Spanning Tree Protocol.....	418
5.2.8.3	MSTP for QinQ SAPs.....	419
5.2.8.4	Enhancements to the Spanning Tree Protocol in MPLS-based VPLS service (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C Only).....	420
5.2.9	VPLS access redundancy.....	422
5.2.9.1	STP-based redundant access to VPLS.....	422

5.2.10	VPLS redundancy on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C using MPLS uplinks.....	423
5.2.10.1	Spoke-SDP redundancy on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	423
5.2.10.2	Spoke-SDP-based redundant access on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	424
5.2.10.3	Inter-domain VPLS resiliency using multi-chassis endpoints on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	424
5.2.11	VPLS access redundancy when using MPLS uplinks on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	425
5.2.11.1	STP-based redundant access to VPLS using PWs on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	425
5.2.11.2	Redundant access to VPLS without STP and using MPLS uplinks on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	426
5.2.11.3	MAC flush message processing in VPLS services with MPLS uplinks on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	427
5.2.11.4	Dual homing to a VPLS service on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	428
5.2.12	VPLS service considerations.....	429
5.2.12.1	SAP encapsulations.....	429
5.2.12.2	VLAN processing.....	429
5.3	BGP-AD VPLS.....	430
5.3.1	BGP-AD and Target LDP (T-LDP) interaction on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	430
5.3.2	BGP-AD VPLS - SDP usage.....	431
5.3.3	BGP-AD VPLS - automatic creation of SDPs.....	432
5.3.4	BGP-AD VPLS - manually provisioned SDP.....	432
5.3.5	BGP-AD VPLS - automatic instantiation of pseudowires (SDP bindings).....	432
5.3.6	BGP-AD VPLS - mixing statically configured and auto-discovered pseudowires in a VPLS service.....	433
5.3.7	BGP-AD VPLS - resiliency schemes.....	433
5.4	Routed VPLS.....	434
5.4.1	IES IP interface binding.....	434
5.4.2	Assigning a service name to a VPLS service.....	434
5.4.3	Service binding requirements.....	435
5.4.3.1	Bound service name assignment.....	435
5.4.3.2	Binding a service name to an IP interface.....	435
5.4.3.3	IP interface attached VPLS service constraints.....	435
5.4.3.4	IP interface and VPLS operational state coordination.....	436

5.4.3.5	IP interface MTU and fragmentation on 7210 SAS-D and 7210 SAS-Dxp.....	436
5.4.3.6	IP interface MTU and fragmentation on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.....	436
5.4.4	ARP and VPLS FIB interactions.....	437
5.4.5	Specific ARP cache behavior.....	437
5.4.5.1	The allow-ip-int-binding VPLS flag.....	438
5.4.6	SAPs only supported on standard Ethernet ports.....	438
5.4.6.1	LAG port membership constraints.....	438
5.4.6.2	VPLS feature support and restrictions.....	438
5.4.7	VPLS SAP ingress IP filter override on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.....	439
5.4.7.1	QoS support for VPLS SAPs and IP interface.....	442
5.4.7.2	Routing related protocols on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	442
5.4.7.3	Spanning Tree and split horizon.....	443
5.4.8	R-VPLS support and caveats.....	443
5.5	Configuring a VPLS service with CLI.....	444
5.5.1	Basic configuration.....	444
5.5.2	Common configuration tasks.....	445
5.5.3	Configuring VPLS components.....	445
5.5.3.1	Creating a VPLS service.....	445
5.5.3.2	Configuring a VPLS SAP.....	452
5.5.4	Configuring VPLS redundancy.....	461
5.5.4.1	Creating a management VPLS for SAP protection.....	462
5.5.4.2	Configuring load balancing with management VPLS.....	463
5.5.4.3	Configuring a BGP Auto-Discovery on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	465
5.5.4.4	Configuration steps.....	466
5.5.5	BGP- AD- LDP signaling.....	468
5.5.6	BGP AD- pseudowire template.....	470
5.5.7	Configuring a VPLS management interface.....	471
5.5.7.1	Creating a management VPLS for SAP protection using MPLS uplinks.....	472
5.5.7.2	Creating a management VPLS for spoke-SDP protection, using MPLS uplinks (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C).....	473
5.5.7.3	Configuring load balancing with management VPLS using MPLS uplinks (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C).....	474
5.5.7.4	Configuring selective MAC Flush, when using MPLS uplinks (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C).....	478

5.5.7.5	Configuring AS pseudowire in VPLS when using MPLS uplinks (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C).....	478
5.6	Service management tasks.....	479
5.6.1	Modifying VPLS service parameters.....	479
5.6.2	Modifying management VPLS parameters.....	480
5.6.3	Deleting a management VPLS.....	480
5.6.4	Disabling a management VPLS.....	480
5.6.5	Deleting a VPLS service.....	480
5.6.6	Disabling a VPLS service.....	481
5.6.7	Re-enabling a VPLS service.....	481
5.7	VPLS services command reference.....	481
5.7.1	Command hierarchies.....	481
5.7.1.1	VPLS service configuration commands.....	482
5.7.1.2	VPLS service xSTP commands.....	484
5.7.1.3	VPLS service SAP xSTP commands.....	485
5.7.1.4	VPLS service SAP DHCP snooping commands.....	486
5.7.1.5	SAP commands.....	486
5.7.1.6	VPLS SAP QoS and filter commands for 7210 SAS-D and 7210 SAS-Dxp.....	487
5.7.1.7	VPLS SAP QoS and filter commands for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.....	488
5.7.1.8	VPLS service SAP IGMP snooping and MVR commands.....	488
5.7.1.9	VPLS SAP meter override commands.....	489
5.7.1.10	VPLS SAP statistics commands for 7210 SAS-D.....	490
5.7.1.11	VPLS mesh SDP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	490
5.7.1.12	VPLS spoke-SDP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	491
5.7.1.13	Routed VPLS commands for 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	493
5.7.1.14	Show commands.....	493
5.7.1.15	Clear commands.....	494
5.7.1.16	Debug commands.....	494
5.7.2	Command descriptions.....	494
5.7.2.1	VPLS service configuration commands.....	494
5.7.2.2	VPLS show commands.....	598
5.7.2.3	IGMP snooping show commands.....	640
5.7.2.4	VPLS clear commands.....	655
5.7.2.5	VPLS debug commands.....	660

6	Internet Enhanced Service on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	666
6.1	IES service overview.....	666
6.2	IES features.....	667
6.2.1	IP interfaces.....	667
6.2.2	IPv6 support for IES IP interfaces associated with access-uplink SAPs (on 7210 SAS-D and 7210 SAS-Dxp).....	667
6.3	SAPs.....	668
6.3.1	Encapsulations.....	668
6.3.2	CPE connectivity check.....	668
6.3.3	QoS policies.....	668
6.3.4	CPU QoS for IES interfaces in access-uplink mode.....	669
6.3.5	Filter policies.....	669
6.4	Configuring an IES service with CLI.....	669
6.4.1	Basic configuration.....	669
6.4.2	Common configuration tasks.....	670
6.4.3	Configuring IES components.....	671
6.4.3.1	Configuring an IES service.....	671
6.4.3.2	Configuring IES interface parameters.....	671
6.4.3.3	Configuring IES SAP parameters.....	671
6.4.3.4	Configuring SAP parameters.....	672
6.4.4	Service management tasks.....	672
6.4.4.1	Modifying IES service parameters.....	673
6.4.4.2	Deleting an IES service.....	673
6.4.4.3	Disabling an IES service.....	673
6.4.4.4	Re-enabling an IES service.....	674
6.5	IES services command reference.....	674
6.5.1	Command hierarchies.....	674
6.5.1.1	Global IES commands for 7210 SAS-D and 7210 SAS-Dxp.....	674
6.5.1.2	IES interface commands.....	674
6.5.1.3	IES interface routed VPLS commands for 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.....	675
6.5.1.4	IES interface SAP commands for 7210 SAS-D and 7210 SAS-Dxp.....	676
6.5.1.5	IES interface SAP statistics commands for 7210 SAS-D and 7210 SAS-Dxp....	676
6.5.1.6	IES interface SAP QoS and filter commands for 7210 SAS-D and 7210 SAS-Dxp.....	676

6.5.1.7	IES interface IPv6 commands (applicable only to access-uplink SAPs on 7210 SAS-D and 7210 SAS-Dxp).....	677
6.5.1.8	Show commands.....	677
6.5.2	Command descriptions.....	678
6.5.2.1	IES service configuration commands.....	678
7	Internet Enhanced Service on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	727
7.1	IES service overview.....	727
7.2	IES features.....	727
7.2.1	IP interfaces.....	728
7.3	SAPs.....	728
7.3.1	IPv6 support for IES IP interfaces.....	728
7.3.2	Encapsulations.....	728
7.3.3	Routing protocols.....	728
7.3.3.1	CPE connectivity check.....	729
7.3.4	QoS policies.....	729
7.3.4.1	CPU QoS for IES access interfaces in network mode.....	729
7.3.5	Filter policies.....	729
7.3.6	VRRP support for IES IP interfaces in network mode.....	730
7.4	Configuring an IES service with CLI.....	730
7.4.1	Basic configuration.....	730
7.4.2	Common configuration tasks.....	730
7.4.2.1	Configuring IES components.....	731
7.4.3	Service management tasks.....	733
7.4.3.1	Modifying IES service parameters.....	733
7.4.3.2	Deleting an IES service.....	734
7.4.3.3	Disabling an IES service.....	734
7.4.3.4	Re-enabling an IES service.....	734
7.5	IES services command reference for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	734
7.5.1	Command hierarchies.....	734
7.5.1.1	Global commands.....	735
7.5.1.2	Interface commands.....	735
7.5.1.3	VRRP commands (applicable only for network mode).....	736
7.5.1.4	Routed VPLS commands.....	736
7.5.1.5	IES SAP configuration — QoS and filter commands.....	737
7.5.1.6	Interface IPv6 commands.....	737

7.5.1.7	Show commands.....	737
7.5.2	Command descriptions.....	738
7.5.2.1	IES service configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	738
7.5.2.2	IES show commands.....	790
8	Virtual Private Network service on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	807
8.1	VPRN service overview.....	807
8.1.1	Routing prerequisites.....	808
8.1.2	BGP support.....	808
8.1.3	Route distinguisher.....	808
8.1.4	Route reflector.....	809
8.1.5	CE to PE route exchange.....	809
8.1.5.1	Route redistribution.....	810
8.1.5.2	CPE connectivity check.....	810
8.1.6	Constrained route distribution.....	811
8.1.6.1	Constrained VPN route distribution based on route targets.....	811
8.1.6.2	Configuring the route target address family.....	811
8.1.6.3	Originating RT constraint routes.....	812
8.1.6.4	Receiving and re-advertising RT constraint routes.....	812
8.1.6.5	Using RT constraint routes.....	813
8.2	VPRN features.....	814
8.2.1	IP interfaces.....	815
8.2.2	SAPs.....	815
8.2.2.1	IPv6 support for VPRN IP interfaces (6VPE) in network mode.....	815
8.2.2.2	Encapsulations.....	815
8.2.2.3	QoS policies for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	815
8.2.2.4	Filter policies.....	816
8.2.2.5	CPU QoS for VPRN interfaces.....	816
8.2.3	CE to PE routing protocols.....	816
8.2.3.1	PE to PE tunneling mechanisms.....	816
8.2.3.2	Per VRF route limiting.....	817
8.2.4	Spoke-SDP termination.....	817
8.2.4.1	Using OSPF in IP-VPNs.....	817
8.2.5	Service label mode of a VPRN.....	818
8.2.5.1	Inter-AS VPRNs.....	818

8.2.5.2	VRRP support for VPRN IP interfaces in network mode.....	820
8.3	Configuring a VPRN service with CLI.....	820
8.3.1	Basic configuration.....	820
8.3.2	Common configuration tasks.....	821
8.3.3	Configuring VPRN components.....	822
8.3.3.1	Creating a VPRN service.....	822
8.3.3.2	Configuring global VPRN parameters.....	822
8.3.3.3	Configuring router interfaces.....	823
8.3.3.4	Configuring VPRN protocols - BGP.....	823
8.3.3.5	Configuring VPRN BGP group and neighbor parameters.....	824
8.3.3.6	Configuring a VPRN interface.....	825
8.3.3.7	Configuring a VPRN interface SAP.....	826
8.3.3.8	Configuring VRRP.....	827
8.3.3.9	Configuring IPv6 parameters for VPRN BGP.....	827
8.3.3.10	Configuring VPRN IPv6 neighbor discovery parameters.....	828
8.3.3.11	Configuring OSPF for VPRN.....	828
8.3.4	Service management tasks.....	828
8.3.4.1	Modifying VPRN service parameters.....	829
8.3.4.2	Deleting a VPRN service.....	830
8.3.4.3	Disabling a VPRN service.....	830
8.3.4.4	Re-enabling a VPRN service.....	831
8.4	VPRN services command reference for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C...	831
8.4.1	Command hierarchies.....	831
8.4.1.1	VPRN service configuration commands.....	832
8.4.1.2	DHCP server commands.....	833
8.4.1.3	IPoE commands.....	834
8.4.1.4	Interface commands.....	835
8.4.1.5	IPv6 interface commands (supported only on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C).....	836
8.4.1.6	IPv6 router advertisement commands (supported only on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C).....	837
8.4.1.7	Interface VRRP commands (for network mode only).....	838
8.4.1.8	Routed VPLS commands for VPRN service.....	838
8.4.1.9	Interface SAP commands.....	838
8.4.1.10	BGP configuration commands.....	839
8.4.1.11	OSPF configuration commands.....	842
8.4.1.12	Show commands.....	844

8.4.1.13	Clear commands.....	846
8.4.1.14	Debug commands.....	846
8.4.2	Command descriptions.....	846
8.4.2.1	VPRN service configuration commands.....	847
8.4.2.2	VPRN interface IPv6 commands.....	969
8.4.2.3	BGP commands.....	978
8.4.2.4	OSPF commands.....	1008
8.4.2.5	VPRN show commands.....	1043
8.4.2.6	VPRN clear commands.....	1133
8.4.2.7	VPRN debug commands.....	1140
9	Common CLI command descriptions.....	1147
9.1	Command descriptions.....	1147
9.1.1	SAP syntax.....	1147
	sap.....	1147
10	Appendix: port-based split horizon.....	1149
10.1	Overview.....	1149
10.1.1	Topology.....	1149
10.2	Configuration guidelines.....	1150
10.2.1	Verification.....	1151
11	Standards and protocol support.....	1153
11.1	BGP.....	1153
11.2	Ethernet.....	1155
11.3	EVPN.....	1156
11.4	Fast Reroute.....	1156
11.5	Internet Protocol (IP) — General.....	1157
11.6	IP — Multicast.....	1159
11.7	IP — Version 4.....	1160
11.8	IP — Version 6.....	1161
11.9	IPsec.....	1162
11.10	IS-IS.....	1162
11.11	Management.....	1164
11.12	MPLS — General.....	1167
11.13	MPLS — GMPLS.....	1167

11.14	MPLS — LDP.....	1168
11.15	MPLS — MPLS-TP.....	1168
11.16	MPLS — OAM.....	1169
11.17	MPLS — RSVP-TE.....	1169
11.18	OSPF.....	1169
11.19	Pseudowire.....	1170
11.20	Quality of Service.....	1171
11.21	RIP.....	1172
11.22	Timing.....	1172
11.23	VPLS.....	1173

List of tables

Table 1: Supported modes of operation and configuration methods.....	31
Table 2: Supported port modes by mode of operation.....	33
Table 3: 7210 SAS platforms supporting port modes.....	34
Table 4: Configuration process.....	35
Table 5: Service and SAP encapsulation for access ports.....	44
Table 6: Port type and encapsulation for access-uplink ports.....	44
Table 7: SAP and service combinations for 7210 SAS-D and 7210 SAS-Dxp.....	49
Table 8: SAP and service combinations for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.....	52
Table 9: L2CP support on the 7210 SAS.....	77
Table 10: Service endpoint and MEP/MIP acronyms.....	88
Table 11: ETH-CC defect condition groups.....	90
Table 12: Output fields: customer.....	140
Table 13: Output fields: FDB MAC.....	142
Table 14: Output fields: SDP-using.....	143
Table 15: Output fields: service-using.....	145
Table 16: Output fields: eth-ring status.....	147
Table 17: Output fields: eth-ring.....	148
Table 18: Output fields: PW template using.....	153
Table 19: Output fields: service SDP.....	160
Table 20: Output fields: association.....	164
Table 21: Output fields: CFM stack-table.....	166

Table 22: Output fields: domain.....	167
Table 23: Output fields: MEP.....	174
Table 24: Output fields: show connection profile.....	179
Table 25: MTU values.....	245
Table 26: Final disposition of the packet based on per-FC and per-SAP policer or meter.....	265
Table 27: Dot1p values extracted from the packet on SAP (ingress) when dot1p-inner and dot1p-outer commands are configured.....	268
Table 28: Dot1p values marked in the packet on SAP egress when dot1p-inner and dot1p-outer commands configured.....	269
Table 29: EVPN routes and usage.....	301
Table 30: Proxy-ARP entry combinations.....	333
Table 31: Output fields: EVPN MPLS tunnel endpoints.....	388
Table 32: Output fields: service ID BGP-EVPN.....	390
Table 33: Output fields: EVPN MPLS.....	392
Table 34: Output fields: proxy-ARP.....	393
Table 35: Output fields: proxy-ND.....	396
Table 36: Output fields: system BGP-EVPN.....	398
Table 37: Output fields: BGP-EVPN multi-homing.....	399
Table 38: Routing behavior in R-VPLS and interaction ARP cache and MAC FIB.....	438
Table 39: ACL lookup behavior with ingress override filter attached to an IES interface in an R-VPLS service.....	440
Table 40: ACL lookup behavior without ingress override filter attached to an IES interface in a R-VPLS service.....	441
Table 41: Routing protocols on IPv4 interfaces bound to a VPLS service.....	442
Table 42: SAP BPDU encapsulation states.....	459

Table 43: MTU values for VC types.....	523
Table 44: Final disposition of the packet based on per-FC and per-SAP policer or meter.....	567
Table 45: Output fields: FDB-info.....	599
Table 46: Output fields: FDB-MAC.....	601
Table 47: Output fields: SAP-using.....	603
Table 48: Output fields: service-using.....	606
Table 49: Output fields: service ID all.....	613
Table 50: Output fields: ARP.....	618
Table 51: Output fields: base.....	620
Table 52: Output fields: BGP-AD.....	622
Table 53: Output fields: FDB.....	623
Table 54: Output fields: L2PT.....	626
Table 55: Output fields: SAP.....	630
Table 56: Output fields: STP.....	634
Table 57: Output Fields: MSTP configuration.....	636
Table 58: Output fields: DHCP statistics.....	638
Table 59: Output fields: DHCP summary.....	639
Table 60: Output fields: IGMP snooping all.....	642
Table 61: Output fields: service MFIB.....	644
Table 62: Output fields: IGMP-snooping mrouter.....	645
Table 63: Output fields: IGMP snooping port DB.....	648
Table 64: Output fields: IGMP snooping proxy-database.....	650
Table 65: Output fields: IGMP snooping queriers.....	651

Table 66: Output fields: IGMP snooping static.....	653
Table 67: Output fields: IGMP-snooping statistics.....	654
Table 68: Final disposition of the packet based on per-FC and per-SAP policer or meter.....	714
Table 69: Final disposition of the packet based on per-FC and per-SAP policer or meter.....	768
Table 70: Output fields: customer.....	793
Table 71: Output fields: SAP-using.....	795
Table 72: Output fields: service-using.....	796
Table 73: Output fields: ID all.....	798
Table 74: Output fields: ARP.....	802
Table 75: Output fields: base.....	803
Table 76: Output fields: service ID interface.....	805
Table 77: Route preference defaults by route type.....	1031
Table 78: Output fields: egress label.....	1044
Table 79: Output fields: ingress label.....	1046
Table 80: Output fields: SAP-using.....	1048
Table 81: Output fields: SDP.....	1050
Table 82: Output fields: SDP-using.....	1054
Table 83: Output fields: service-using.....	1056
Table 84: Output fields: service ID all.....	1066
Table 85: Output fields: ARP.....	1073
Table 86: Output fields: base.....	1074
Table 87: Output fields: DHCP statistics.....	1076
Table 88: Output fields: interface.....	1079

Table 89: Output fields: SAP.....	1082
Table 90: Output fields: SDP.....	1087
Table 91: Output fields: aggregate.....	1089
Table 92: Output fields: router ARP.....	1091
Table 93: Output fields: BGP damping.....	1096
Table 94: Output fields: BGP group.....	1098
Table 95: Output fields: BGP neighbor.....	1104
Table 96: Output fields: neighbor received routes.....	1108
Table 97: Output fields: show neighbor add-path.....	1109
Table 98: Output fields: BGP paths.....	1114
Table 99: Output fields: BGP routes.....	1118
Table 100: Output fields: BGP summary.....	1120
Table 101: Output fields: detailed IP interface.....	1123
Table 102: Output fields: interface summary.....	1125
Table 103: Output fields: route table.....	1127
Table 104: Output fields: static ARP.....	1129
Table 105: Output fields: static route.....	1132
Table 106: Output fields: tunnel table.....	1133
Table 107: SAP-ID formats.....	1147
Table 108: Encapsulation types.....	1148

List of figures

Figure 1: Multiple SAPs in a service using QinQ uplinks in 7210 SAS configured in access-uplink mode...	42
Figure 2: Multiple SAPs for 7210 SAS configured with MPLS uplinks.....	42
Figure 3: Multiple SAPs in a service using QinQ uplinks in 7210 SAS configured in access-uplink mode...	43
Figure 4: Multiple SAPs on a single port.....	44
Figure 5: Epipe emulation in a ring using VPLS with G.8032.....	58
Figure 6: MPLS SDP pointing from ALA-A to ALA-B.....	61
Figure 7: G.8032 ring in the initial state.....	66
Figure 8: G.8032 ring in the protecting state.....	66
Figure 9: 0 to 3 ring example.....	68
Figure 10: Major ring and sub-ring scenario.....	70
Figure 11: G.8032 sub-ring.....	71
Figure 12: 0 to 6 sub-ring homed to VPLS.....	73
Figure 13: Service creation and implementation flow.....	79
Figure 14: Service creation and implementation flow.....	82
Figure 15: Ethernet OAM model for Ethernet access - business.....	89
Figure 16: Ethernet OAM model for Ethernet access - wholesale.....	89
Figure 17: Local Epipe/VLL service using access or access-uplink ports on the 7210 SAS.....	181
Figure 18: Distributed Epipe/ VLL service using network ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.....	181
Figure 19: Dynamic MS-PW operation.....	183
Figure 20: MS-PW addressing using FEC129 All Type 2.....	183
Figure 21: Advertisement of PE addresses by PW routing.....	184

Figure 22: Signaling of dynamic MS-PWs using T-LDP.....	184
Figure 23: Mapping of All to SAP.....	185
Figure 24: VLL using dynamic MS-PWs, Inter-AS scenario.....	186
Figure 25: Pseudowire redundancy.....	190
Figure 26: Dynamic MS-PW example.....	192
Figure 27: Master-slave pseudowire redundancy.....	195
Figure 28: VLL resilience path.....	196
Figure 29: VLL resilience with PW redundancy and PW switching.....	198
Figure 30: VLL resilience with pseudowire switching.....	199
Figure 31: Access node resilience using MC-LAG and PW redundancy.....	201
Figure 32: VLL resilience with pseudowire redundancy and switching.....	204
Figure 33: Pseudowire switching TLV.....	205
Figure 34: VLL resilience.....	207
Figure 35: Master-slave pseudowire redundancy.....	208
Figure 36: VLL resilience.....	209
Figure 37: VLL resilience with pseudowire switching.....	210
Figure 38: Redundant VLL endpoint objects.....	212
Figure 39: SDPs — unidirectional tunnels.....	222
Figure 40: VLL resilience.....	225
Figure 41: VLL resilience with pseudowire switching.....	226
Figure 42: Default QinQ SAP for transit traffic in a ring scenario.....	228
Figure 43: EVPN for MPLS in VPLS services.....	300
Figure 44: EVPN routes type 1 and 4.....	301

Figure 45: EVPN-VPLS integration.....	308
Figure 46: DF election.....	311
Figure 47: Split-horizon.....	312
Figure 48: Aliasing.....	312
Figure 49: ES discovery and DF election.....	315
Figure 50: All-active multi-homing ES failure.....	321
Figure 51: Black hole caused by SAP/SVC shutdown.....	322
Figure 52: Transient issues caused by "slow" MAC learning.....	323
Figure 53: Backup PE.....	324
Figure 54: Single-active multi-homing ES failure.....	328
Figure 55: Proxy-ARP example usage in an EVPN network.....	330
Figure 56: Network topology of an access ring.....	340
Figure 57: VPLS service architecture.....	404
Figure 58: Access port ingress packet format and lookup.....	404
Figure 59: Network port egress packet format and flooding.....	405
Figure 60: VPLS service architecture.....	406
Figure 61: Access port ingress packet format and lookup.....	407
Figure 62: Network port egress packet format and flooding.....	407
Figure 63: MVR and MVR by proxy.....	413
Figure 64: Access resiliency.....	419
Figure 65: Dual homed 7210 SAS-D acting as MTU-s in two-tier hierarchy H-VPLS.....	422
Figure 66: H-VPLS with spoke redundancy.....	424
Figure 67: H-VPLS resiliency based on AS pseudowires.....	425

Figure 68: Dual homed MTUs in Two-Tier hierarchy H-VPLS.....	426
Figure 69: Dual homed CE connection to VPLS.....	428
Figure 70: BGP-AD and T-LDP Interaction.....	431
Figure 71: Example configuration for protected VPLS SAP.....	462
Figure 72: Example configuration for load balancing across with management VPLS.....	464
Figure 73: BGP AD configuration example.....	467
Figure 74: BGP AD triggering LDP functions.....	469
Figure 75: Show router LDP session output.....	469
Figure 76: Show router LDP bindings FEC-type services.....	470
Figure 77: PW-template-binding CLI syntax.....	471
Figure 78: Sample AS pseudowire.....	478
Figure 79: Internet Enhanced Service.....	666
Figure 80: Internet Enhanced Service.....	727
Figure 81: Virtual Private Routed Network.....	808
Figure 82: Route distinguisher.....	809
Figure 83: Directly connected to IP target.....	810
Figure 84: Multiple hops to IP target.....	811
Figure 85: Inter-AS option-A: VRF-to-VRF model.....	818
Figure 86: Inter-AS option-B.....	819
Figure 87: Option C example.....	819
Figure 88: OSPF areas.....	1028
Figure 89: Split horizon group example.....	1149

1 Getting started

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

1.1 About this guide

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.



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.

- 7210 SAS-D
- 7210 SAS-Dxp 12p (2SFP+ 4SFP 6Tx)
- 7210 SAS-Dxp 16p (2SFP+ 4SFP 10Tx)
- 7210 SAS-Dxp 24p (2SFP+ 6SFP 16Tx)
- 7210 SAS-K 2F1C2T
- 7210 SAS-K 2F6C4T
- 7210 SAS-K 3SFP+ 8C

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.

- **access-uplink mode of operation**

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

- **network mode of operation**

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

When configured in the access-uplink or network mode of operation, the 7210 SAS platform implicitly operates in the standalone mode.

1.1.1 Document structure and content

This guide uses the following structure to describe features and configuration content.



Note:

This guide generically covers Release 23.x.Rx content and may include some content that will be released in later maintenance loads. See the *7210 SAS Software Release Notes 23.x.Rx*, part number 3HE 19296 000x TQZZA, for information about features supported in each load of the Release 23.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.
- 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.
- 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 23.x.Rx*, part number 3HE 19296 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.

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 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. 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 service considerations
	EVPN	Ethernet Virtual Private Networks
	VPLS	Virtual Private LAN Service
	IES (7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T)	Internet Enhanced Service on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T
	IES (7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)	Internet Enhanced Service on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
	VPRN service	Virtual Private Network service on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
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-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C subscriber services, service model, and service entities. Additional information about the individual subscriber services supported on different 7210 SAS platforms and their configuration options is 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.

On the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T platforms, services can provide Layer 2/bridged service between a service access point (SAP) and another service access point (a SAP is where traffic enters and exits the service) on the same (local) router. It cannot support distributed services using MPLS uplinks.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C platforms, services can provide Layer 2/bridged service or Layer 3/IP routed connectivity between a service access point (SAP) on one router and another SAP (which is where traffic enters and exits the service) on the same (local) router or another router (distributed). The use of either MPLS uplinks or Ethernet uplinks is supported.

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C support both local and distributed service. A distributed service spans more than one router. Distributed services use service distribution points (SDPs) to direct traffic through a service tunnel to another Nokia router. 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).

2.1.1 Service types on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

The 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T offer the following types of subscriber services, described in more detail in the referenced chapters:

- **Virtual Leased Line (VLL) services**

- **Ethernet pipe (Epipe)**

- A Layer 2 point-to-point VLL service for Ethernet frames. See [Epipe](#) for more information about Epipe.

- **Virtual Private LAN Service (VPLS)**

- A Layer 2 multipoint-to-multipoint VPN bridging service or VPN (using QinQ uplinks). See [Virtual Private LAN Service](#) for more information about VPLS.

2.1.2 Service types on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C offer the following types of subscriber services, described in more detail in the referenced chapters:

- **Virtual Leased Line (VLL) services**

- **Ethernet pipe (Epipe)**

- A Layer 2 point-to-point VLL service for Ethernet frames. See [Epipe](#) for more information.

- **Virtual Private LAN Service (VPLS)**

- VPLS is a Layer 2 multipoint-to-multipoint VPN bridging service or VPN (using QinQ uplinks). See [Virtual Private LAN Service](#) for more information.

- **Routed VPLS Service (R-VPLS)**

- An R-VPLS integrates routing and bridging into a single service. It provides the capability to provide Layer 2 multipoint-to-multipoint services locally on the node, typically using an IP interface for uplink connectivity. See [Virtual Private LAN Service](#) for more information.

- **Internet Enhanced Services (IES)**

- IES is a direct Internet access service where the customer is assigned an IP interface for Internet connectivity. See the [Internet Enhanced Service on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T](#) for more information about IES.

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

- VPRN is a Layer 3 IP multipoint-to-multipoint VPN service as defined in RFC 2547bis. See [Virtual Private Network service on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#) for more information.

2.1.3 Service policies on 7210 SAS-D and 7210 SAS-Dxp

Common to connectivity services on 7210 SAS-D and 7210 SAS-Dxp platforms are policies 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 following types of policies are common to all 7210 SAS-series connectivity services:

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

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

- Filter policies allow selective blocking of traffic matching criteria from ingressing or egressing a SAP.

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

- Scheduler policies define the operating parameters (such as scheduling algorithm, weights per priority). Depending on the platform, these are either associated with SAPs or physical ports.
- Accounting policies define how to count the traffic usage for a service for billing purposes.

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

2.2 Service policies on 77210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

Common to connectivity services on the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, are policies 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 service enhancements. The following types of policies and their functions are common to all 7210 SAS connectivity services:

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

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

- Filter policies allow selective blocking of traffic matching criteria from ingressing and egressing a SAP.

Filter policies, also referred to as access control lists (ACLs), control the traffic allowed in or out of a SAP, based on MAC or IP match criteria. Associating a filter policy 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.

- Accounting policies define how to count the traffic usage for a service, for billing purposes.

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

2.3 Nokia service model on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

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 and IP/MPLS provider core network in encapsulation tunnels created using generic router encapsulation MPLS label-switched paths (LSPs). The 7210 SAS-D, 7210 SAS-Dxp, and 77210 SAS-K 2F1C2T support only QinQ and dot1q Layer 2 uplinks, which are used to transport the services to the provider edge in a hierarchal configuration. The platforms do not support transport tunnels that use MPLS LSPs.

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 the following:

- 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.

2.4 Nokia service model on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 and IP/MPLS provider core network in encapsulation tunnels created using generic router encapsulation MPLS label-switched paths (LSPs). The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C support transport tunnels that use MPLS LSPs or QinQ/dot1q Layer 2 uplinks. These tunnels are used to transport the services to the provider edge in a hierarchical configuration.

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

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

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

2.5 Service entities on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

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

- [Customers](#)
- [SAPs on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T](#)

2.6 Service entities on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

- [Customers](#)
- [SAPs on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T](#)
- [SAPs on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [SAP encapsulation types and identifiers](#)
- [SDPs on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [G.8032/Ethernet ring protection switching](#)
- [Layer 2 control processing](#)
- [Service creation process overview](#)

2.6.1 Customers

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

2.7 SAPs on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

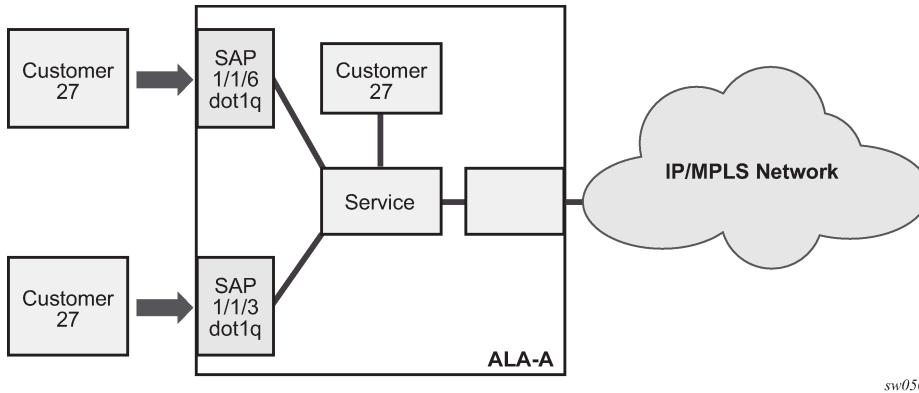
Each subscriber service type is configured with at least one SAP. [Figure 1: Multiple SAPs in a service using QinQ uplinks in 7210 SAS configured in access-uplink mode](#) shows how a SAP identifies the customer interface point for a service on a 7210 SAS router. The SAP configuration requires that slot, 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-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C 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" or "access uplink" in the physical port configuration. SAPs can be created on ports designated as core facing "access uplink" ports. These ports have a different set of features enabled in software.

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



The preceding figure shows 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).

2.8 SAPs on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

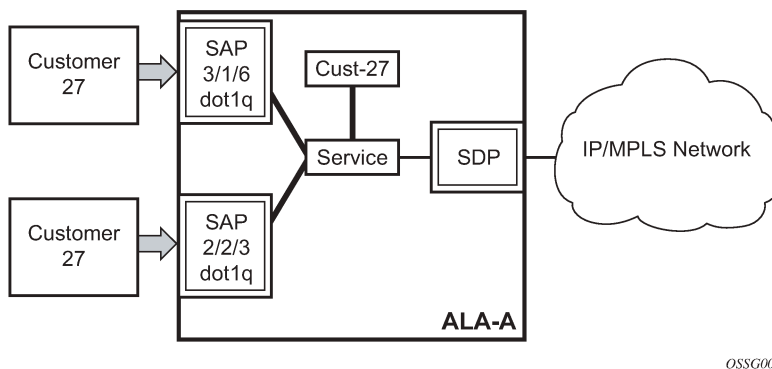
Each subscriber service type is configured with at least one SAP. A SAP identifies the customer interface point for a service on a 7210 SAS router (Figure 2: Multiple SAPs for 7210 SAS configured with MPLS uplinks). The SAP configuration requires that slot, MDA, and port information be specified. The slot, MDA, and port parameters must be configured before provisioning a service (see the Cards, MDAs, and Ports sections of the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C 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.

Figure 2: Multiple SAPs for 7210 SAS configured with MPLS uplinks

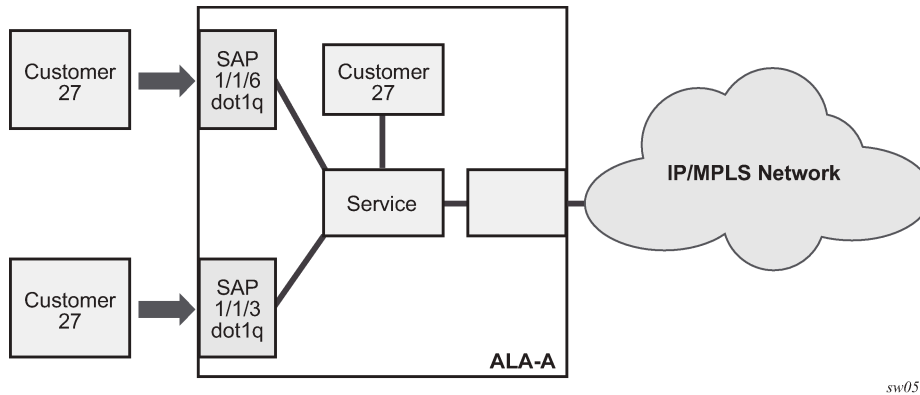


The preceding figure shows SAPs used for customer service delivery, with SDP used for service transport on 7210 SAS devices that support MPLS uplinks.

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

The following figure shows SAPs used for customer service delivery with access-uplink SAPs (also known as QinQ SAPs) used for service transport on 7210 SAS devices when using Layer 2 uplinks.

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



2.9 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.

2.9.1 Ethernet encapsulations supported on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

The following lists encapsulation service options on Ethernet access ports:

- **null**

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

- **dot1q**

Supports multiple services for one customer or services for multiple customers ([Figure 4: Multiple SAPs on a single port](#)). For example, the port is connected to a customer who needs multiple services. The encapsulation ID used to distinguish an individual service is the VLAN ID in the IEEE 802.1Q header.

- **QinQ**


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

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

- **QinQ**

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

Figure 4: Multiple SAPs on a single port



OSSG002

2.9.2 Services and SAP encapsulations

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

Table 5: Service and SAP encapsulation for access ports

Port type	Encapsulation	7210 SAS platforms support
Ethernet	null	All
Ethernet	Dot1q	All
Ethernet	QinQ	7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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

Table 6: Port type and encapsulation for access-uplink ports

Port type	Encapsulation	7210 platforms
Ethernet access-uplink	QinQ	All

2.9.3 Default SAPs 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.

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 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.
- 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.

2.9.4 Default SAPs on a QinQ port

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. 7210 SAS-K, accepts untagged and tagged packets on a default QinQ SAP.

When an Epipe service with default QinQ SAPs on the ring ports is used for transit traffic in a ring deployment, no protection mechanism (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.

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).

For default QinQ SAPs on either access ports or access-uplink ports the following is true:

- On the 7210 SAS-D and 7210 SAS-Dxp, a default QinQ SAP accepts only tagged packets. Untagged packets or priority-tagged packets are not accepted on default QinQ SAPs.
- On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, a default QinQ SAP accepts both untagged and tagged packets.
- On the 7210 SAS-D and 7210 SAS-Dxp, default QinQ SAP is available for use only in an Epipe and a VPLS service created with **svc-sap-type** parameter set to "null-star". Default QinQ SAP can be configured along with other SAPs allowed in the same service (that is, service with **svc-sap-type** parameter set to "null-star").
- On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, default QinQ SAP is available for use only in an Epipe and a VPLS service created with **svc-sap-type** parameter set to "any". Default QinQ SAP can be configured along with other SAPs allowed in the same service (that is, service with **svc-sap-type** parameter set to "any").
- MAC learning and aging is available for use in a VPLS service.
- Per-SAP MAC limit is available for use in a VPLS service.
- Mac-move detection and Mac-pinning is available for use in a VPLS service.
- Discard-unknown and discard-unknown-source is available for use in a VPLS service.
- Ethernet Connectivity Fault Management (ETH-CFM) and Y.1731 are not available for use.
- STP (and all its different flavors) cannot be enabled in the service with default QinQ SAPs.
- M-VPLS with xSTP can be used for loop prevention. The default QinQ SAPs inherit the state from the associated M-VPLS 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 needs to be configured for use on the ring ports, and the service with default QinQ ports needs to be associated with this G.8032 control instance.
- IGMP snooping is not available for use.
- L2PT and BPDU translation is not available for use.
- On the 7210 SAS-D and 7210 SAS-Dxp, IP interface in a VPLS service is not supported in a service priority-tagged using this SAP. On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, IP interface in a VPLS service is not supported (irrespective of SAP encapsulations).

For default QinQ SAPs created on an access-uplink port, the following is true:

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

For default QinQ SAPs created on access ports, the following is true:

- A SAP ingress QoS policy is available for use.
- On the 7210 SAS-D and 7210 SAS-Dxp, an egress QoS policy applied on an access port is available for egress shaping, scheduling, and marking.
- On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, an egress QoS policy applied on an access SAP is available for egress shaping, scheduling and marking.
- SAP ingress ACLs are available for use.
- SAP egress ACLs are not available for use.
- On the 7210 SAS-D and 7210 SAS-Dxp, SAP ingress meter counters, SAP ingress received count, and SAP egress forwarded counter are available for use (appropriate accounting records can be used).
- On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, SAP ingress queue counters, SAP ingress received count, and SAP egress forwarded counter are available for use (appropriate accounting records can be used).

2.9.4.1 Configuration notes for default QinQ SAPs for transit service in a ring deployment

- 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/loop-detection mechanisms can be implemented for VPLS service configured in the ring nodes, by using M-VPLS 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 M-VPLS/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 M-VPLS with xSTP configured for protection, load-balancing of the traffic based on the VLAN IDs is not possible. If load-balancing is needed, then it is better to use Epipe service with default QinQ SAPs as the transit service.

2.9.5 SAP configuration considerations applicable for SAPs configured on access-uplink ports



Note:

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the simultaneous configuration of access-uplink and network operating modes, without explicit BOF configuration, is supported. A mix of access-uplink and network ports can be simultaneously configured on these platforms.

When configuring a SAP, consider the following:

- 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 device.

- On the 7210 SAS-D and 7210 SAS-Dxp, 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.
- On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, multiple SAPs configured on the same port can be part of the same service.
- The 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C support the use of a Q1.0 SAP. This SAP matches packets received on a port with the outermost tag being Q1 and the inner tag being absent (that is, no tag) or the inner tag is a priority tag. It does not accept packets with any other VLAN tag value as the inner tag.
- There are no default SAPs configured on the node. All SAPs in subscriber services must be created.
- The default administrative state for a SAP at creation time is administratively enabled.
- When a SAP is deleted, all configuration parameters for the SAP will also be deleted.
- An access SAP is owned by and associated with the service in which it is created in each router.
- A port with a dot1q encapsulation type means the traffic for the SAP is identified based on a specific IEEE 802.1Q VLAN ID value. The VLAN ID is stripped off at SAP ingress and the appropriate VLAN ID placed on at SAP egress. As a result, VLAN IDs only have local significance, so the VLAN IDs for the SAPs for a service need not be the same at each SAP.



Note:

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

- If a port is administratively shut down, all SAPs on that port are operationally out of service.
- A SAP cannot be deleted until it has been administratively disabled (shut down).
- Each SAP can have one each of the following policies assigned:
 - ingress filter policy
 - egress filter policy
 - ingress QoS policy
 - accounting policy
 - egress QoS policy on the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C
- An ingress QoS policy and accounting policy is assigned per access-uplink port and cannot be assigned per access-uplink SAP. That is, all access-uplink SAPs on a specific port share the QoS policy.
- Ingress filter policy and egress filter policy is available per access-uplink SAP.
- The **svc-sap-type** parameter value determines the type of SAPs provisioned in a service:
 - provides details of SAP and service combinations allowed in 7210 SAS-D and 7210 SAS-Dxp devices
 - provides details of SAP and service combinations allowed in 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C devices
- 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.

- For dot1q range SAP, the outermost VLAN tag is compared against the range of VLAN IDs configured and mapped to the SAP if there is a match. The outermost tag is not stripped on SAP ingress and is forwarded through the other service endpoint.
- L2PT cannot be configured for use on all the configured SAPs simultaneously. The number of SAPs that can use this simultaneously is less than the maximum number of SAPs supported by the node.

2.9.5.1 Configuration guidelines for 7210 SAS-D and 7210 SAS-Dxp

The following considerations apply to SAP configurations:

- Ensure that egress SAP counters are enabled before associating accounting records that count egress forwarded packets.
- Before modifying the counter, disable the account log generation and run the **no collect-stats** command.
- Egress SAP statistics are not available on any of the SAPs of a port, on which a dot1q SAP and dot1q default SAP configuration are present at the same time. This limitation also applies for egress ACLs. That is, egress ACLs are not supported on either dot1q SAPs or dot1q default SAPs when both these SAPs are configured on a port simultaneously.
- Egress SAP statistics cannot be configured for use simultaneously on all the configured SAPs. The number of SAPs that can use this feature simultaneously is less than the maximum number of SAPs supported by the node.
- Service MTU is not supported on the 7210 SAS-D.
- QinQ access SAPs of type Q1.0 are supported only for IES and R-VPLS services. They are not supported for Layer 2 services.

The following table lists the SAPs allowed with different values for the **svc-sap-type** on 7210 SAS-D and 7210 SAS-Dxp.

Table 7: SAP and service combinations for 7210 SAS-D and 7210 SAS-Dxp

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

svc-sap-type	Access SAPs	Access-uplink SAPs
	Q1.Q2 SAP, Q.* SAP, 0.* SAP	
dot1q-range	Dot1q SAP (Dot1q VLAN tag not stripped on ingress), Q1.* SAP	Q1.* SAP

- A dot1q default SAP cannot be configured when the **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 the **svc-sap-type** set to **null-star**, to process and forward packets with one or more tags (including priority tag) on a null SAP.
- The default QinQ SAP processes only tagged packets received on a QinQ port. All tagged packets that do not match the specific SAP tags configured on the same port are processed by this SAP. The default QinQ SAP cannot process untagged packets, even if 0.* SAP is not configured for use on that port.
- The default QinQ SAPs are 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 the **svc-sap-type** parameter set to **null-star**.
- SAPs using connection profiles (to specify dot1q VLAN ranges) can be configured in a service only when the **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](#) for information about processing behavior for this type of service.
- Service MTU is not supported on 7210 SAS-D.
- The following counters are available on 7210 SAS-D and 7210 SAS-Dxp devices:
 - ingress and egress counters per SAP
 - ingress policer counters per SAP
 - egress queue counters per access port
 - ingress and egress counters per access-uplink port
- The number of counters available to count total received packets or octets on an access-uplink SAP (in a VPLS, VLL or IES service) ingress is limited; therefore a count of received packets or octets cannot be obtained for all the SAPs simultaneously. By default, these counts are not available.
- The **config>service>epipe>sap>statistics>ingress** command is available to associate a counter with the SAP and obtain the counts.
- The number of counters available to count total forwarded packets or octets on an access SAP egress and access-uplink SAP egress is limited and therefore a count of received packets or octets cannot be obtained for all the SAPs simultaneously. By default, these counts are not available.

- The `config>service>epipe>sap>statistics>ingress>forwarded-count` command is available to associate a counter with the SAP and obtain the counts.

2.9.5.2 Configuration guidelines for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

The following guidelines apply to SAP configurations:

- The 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C support the following types of SAPs:
 - **access SAPs**
null, dot1q (dot1q, dot1q default, dot1q range, dot1q explicit null), QinQ SAPs (Q1.Q2, Q1.* , Q1.0 QinQ default SAP (that is, *.* SAP), 0.*) are supported.
 - **access-uplink SAPs**
QinQ SAPs (various SAP types, such as Q1.Q2, QinQ default SAP (*. * SAP), 0.* , Q1.* , Q1.0) are supported.
- The 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C support Q1.0 SAP. This SAP accepts a packet with the outermost tag as Q1 or a packet with outermost tag as Q1 and the following inner tag is a priority tag. Unlike 7x50, it does not accept packets with two tags with the outermost tag being Q1 and the inner tag being any tag other than priority tag.
- The 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C allow any port to be configured in either access-uplink mode or access mode. Additionally the ports can be in access-uplink mode or access mode or they can be mix of the ports using either modes. There is no limit to the number of access ports allowed to be configured. That is, all ports can be configured as access ports. The number of access-uplink ports depends on the number of QoS resources allocated per port. That is, not all the ports can be configured as access-uplink ports at a specific time.
- The 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C support service MTU:
 - A received frame/packet length is checked against the configured service MTU after subtracting the length of the SAP encapsulation (including the Layer 2 header) from the received frame length. The packet is further processed in the context of the service, if the computed length is less than equal to the configured service MTU or the packet is dropped.
 - The user must configure the correct service MTU across all the nodes, if support is available, through which the service is transported.
- The 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C support only the following svc-sap type parameters:
 - **'any'**
A service configured with this value for svc-sap-type allows for configuration of all combination of access SAPs and access-uplink SAPs in the same service, except for dot1q range SAPs. A packet that is received with tags more than the number of SAP tags to which it is mapped to, is forwarded transparently in the service (the processing behavior is similar to any other packet mapped to the SAP).
 - **'dot1q-range'**
A service configured with this value for svc-sap-type allows for configuration of dot1q range SAPs and Q1.* access-uplink SAP in the same service.

The following table lists the SAPs allowed on the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C with different values of svc-sap-type.

Table 8: SAP and service combinations for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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

2.10 VLAN range for SAPs (or dot1q range SAPs) in an Epipe or VPLS service

The 7210 SAS VLAN ranges (or dot1q range SAPs) provide a mechanism to group a range of VLAN IDs as a single service entity. This allows the operator to provide the service treatment (forwarding, ACL, QoS, accounting, and others) to the group of VLAN IDs as a whole.



Note:

- Dot1q range SAPs are supported on all 7210 SAS platforms as described in this document.
- Grouping a range of VLAN IDs to a SAP is supported only for Epipe and VPLS Ethernet services.

2.10.1 Epipe processing behavior for dot1q range SAPs in access-uplink mode on the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

The access SAPs that specify VLAN range values using the connection profile (also known as dot1q range SAPs) are allowed in the Epipe and VPLS service. For more information about the supported functionality, see [Epipe VLAN range SAPs feature support and restrictions](#). The system allows only one range SAP in an Epipe service; attempts to configure more than one range SAP fail. Range SAPs can only be configured on access ports. The other endpoint in the Epipe service must be a Q.* SAP in access-uplink mode. The processing and forwarding behavior for packets received on range SAPs is as follows:

- No VLAN tags are 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**.

2.10.1.1 Epipe VLAN range SAPs feature support and restrictions

The feature support and restrictions of VLAN range SAPs in an Epipe service are the following:

- 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.
- On 7210 SAS-D and 7210 SAS-Dxp, the dot1q range SAP can only be configured in a service with **svc-sap-type** set to **dot1q-range**.
- On 7210 SAS-K 2F1C2T, the dot1q range SAP can only be configured in a service with **svc-sap-type** set to **dot1q-range**.
- The access SAPs using VLAN range values are only allowed for ports or a LAGs configured for dot1q encapsulation.
- A connection profile is used to specify either range of VLAN IDs or individual VLANs to be grouped together in a single SAP.
- On 7210 SAS-D and 7210 SAS-Dxp, dot1q default SAP and dot1q range SAP is not allowed to be configured on the same access port. That is, they are mutually exclusive. This restriction does not apply to 7210 SAS-K 2F1C2T.
- Multiple connection-profiles can be used per port or LAG if the VLAN value specified by each of them does not overlap. The number of VLAN ranges available per port/LAG is limited. The available number must be shared among all the SAPs on the port/LAG.
- The connection profile associated with a SAP cannot be modified. To modify a connection profile, it must be removed from all SAPs using it.
- ACL support – Filter policies are supported on SAP ingress. IP criteria and MAC criteria-based filter policy is available for use with access SAPs. For more information about ACL on range SAPs, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.
- On 7210 SAS-K 2F1C2T, the outermost VLAN tag and Inner VLAN tag matches are supported for both ingress/egress ACL and ingress QoS classification.
- On 7210 SAS-D and 7210 SAS-Dxp, access SAP egress filters are not supported for SAPs configured in a dot1q-range service. This restriction does not apply to 7210 SAS-K 2F1C2T.
- On 7210 SAS-D and 7210 SAS-Dxp, access-uplink SAP egress filters are not supported for SAPs configured in a dot1q-range service. This restriction does not apply to 7210 SAS-K 2F1C2T.
- QoS on 7210 SAS-D and 7210 SAS-Dxp, Ingress classification, metering with hierarchical metering support for SAP ingress is available:
 - 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.

- QoS on 7210 SAS-K 2F1C2T- Ingress classification, queuing with hierarchical shaping support for SAP ingress is available:
 - 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.
- 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 can match a range of VLAN values, and therefore uses X resources for a SAP using a VLAN range, instead of $X * n$, where n is the number of VLANs specified in the range and X is the amount of QoS or ACL resources needed.
- Ingress accounting support is like the support available for other SAPs in an Epipe service. 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.

2.10.1.2 Epipe processing behavior for dot1q range SAPs on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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; attempts to configure more than one range SAP fail. 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 Q.* access-uplink SAP or a spoke-SDP (PW) in network mode.

The spoke-SDP processing and forwarding behavior for packets received on range SAPs are as follows:

- No VLAN tags are stripped on ingress of the access dot1q SAPs using VLAN range connection profile.
- When the other endpoint in the service is configured to be an Q1.* access SAP or a Q1.* access-uplink SAP, the 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.* access SAP or a Q1.* access-uplink 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 range SAP. 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. That is, when forwarding the traffic out of the range SAPs, the outermost VLAN tag value is not matched against the range configuration.

2.10.1.2.1 Epipe VLAN range SAPs feature support and restrictions on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The following information describes feature support and restrictions for VLAN range SAPs configured in an Epipe service:

- The access SAPs that specifies VLAN range values (using connection profile) is only allowed in Epipe service. The system allows only one range SAP in an Epipe service. It will fail any attempt to configure more than one range SAP in an Epipe service. Range SAP can be configured only on access ports.
- The dot1q range sap can be configured only in a service with **svc-sap-type** set to **dot1q-range**.
- The access SAPs using VLAN range values are only allowed for dot1q encapsulation port.
- A connection profile is used to specify either range of VLAN IDs or individual VLANs to be grouped together in a single SAP.
- The dot1q default SAP and dot1q range SAP can co-exist on the same access port.
- 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 match resources available per port/LAG is limited. The available number must be shared among all the dot1q range SAPs configured on the port/LAG.
- The connection profile associated with a SAP cannot be modified. To modify a connection profile, it must be removed from all SAPs using it.
- ACL support – Filter policies are supported on SAP ingress and SAP egress. Only MAC criteria-based filter policy is available for use with access SAPs. For more information about ACL on range SAPs, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.
- The outermost VLAN tag and Inner VLAN tag matches are supported for both ingress/egress ACL and ingress QoS classification.
- SDP egress and ingress filter are not supported.
- QoS – ingress classification, queuing with hierarchical shaping support for SAP ingress:
 - SAP ingress classification criteria that is available for use with VLAN range SAPs is like that available for other SAPs supported in an Epipe service. Dot1p-based ingress classification uses the dot1p bits in the outermost VLAN tag for matching. On access egress, dot1p received from the SDP (on a network port) from another access port is preserved.
- The amount of hardware resources (such as CAM entries used for matching in QoS classification and ACL 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 can match a range of VLAN values, and therefore uses X resources for a SAP using a VLAN range, instead of $X * n$, where n is the number of VLANs specified in the range and X is the amount of QoS or ACL resources needed.
- Ingress accounting support is like the support available for other SAPs in an Epipe service. 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.
- Service resiliency mechanisms, such as Epipe PW redundancy, are supported.

2.10.2 VLAN range SAPs in a VPLS service on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

VLAN ranges on the 7210 SAS 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 so on) to the group of VLAN IDs as a whole.

2.10.2.1 VPLS processing behavior for dot1q range SAPs on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The access SAPs that specify VLAN range values (using the connection profile) are allowed in VPLS services. The system allows support for one or more dot1q range access SAPs in a VPLS service, with the other endpoint being only a spoke-SDP, mesh-SDP, or access-uplink SAP. No other SAP encapsulations are allowed when the access SAP is a dot1q range SAP (other than access uplink SAP as specified above).

The processing and forwarding behavior for packets received on range SAPs are as follows:

- No VLAN tags are stripped on the ingress of the access dot1q SAPs using the VLAN range connection profile.
- When the other endpoint in the service is configured to be a Q1.* access SAP or a Q1.* access-uplink SAP, the 7210 SAS 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 and mesh-SDP, 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.* access SAP or a Q1.* access-uplink 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 range SAP. When the other endpoint is a spoke-SDP and mesh-SDP, the 7210 SAS removes the MPLS PW and LSP encapsulation and forwards the packet out of the access dot1q range SAP.
- The system does not check if the outermost VLAN tag in the packet matches the VLAN IDs of the dot1q range SAP configured in the service. That is, when forwarding the traffic out of the dot1q range SAPs, the outermost VLAN tag value is not matched against the range configuration.

2.10.2.2 VPLS VLAN range SAP feature support and restrictions

The following information describes dot1q range SAPs feature support and restrictions:

- Dot1q range SAPs are only allowed to be configured in a service with the **svc-sap-type** set to **dot1q-range**.
- The access SAPs using VLAN range values are only allowed for a dot1q encapsulation port or LAG.
- A connection profile is used to specify either the range of VLAN IDs or individual VLANs to be grouped together in a single SAP.
- The system allows support for one or more dot1q range access SAPs in an VPLS service, with the other endpoint being only a spoke-SDP, mesh-SDP, or access-uplink SAP. No other SAP encapsulations are allowed when the access SAP is a dot1q range SAP (other than access uplink SAP as specified above).
- Dot1q default SAPs are not allowed on the same access port as the one on which a SAP with a range is configured.
- 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 or LAG is limited. The available number must be shared among all the SAPs on the port or 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.

- L2 processing (Learning, Unicast forwarding, BUM flooding, and replication) completed for traffic received on a dot1q range SAP is similar to the processing for a NULL/dot1q default SAP in a VPLS service. In other words, no special handling is required for dot1q range SAPs, other than the requirement to not add and or remove the outermost VLAN tag for a dot1q range SAP.
- For BUM traffic flooding (including L2 multicast if any), a packet is replicated to all dot1q range SAPs in the VPLS service with no modification or addition of VLAN tags. In other words, a single copy of BUM traffic is sent out of the SAP and the dataplane does not replicate a copy per VLAN ID configured in the dot1q range SAP for BUM traffic. It is assumed that the downstream device can do the required replication for L2 multicast/BUM traffic if needed.
- No egress checks are implemented for all traffic (unicast and BUM) to check if outermost VLAN tag in the packet matches the VLAN IDs of the range of the egress SAP that the packet is sent out of.
- L2 forwarding accounts for service-based SHG or port-based SHG while forwarding.
- Learning is enabled on the service by default. The user has an option to disable learning per-service. Learning enable and disable per-SAP is not supported.
- MAC limiting is available per-service. MAC limiting per-SAP is not supported.
- ACL support:
 - Filter policies are supported on SAP ingress. IP criteria and MAC criteria-based filter policies are available for use with access SAPs.
 - Egress filters are supported for access SAPs and access-uplink SAPs configured in a dot1q range service.
 - The outermost VLAN tag and inner VLAN tag matches are supported for both ingress/egress ACL and ingress QoS classification.

For more information about ACL on range SAPs, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.
- QoS on the 7210 SAS-K 2F6C4T: ingress classification and queuing with hierarchical shaping support for SAP ingress is available:
 - SAP ingress classification criteria that is available for use with dot1q range SAPs is similar to that available for other SAPs supported in a VPLS service. Dot1p-based ingress classification uses the dot1p bits in the outermost VLAN tag for matching.
 - SAP egress queueing and shaping is supported for dot1q range SAPs.
- 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 so on) 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 can match a range of VLAN values, and therefore uses X resources for a SAP using a VLAN range, instead of $X * n$, where n is the number of VLANs specified in the range, and X is the amount of QoS or ACL resources needed.
- Ingress accounting support is like the support available for other SAPs in a VPLS service. The packets or octets count received from individual VLANs configured in the connection profile is not available.
- Egress SAP statistics and accounting are not supported.
- Mirroring is supported.
- IGMP snooping and MVR are not supported.
- DHCP snooping is not supported.
- xSTP is not supported.

- G8032 is not supported.
- CFM OAM is not supported.
- A/S PW redundancy is not supported.

2.10.3 Epipe emulation using dot1q VLAN range SAP in VPLS with G.8032




Note:

This feature is only supported on the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C 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 for breaking the loop in the ring and VPLS service protection. The intermediate nodes on the ring need 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 5: Epipe emulation in a ring using VPLS with G.8032



The following are the requirements to provide for an Epipe service connectivity between two business sites:

- Transport all the VLANs used by the internal enterprise network of the businesses.
- Support high availability for the service between the business sites by protecting against failure of the links or nodes in the ring.

To achieve connectivity between two business sites in access-uplink/L2 mode is to configure SAPs for each of the individual VLANs used in the enterprise network in a VPLS service and use G.8032 for protection. The number of VLANs that was supported is limited by the number of SAPs supported on the platform.

The 7210 SAS platforms, currently support the use of Dot1q range SAPs with only Epipe services in either network/MPLS mode or access-uplink/L2 mode. Dot1q range SAPs allows operators to transport a range

of VLANs by providing similar service treatment (service treatment refers to forwarding decision along with encapsulation used, QoS and ACL processing, accounting, 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 L2 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 L2 mode, by emulating an Epipe service using a VPLS service with G.8032 protection and at the same time provides the benefits of using dot1q range SAPs. The user should ensure that the VPLS service is a point-to-point service. This is achieved by configuring a VPLS service with an access dot1q range SAP used at the customer hand-off on one node in the ring and an access dot1q range SAP in a customer hand-off of a VPLS service on another node (that is, at the other end of the Epipe), such that there are only two endpoints for the service in the network.

On the node where the service originates, in addition to the access dot1q range SAP, the service needs to be configured with access-uplink SAPs on the two G.8032 ring ports. G.8032 mechanism is used to for breaking the loop in the ring and VPLS service protection. The intermediate nodes on the ring needs to use VPLS service with access-uplink SAPs on the ring ports and use the same G.8032 instance for protection, as one is used for service protection on the originating node.

2.10.3.1 Epipe emulation using dot1q VLAN range SAP in VPLS with G.8032 — configuration guidelines and restrictions

The VPLS service with dot1-range SAPs use svc-sap-type of dot1q-range and supports limited functionality in comparison to a normal VPLS service, The following paragraph provide more details of the feature functionality, configuration guidelines and restrictions:

- The user can define access dot1q range SAPs, which specifies a group of VLANs which receive similar service treatment, that is, forwarding behavior, SAP ingress QoS treatment and SAP (behavior like that available in Epipe service) and allows it to be configured in a VPLS service:
 - On the node, where the service originates, in addition to the access dot1q range SAP, the service should be configured with Q1.* SAPs on the two G.8032 ring ports. The access or access-uplink Q1.*SAPs can be used, but the access-uplink SAPs are recommended for use. The user cannot configure any other SAPs in the same VPLS service.
 - There is no special configuration required on intermediate nodes, that is, the ring nodes which do not terminate or originate the service. The nodes should be configured for providing transit VPLS service and the VPLS service must use the same G.8032 instance for protection as is used by the service on originating and terminating node.
 - The Epipe service on 7210, currently does not check if the inner tag received on a Q1.* SAP is within the range of the configured VLANs. VPLS service too has the same behavior.
- Support for SAP Ingress QoS, Ingress and Egress ACLs, accounting, and other services, for dot1q range SAP configured in a VPLS service matches the support available in Epipe service.
- G.8032 mechanism is used for loop detection in ring network and service protection. A separate VPLS service representing the G.8032 control instance must be configured and the state should be associated with this service:

- Use of dot1q range SAPs to provide service on the interconnection node, in a G.8032 major-ring/sub-ring deployment, when using the virtual channel, is not supported. This restriction is not applicable when the interconnection node in a G.8032 major-ring/sub-ring is configured without a virtual channel.
- mVPLS/xSTP support is available for use with Q1.* SAP on the ring ports to break the loop. This is an add-on to the G.8032 support.
- Broadcast, Unknown Unicast and Multicast (BUM) traffic is flooded in the service.
- Learning is enabled on the service by default, to avoid the need to flood the service traffic out of one of the ring ports, after network MAC addresses are learnt. The user has an option to disable learning per service. Learning enable/disable per SAP is not supported.
- MAC limiting is available per service. MAC limiting per SAP is not supported.
- CFM OAM is supported. The support for UP MEPs on the dot1q range SAP in the service to be used for fault management and performance management using the CFM/Y.1731 OAM tools is available:
 - On 7210 SAS-D and 7210 SAS-Dxp, 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 dot1q range SAP using the **config>eth-cfm>domain>association>bridge-identifier** VLAN command and enable the use of using the CLI command **primary-vlan-enable** under the MEP CLI context. It is used as the VLAN tag in the packet header for all the CFM/Y.1731 messages sent out in the context of the UP MEP. This is not supported on 7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T.
 - 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.

2.11 SDPs on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C



Note:

SDPs are only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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 router, which directs packets to the correct service egress SAPs on that router. 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 router to a far-end router requires a return path SDP from the far-end router 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.

2.11.1 SDP binding

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

The following figure shows MPLS service distribution pointing from ALA-A to ALA-B.

Figure 6: MPLS SDP pointing from ALA-A to ALA-B



2.11.2 Spoke and mesh SDPs

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

2.11.3 SDP using BGP route tunnel

SDPs are enhanced to use BGP route tunnel 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 due to 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.

2.11.4 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 OAM 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 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.

For information about configuring keepalive parameters, see [Configuring an SDP](#).

2.11.5 SDP administrative groups

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

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

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

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

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

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

```
config>service>sdp>sdp-group group-name
```

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

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

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

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

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

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

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

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

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

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

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

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

2.11.6 Mixed-LSP mode of operation

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

An LDP LSP can be configured as a primary LSP type, which can then be backed up by a BGP LSP type. At any 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 can be used per SDP. The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C support LDP ECMP.

3. BGP LSP type

One RFC 3107-labeled BGP prefix, which is programmed by the service manager, is used.

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



Note:

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

Use the **configure>router>ldp>tunnel-down-damp-time** command. For more information, see the *7210 SAS-K 2F6C4T, K 3SFP+ 8C MPLS Guide*.

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

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



Note:

The following are differences in behavior of the LDP/BGP SDP compared to that of 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. The impact of this is that the tunnel table needs to be reprogrammed each time a route is deactivated and the other is activated.
- The SDP revert-time cannot be used, because there is no situation where both LSP types are active for the same /32 prefix.

2.12 G.8032/Ethernet ring protection switching



Note:

On the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, CCMs used for G.8032 Ethernet ring protection service are implemented in hardware.

Ethernet ring protection switching (Eth-ring) provides ITU-T G.8032 specification compliance to achieve resiliency for Ethernet Layer 2 networks. 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 Rings SAPs can connect to other rings and Ethernet service using 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. Due to the symmetry and the simple topology, rings are viewed a good solution for access and core networks where resilient LANS are required. The Nokia implementation of G.8032 Eth-ring can be used for interconnecting access rings and to provide traffic engineered backbone rings.

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

2.12.1 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, 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, LOS (Loss of Signal) when the nodes are directly connected.

Initially, each Ring Node blocks one of its links and notifies other nodes in the ring about the blocked link. When a ring node in the ring learns that another link is blocked, the node unblocks its blocked link, possibly causing FDB flush in all links of the ring for the affected service VLANs, controlled by the ring control instance. This procedure results in unblocking all links 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 or nodes detecting the failure (enabled by Y.1731 OAM CC monitoring) sends 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: 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 R-APS message, when received by the nodes at the recovered link cause them to unblock that link and restore connectivity (again all nodes in the ring perform 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 as optionally 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 when 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.

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

In the 7210 SAS implementation, the Ethernet ring is built from a VPLS service on each node with VPLS SAPs that provides ring path with SAPs. As a result, most of the VPLS SAP features are available on Ethernet rings, if 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 diagram shows an example G.8032 ring. The following [Figure 9: 0 to 3 ring example](#) shows a resilient ring service. In the ring example, a QinQ ring (solid line) using VID 500 carries two customer VLANs dot1q 100 and QinQ 400.1, respectively). The RPL for the G.8032 ring is between A and B, where B is the RPL owner. [Figure 9: 0 to 3 ring example](#) is also a QinQ service on the (dotted line) ring that uses dot1q VID 600 for the ring to connect service VLAN 100.50.


The two rings have RPLs on different nodes which allow a form of load balancing. The example serves to illustrate that service encapsulations and ring encapsulation can be mixed in various combinations.



Note:

Neither of the rings is a closed loop. A ring can restore connectivity when any one node or link fails to all remaining nodes within a small amount of transfer time (signaling time after detection).

Figure 9: 0 to 3 ring example



Example

The following are sample configuration outputs for G.8032 ring for the preceding figure.

1. Configure G.8032 ring Paths and Control MEPs used for failure detection on ring ports.

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

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

2. Configure VPLS service used for G.8032 control instances (identified with VID tag 100 and 500) used to exchange R-APS messages for the two control instances created on the ring.

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

3. Configure VPLS data services that will use G.8032 for protection.

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

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

  no shutdown
exit

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

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

  no shutdown
exit
```

2.12.2 G.8032 Ethernet ring sub-rings

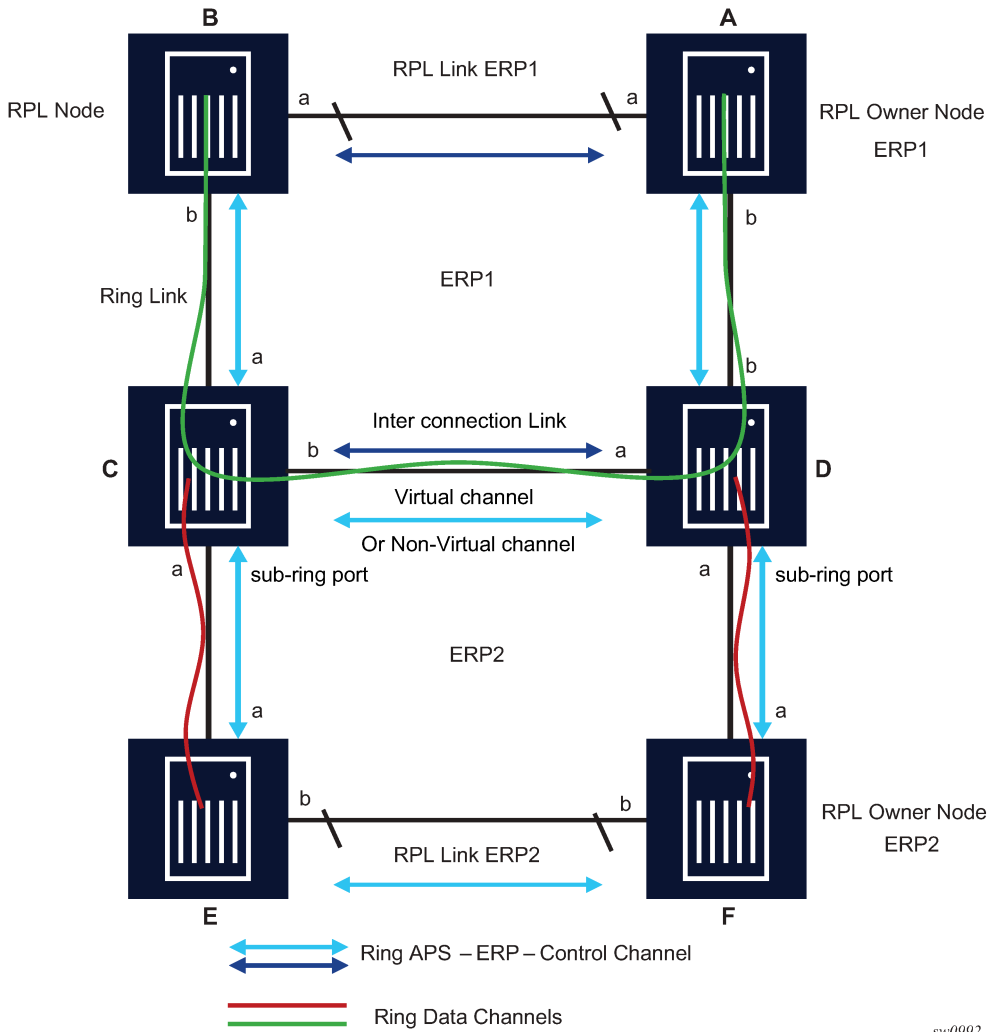
Ethernet sub-rings offer a dual redundant way to interconnect rings. The 7210 SAS supports sub-rings connected to major rings, and a sub-ring connected to a VPLS (LDP based) for access ring support in VPLS networks, as shown in the following figure.



Note:

The platforms as described in this document cannot be used as the interconnection nodes. They can be used only as the ring nodes in the sub-ring.

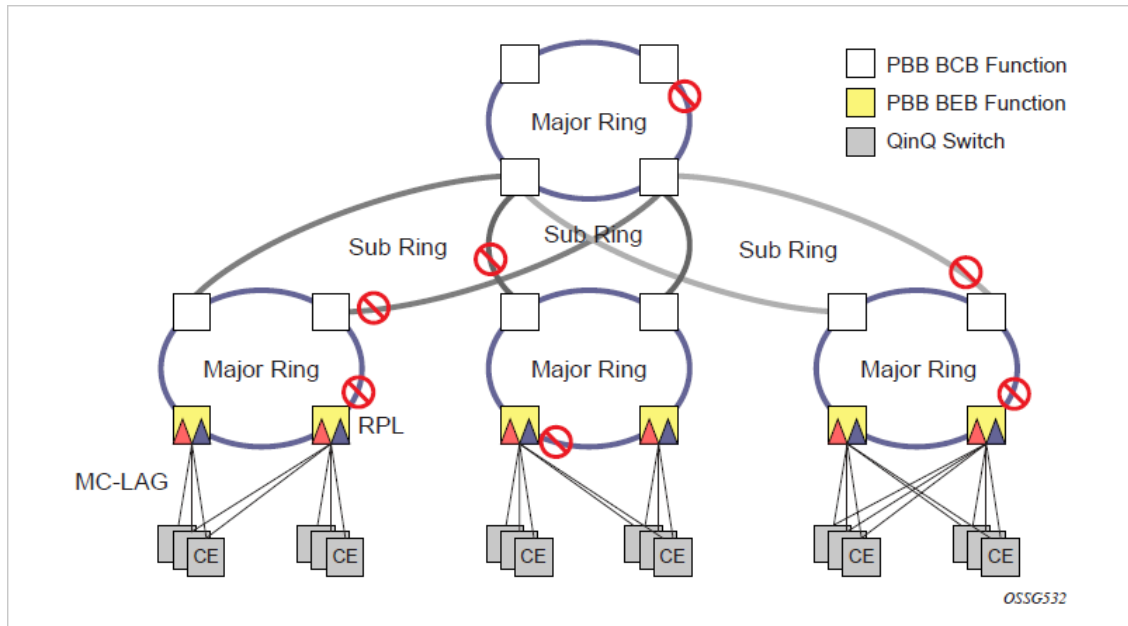
Figure 10: Major ring and sub-ring scenario



sw0992

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 11: 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.

2.12.3 G.8032 virtual and non-virtual channel

Example

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

```
eth-ring 2
  description "Ethernet Sub Ring on Ring 1"
  interconnect ring-id 1 // Link to Major Ring 1
  propagate-topology-change
  exit
exit
path a 1/1/3 raps-tag 100 // Ring control uses VID 100
  eth-cfm

  mep 9 domain 1 association 4
  ccm-enable
  control-mep
  no shutdown
  exit
exit
no shutdown
```

```
        exit
        no shutdown
    exit

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

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

# Data configuration for the Sub-Ring

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

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

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


2.12.3.1 G.8032 Ethernet ring sub-ring using non-virtual link

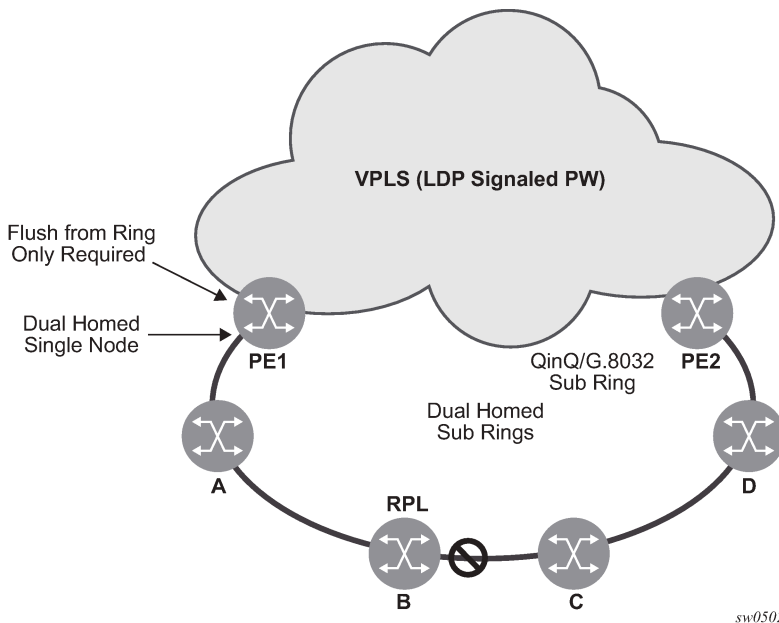
The following figure shows 0 to 6 sub-ring homed to VPLS.



Note:

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

Figure 12: 0 to 6 sub-ring homed to VPLS



Example: Configuration Outputs for sub-rings using non-virtual link configurations

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

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

```
exit
```

All the sub ring nodes part of a sub-ring with non-virtual link should be configured with the "sub-ring non-virtual-link" option.

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

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

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

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

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

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

2.12.4 G.8032 OAM considerations

Ethernet CFM can be enabled on each individual path under an Ethernet ring. Only Down MEPs can be configured on each of them and CCM sessions can be enabled to monitor the liveness of the path using an interval of 100 ms. Different CCM intervals can be supported on path A and path B in an Ethernet ring. CFM is optional if 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.

2.12.5 G.8032 QoS considerations

Ethernet ring CC messages transmitted over the SAP queues using the default egress QoS policy will use NC (network class) as a forwarding class. If user traffic is assigned to the NC forwarding class, it will compete for the same bandwidth resources with the Ethernet CCMs. 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.



Note:

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

2.12.6 G.8032 service and solution combinations

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

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

2.12.7 G.8032 configuration guidelines

The following are the configuration guidelines for G.8032:

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

2.13 Layer 2 control processing

Operators providing Epipe and VPLS services need to be able to transparently forward Layer 2 control processing (L2CP) 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. For more information, see the [L2PT and BPDU translation](#).



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 L2CP 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 SAP/SDPs for redundant connectivity). That is, if the VPLS service is used for multi-point connectivity, Nokia does not recommend using 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 preceding protocols, protocols not supported on 7210 (for example, GARP, GVRP, ELMI, and others) are transparently forwarded in case of a VPLS service. These protocols are transparently forwarded if a null SAP, dot1q default SAP, dot1q explicit null SAP or 0.* SAP is configured on the port and 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 lists the L2CP support on 7210 SAS platforms.

Table 9: L2CP support on the 7210 SAS

Packet Type	7210 SAS-D, 7210 SAS-Dxp	7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C
LACP	Option to tunnel or discard or peer	Option to tunnel or discard or peer


Packet Type	7210 SAS-D, 7210 SAS-Dxp	7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C
Dot1x	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
EFM	Option to tunnel or discard or peer	Option to tunnel or discard or peer
L2PT	Supported	Supported
BPDU Tunneling	Supported	Supported
xSTP	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

2.14 Service creation process overview

This section provides an overview of the service creation process with access-uplink ports. The following figure shows the overall process to provision core and subscriber services.

⁹ See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Interface Configuration Guide* for more information about the options available for LLDP tunneling.

Figure 13: Service creation and implementation flow



2.14.1 Deploying and provisioning services

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

- [Phase 1: core network construction](#)
- [Phase 2: service administration](#)
- [Phase 3: service provisioning](#)

2.14.1.1 Phase 1: core network construction

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

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

2.14.1.2 Phase 2: service administration

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

- Configure group and user access privileges.

- Build templates for QoS, filter and/or accounting policies needed to support the core services.

2.14.1.3 Phase 3: service provisioning

For service provisioning, the following tasks should be completed:

- 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.

2.14.2 Configuration notes

This section describes service configuration caveats.

2.14.2.1 General

Service provisioning tasks can be logically separated into two main functional areas, core tasks and subscriber 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.

Subscriber services tasks include the following:

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

2.15 Configuring global service entities with CLI

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

2.15.1 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.

2.16 Basic configuration

The most basic service configuration must have the following:

- customer ID

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

2.17 Common configuration tasks

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

2.17.1 Configuring customer accounts

Subscribers

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

- description
- contact name
- telephone number

2.17.1.1 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 "Alcatel Customer"
      contact "Technical Support"
      phone "650 555-5100"
    exit
...
-----
A:A:ALA-12>config>service#
```

2.18 Service creation process overview on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C


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



Note:

Service creation with MPLS uplinks using SDPs for interconnecting routers is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

Figure 14: Service creation and implementation flow



2.18.1 Deploying and provisioning services on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP + 8C

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.

2.18.1.1 Phase 1: core network construction

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

- Build the IP or IP/MPLS core network.
- Configure routing protocols.
- Configure MPLS LSPs (if MPLS is used).
- Configure access-uplink SAPs (if only Layer 2 uplinks are needed).

2.18.1.2 Phase 2: service administration

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

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

2.18.1.3 Phase 3: service provisioning

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

2.18.2 Configuration notes

This section describes service configuration caveats.

2.18.2.1 General

Service provisioning tasks can be logically separated into two main functional areas, core tasks and subscriber 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 (if using MPLS uplinks)

Subscriber services tasks include the following:

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

2.18.3 Configuring global service entities with CLI

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

2.18.3.1 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.

2.18.4 Basic configuration

The most basic service configuration must have the following:

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

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

2.18.5 Common configuration tasks

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

2.18.5.1 Configuring customer accounts

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

- description

- contact name
- telephone number

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

2.18.5.3 Configuring an SDP

The most basic SDP must have the following:

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

2.18.5.4 SDP configuration tasks

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

Consider the following SDP characteristics:

- SDPs can be created as MPLS.
- Each distributed service must have an SDP defined for every remote router to provide VLL, VPLS, IES and VPRN services.
- A distributed service must be bound to an SDP. By default, no SDP is associated with a service. 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.



Note:

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

To configure a basic SDP, perform the following steps:

- Specify an originating node.
- Create an SDP ID.
- Specify an encapsulation type.
- Specify a far-end node.

2.18.6 Configuring an SDP

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



Note:

When you specify the far-end IP address, you are creating the tunnel; 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-K 2F6C4T, K 3SFP+ 8C MPLS Guide* for configuration and command information.

Use the following syntax to create an MPLS SDP.

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

Example

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

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

```
        timeout 5
        no message-length
        exit
-----
*A:K-SASK12>config>service>sdp$
```

2.18.7 Configuring a mixed-LSP SDP

The following shows 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.

2.19 ETH-CFM

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 base ETH-CFM configuration, that 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-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Services Guide* for more 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.

The supported features vary across services and depending on different 7210 SAS platforms.

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-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C 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 10: Service endpoint and MEP/MIP acronyms

Acronym	Expansion
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)
GNM	Generic Notification Message
LBM	Loopback Message
LBR	Loopback Reply
LTM	Linktrace Message
LTR	Linktrace Reply
ME	Maintenance Entity
MA	Maintenance Association
MA-ID	Maintenance Association Identifier
MD	Maintenance Domain
MEP	Maintenance Association 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)

Acronym	Expansion
SLM	Synthetic Loss Message (Y.1731)
SLR	Synthetic Loss Reply (Y.1731)

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

Figure 15: Ethernet OAM model for Ethernet access - business

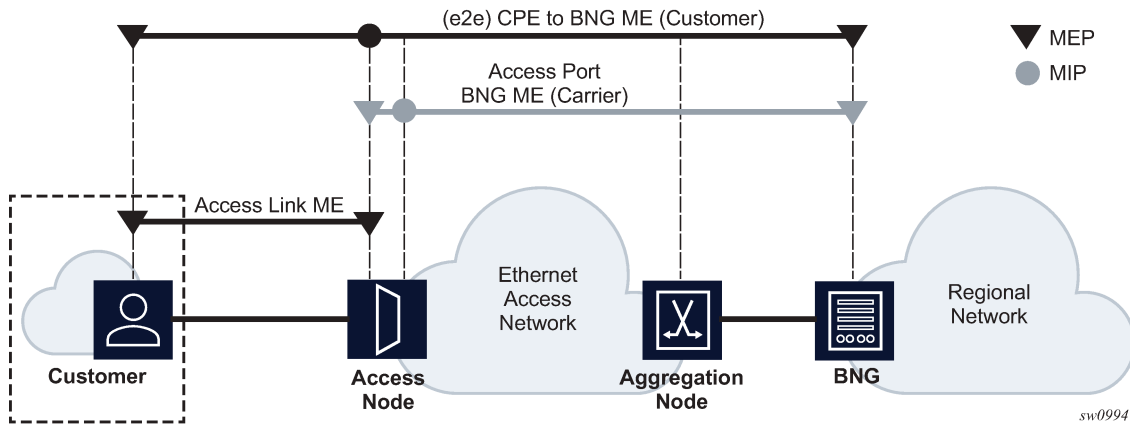
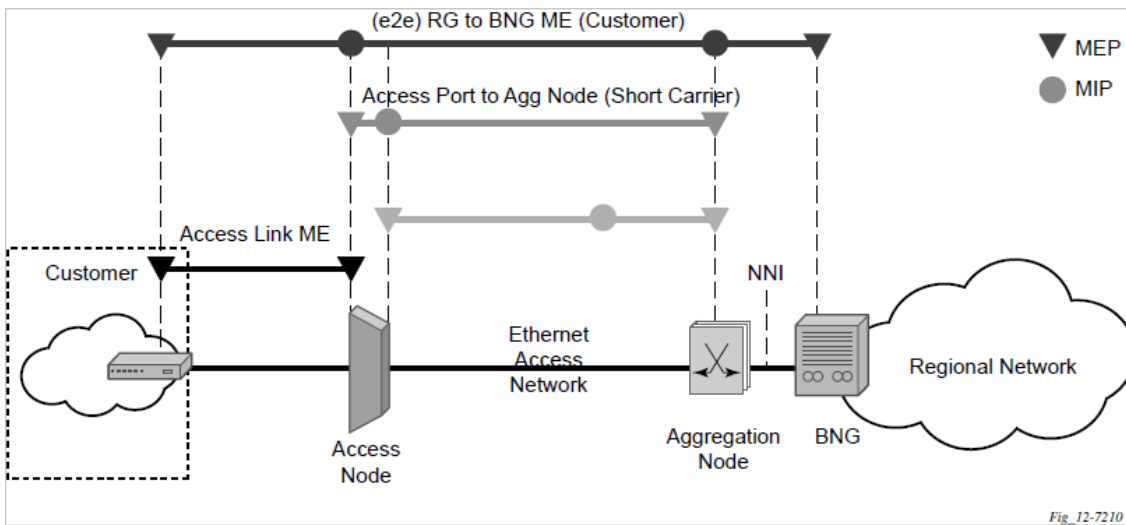


Figure 16: Ethernet OAM model for Ethernet access - wholesale



The following functions are supported:

- On the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T, CFM can be enabled or disabled on a per-SAP basis.

- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, CFM can be enabled or disabled on a per-SAP and per-SDP binding basis.
- The eight ETH-CFM levels are suggested to be broken up numerically between customer 7 to 5, service provider 4 to 3 and operator 2 to 1:
 - Level 0 typically is meant to monitor direct connections without any MIPs and should be reserved for port-based G8032 MEPs.
- Down MEP and Up MEP with an MEP-ID on a SAP/SDP binding for each MD level can be configured, modified, or deleted. Each MEP is uniquely identified by the MA-ID and MEP-ID tuple:
 - MEP creation on a SAP is only allowed 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 platforms. The following table lists the support available for MEP on different 7210 platforms.
 - On the 7210 SAS-D and 7210 SAS-Dxp, Up MEPs cannot be created by default on system bootup. See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C OAM and Diagnostics Guide* for more information about ETH-CFM configuration guidelines.
- 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.
 - The 7210 SAS-D and 7210 SAS-Dxp have the notion of ingress and egress MIPs. That is, on these platforms, MIPs are not bidirectional by default. 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 11: ETH-CC defect condition groups](#). See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C OAM and Diagnostics Guide* for more information.

2.19.1 Common actionable failures

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 11: ETH-CC defect condition groups

Defect	Low priority defect	Description	Causes	Priority
DefNone	N/A	No faults in the association	Normal operations	N/A
DefRDICCM	allDef	Remote Defect Indication	Feedback mechanism to inform that unidirectional faults exist. It provides the feedback loop to the node with the unidirectional failure conditions.	1

Defect	Low priority defect	Description	Causes	Priority
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 service, incorrect association identification.	5

2.19.2 MEP and MIP support

See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C OAM and Diagnostics Guide* for more information about ETH-CFM support for the different services and endpoints.

2.19.3 Configuring ETH-CFM parameters



Note:

See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C 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#
```

Defining the MEP and configuring service-specific ETH-CFM parameters is performed within the service on the specific SAP or SDP binding.

Example

The example using the service VPLS 100 shows this configuration on the SAP.

```
##*A:ALU-7_A>config>service# info
-----
      vpls 100 customer 1 create
      description "VPLS service 100 - Used for MEP configuration example"
      sap 2/2/1:20 create
      description "2/2/1:20"
      eth-cfm
      mep 1 domain 1 association 1 direction down
      no shutdown
      exit
      exit
      exit
      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 samples of the global ETH-CFM configuration output and the advanced Y.1731 functions that can be configured. The configuration will reject the configuration of Y.1731 functions within an IEEE 802.1ag context.

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

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



Note:

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

2.20 Service management tasks

This section describes the service management tasks.

2.20.1 Modifying customer accounts

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

To display a list of customer IDs, use the **show service customer** command.

Enter the parameters (*description*, *contact*, *phone*), 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"
```

2.20.2 Deleting customers

The **no 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
```

```
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
```

2.21 Global services command reference

2.21.1 Command hierarchies

- [Customer commands](#)
- [Pseudowire \(PW\) commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [SDP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [SAP commands for 7210 SAS-D and 7210 SAS-Dxp](#)
- [SAP commands for 7210 SAS-K 2F1C2T](#)
- [SAP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
- [ETH-CFM configuration commands](#)
- [Show commands](#)

2.21.1.1 Customer commands

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

2.21.1.2 Pseudowire (PW) commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
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
- hash-label [signal-capability]
- no hash-label
- [no] force-vlan-vc-forwarding
- igmp-snooping
- [no] disable-router-alert-check
- import policy-name
- no import
- last-member-query-interval 1/10 seconds
- no last-member-query-interval
- max-num-groups max-num-groups
- no max-num-groups
- query-interval seconds
- no query-interval
- query-response-interval seconds
- no query-response-interval
- robust-count robust-count
- no robust-count
- [no] send-queries
- version version
- no version
- limit-mac-move {blockable | non-blockable}
- no limit-mac-move
- [no] mac-pinning
- max-nbr-mac-addr table-size
- no max-nbr-mac-addr
- split-horizon-group group-name
- no split-horizon-group
- description description-string
- no description
- vc-type {ether | vlan}
- vlan-vc-tag 0..4094
- no vlan-vc-tag
```

2.21.1.3 Pseudowire (PW) routing commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
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
```

2.21.1.4 SDP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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


2.21.1.5 SAP commands for 7210 SAS-D and 7210 SAS-Dxp

```
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
- interface address {ip-address/mask | ip-address netmask} [broadcast {all-ones |
host-ones}]
- sap sap-id [create]
- no sap sap-id
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | any | dot1q-preserve}] [customer-vid vlan-id]
- no vpls service-id
- sap sap-id [create] [eth-ring ring-index]
- no sap sap-id
```

2.21.1.6 SAP commands for 7210 SAS-K 2F1C2T

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

2.21.1.7 SAP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

2.21.1.8 ETH-CFM configuration commands



Note:

For command descriptions, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C 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-id-identifier bridge-id
- id-permission {chassis}
- no id-permission
- mhf-creation {default | none | explicit| static}
- no mhf-creation
- mip-ltr-priority priority
- vlan vlan-id
- no vlan
- ccm-interval {10ms | 100ms | 1 | 10 | 60 | 600}
- no ccm-interval
- [no] remote-mepid mep-id
- slm
- inactivity-timer timer
- no inactivity-timer
- system
- enable-dmm-version-interop
- no enable-dmm-version-interop
- sender-id local local-name
- sender-id system
- no sender-id
    
```

2.21.1.9 Show commands

```

show
- service
- customer [customer-id] [site customer-site-name]
- id service-id
- sdp sdp-id keep-alive-history
- sdp far-end ip-address keep-alive-history
- sdp [sdp-id] [detail]
- sdp far-end ip-address [detail]
- pw-template-using sdp-id pw-port [pw-port-id]
- pw-template-using [consistent | inconsistent | na] egressifs
- pw-template-using sdp-id keep-alive-history
- pw-template-using far-end ip-address keep-alive-history
- pw-template-using [sdp-id] [detail]
- pw-template-using far-end ip-address [detail]
- sdp-using [sdp-id[:vc-id] | far-end ip-address]
- service-using [epipe] [ies] [vpls] [mirror] [customer customer-id]
- eth-ring [status]
- eth-ring ring-index hierarchy
- eth-ring ring-index [path {a | b}]
- eth-cfm
- association [ma-index] [detail]
- cfm-stack-table [port [port-id [vlan vlan-id]]][level 0..7] [direction down]
    
```

```
- mep mep-id domain md-index association ma-index [loopback] [linktrace]
- mep mep-id domain md-index association ma-index remote-mepid mep-id | all-remote-
mepids
- mep mep-id domain md-index association ma-index eth-test-results [remote-peer mac-
address]
- mep mep-id domain md-index association ma-index one-way-delay-test [remote-peer mac-
address]
- mep mep-id domain md-index association ma-index two-way-delay-test [remote-peer mac-
address]
- pw-routing {local-prefix | static-route | paths | all}
- pw-routing route-table [all-routes]
- pw-routing route-table summary
```

2.21.1.10 Tools perform commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
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]]
```

2.21.2 Command descriptions

2.21.2.1 Global service configuration commands

2.21.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

config>eth-cfm>mep

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.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up, 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

Service Admin State

While the service is shut down, all customer packets are dropped and counted as discards for billing and debugging purposes.

description

Syntax

description *description-string*

no description

Context

config>service>customer

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.

This 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.

2.21.2.1.2 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

Description

This command configures 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.

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

Description

This command configures contact information for a customer.

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

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

Description

This command adds telephone number information for a customer ID.

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

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.

2.21.2.1.3 Pseudowire commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

pw-routing

Syntax

pw-routing

Context

config>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

boot-timer

Syntax

boot-timer *secs*

no boot-timer

Context

config>service>pw-routing

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 the default.

Default

10

Parameters

secs

Specifies 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 SAS node. The global ID can contain the 2-octet or 4-octet value of the provider Autonomous System (AS) number. The presence of a global ID based on the provider AS number ensures that the All for spoke-SDPs configured on the node is globally unique.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

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

Default

no advertise-bgp

Parameters

rd

Specifies a 32-bit prefix for the 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 AS number. The presence of a global ID based on the provider AS number ensures that the All for spoke-SDPs configured on the node are globally unique.

Values (6 bytes, other 2 bytes of type 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

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

Values	<i>community:</i>	{2-byte-as-number:comm-va1}
	<i>2-byte-asnumber:</i>	1 to 65535
	<i>comm-va1:</i>	0 to 65535

path

Syntax

path name [create]

no path name

Context

config>service>pw-routing

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures an explicit path between this 7210 SAS T-PE and a remote 7210 SAS T-PE. For each path, one or more intermediate S-PE hops must be configured. A path can be used by multiple multi-segment pseudowires. Paths are used by a 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

path-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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 shut down before deleting the hop from the hop list. The **no hop** *hop-index* command does not result in any action, except for a warning message on the console indicating that the path is administratively up.

Default

no hop

Parameters

hop-index

Specifies a locally significant numeric identifier for the hop. The hop index is used to order the specified hops. 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 *count*

no retry-count

Context

config>service>pw-routing

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This optional command specifies the number of attempts that are made 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 to the default value.

Default

30

Parameters

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 to the default value.

Default

30

Parameters

secs

Specifies the initial retry-timer value in seconds.

Values 10 to 480

spe-address

Syntax

spe-address *global-id:prefix*

no spe-address

Context

config>service>pw-routing

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description



Note:

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C can only be configured as T-PE nodes and not as S-PE nodes. This command allows the 7210 SAS-K 2F6C4T or 7210 SAS-K 3SFP+ 8C to be configured as a T-PE node to determine the S-PE node to use for multi-segment PWs.

This command configures a single S-PE address for the node to be used for dynamic MS-PWs. This value is used for the PW switching point TLV used in LDP signaling, and is the value used by PW status signaling to indicate the PE that originates a PW status message. Configuration of this command 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 local prefixes cannot be configured on a T-PE. A 7210 SAS node sends a label release for any label mappings received for FEC129 All type 2.

The S-PE address cannot be changed unless the dynamic ms-pw configuration is removed.
Changing the S-PE address results in all dynamic MS-PWs for which this node is an S-PE being released. Nokia recommends that the S-PE address be configured for the life of an MS-PW configuration after rebooting the 7210 SAS.
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 AS number.

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

Values		
	<i>global-id:</i>	1 to 4294967295
	<i>prefix:</i>	1 to 4294967295
	<i>ipaddress:</i>	a.b.c.d

static-route

Syntax

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

Context

config>service>pw-routing

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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_addresses
```

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		
	<i>route-name:</i>	<i>global-id:prefix:next-hop-ip_addr</i>
	<i>global-id:</i>	0 to 4294967295
	<i>prefix:</i>	a.b.c.d. 0 to 4294967295
	<i>next-hop-ip_addr:</i>	a.b.c.d

pw-template

Syntax

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

Context

config>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures an SDP template.

Parameters

policy-id

Specifies a number that uniquely identifies a template for the creation of an SDP.

Values 1 to 2147483647

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 for the instantiation of the SDP.

create

Keyword required 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 value.

Default

no control-word

2.21.2.1.4 SDP commands

sdp

Syntax

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

no sdp *sdp-id*

Context

config>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 service without having to specifically define far-end SAPs. Each SDP represents a method to reach a 7210 SAS router.

The 7210 SAS supports only Multi-Protocol Label Switching (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, 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, generating an error message specifying the first bound service found during the deletion process. If the specified *sdp-id* does not exist, an error is generated.

Parameters

sdp-id

Specifies the SDP identifier.

Values 1 to 17407

mpls

Keyword to specify that the SDP uses MPLS encapsulation.

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

config>service>sdp

config>service>pw-template

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 an 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 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

When the **no collect-stats** command is issued, the statistics are still accumulated by the IOM cards.

However, the CPU 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

disable-learning

Syntax

[no] **disable-learning**

Context

config>service>pw-template

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables learning of new MAC addresses.

This parameter is mainly used in conjunction with the **discard-unknown** command.

The **no** form of this command enables learning of MAC addresses.

Default

no disable-learning (normal MAC learning is enabled)

discard-unknown-source

Syntax

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

Context

config>service>pw-template

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures that packets received with an unknown source MAC address are dropped only if the maximum number of MAC addresses has 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 address to be forwarded by destination MAC addresses.

Default

no discard-unknown

egress

Syntax

egress

Context

config>service>pw-template

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure spoke-SDP binding egress filter parameters.

ingress

Syntax

ingress

Context

config>service>pw-template

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure spoke-SDP binding ingress filter parameters.

hash-label

Syntax

hash-label [signal-capability]

no hash-label

Context

config>service>pw-template

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the hash label on VLL or VPLS services that are bound to RSVP SDP, 3107 BGP SDP, segment routing, or LDP SDP, using the **auto-bind** 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. The ingress datapath appends the hash label at the bottom of the stack (BoS) and sets the S-bit to one (1).



Note:

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

Packets generated in the CPM that are forwarded with a label 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 0.

Signaling of the hash label capability is enabled by adding the **signal-capability** option under the VLL spoke-SDP, VPLS spoke-SDP or mesh SDP interface, or PW template instance. In this case, the decision of the local PE to insert the hash label on the user and control plane packets is determined by the outcome of the signaling process and can override the local PE configuration. The following process flow applies when the **hash-label** and **signal-capability** options 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 the specific spoke-SDP or mesh SDP.
- If a 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 the user and control plane packets that it forwards on the spoke-SDP or mesh SDP. The local PE also drops user and control plane packets received from a remote PE if they include a hash label. The dropped packets may be caused by any of the following:
 - a remote 7210 SAS PE that does not support the **hash-label** command
 - a remote 7210 SAS PE that has the **hash-label** command enabled but does not support the **signal-capability** option
 - a remote 7210 SAS PE that supports the **hash-label** command and the **signal-capability** option, but the user did not enable them due to a misconfiguration
- If the remote PE sends a flow label sub-TLV in the PW ID FEC element with T=TRUE and R=TRUE, the local PE enables the hash label capability. Consequently, the local PE node inserts a hash label in the user and control plane packets that it forwards on the spoke-SDP or mesh SDP. The local PE node also accepts user and control plane packets from the remote PE with a hash label. The local PE node drops user and control plane packets from the remote PE without a hash label.

If the **hash-label** command is enabled on the local PE with the **signal-capability** option configured and on the remote PE without the **signal-capability** option configured on the spoke-SDP or mesh-SDP, the hash label is included in the pseudowire packets received by the local PE. These packets must be dropped. To resolve this situation, you must disable the **signal-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 is not enabled on the local configuration of the spoke-SDP or mesh-SDP at the remote PE, the hash label is not included in the pseudowire received by the local PE.

If the **signal-capability** option is enabled or disabled in the CLI, 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 7750 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 to 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. 7210 SAS devices do not set the MSB in the hash label value for service traffic. Therefore, the user must ensure that both ends are correctly configured to either process hash labels or disable them. The MSB bit is set for MPLS/OAM traffic on 7210 SAS devices.
- The **cpe-ping**, **mac-ping**, and **svc-ping** commands are not supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C when the **hash-label** command is enabled.

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

Default

no hash-label

Parameters

signal-capability

Keyword that specifies to enable the signaling and negotiation of the use of the hash label between the local and remote PE nodes.

force-vlan-vc-forwarding

Syntax

[no] force-vlan-vc-forwarding

Context

config>service>pw-template

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command forces vc-vlan-type forwarding in the datapath for spoke and mesh SDPs that 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

no force-vlan-vc-forwarding

split-horizon-group

Syntax

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

Context

config>service>vpls

config>service>pw-template

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

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

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

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

Parameters

group-name

Specifies the name of the split horizon group to which the SAP or spoke-SDP belongs.

create

Mandatory keyword to create a split-horizon group.

vc-label

Syntax

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

Context

config>service>pw-template>ingress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the ingress VC label.

Parameters

vc-label

Specifies a VC ingress value that indicates a specific connection.

Values 2048 to 18431

limit-mac-move

Syntax

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

no limit-mac-move

Context

config>service>pw-template

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

Default

blockable

Parameters

blockable

Keyword to specify the agent will monitor the MAC relearn rate and will block it when the relearn rate is exceeded.

non-blockable

Keyword to specify that a SAP will not be blocked, and another blockable SAP will be blocked instead.

vc-type

Syntax

vc-type {ether | vlan}

Context

config>service>pw-template

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 that represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment.

A change of the 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-mps*.

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

Parameters

ether

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

vlan

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

vlan-vc-tag

Syntax

vlan-vc-tag *vlan-id*

no vlan-vc-tag

Context

config>service>pw-template

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

vlan-id

Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

Values 0 to 4094

adv-mtu-override

Syntax

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

Context

config>service>sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 whose 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/VP RN spoke interface using this SDP ID. The router continues to signal a VC MTU equal to the net IP interface MTU, which is $\min(\text{ip-mtu}, \text{sdp operational path mtu} - \text{Layer 2 headers})$. The router also continues to ensure that the advertised MTU values of both PE routers match or the spoke-SDP goes operationally down.

The **no** form of this 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T, K 3SFP+ 8C Routing Protocols Guide*.

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

Default

no bgp-tunnel

far-end

Syntax

far-end *ip-address* **node-id** *node-id*

no far-end

Context

config>service>sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the system IP address of the far-end destination router for the SDP that is the termination point for a service.

The far-end IP address must be explicitly configured. The destination IP address must be a 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 will generate an error if used in the **no far-end** command. The SDP must be administratively disabled using the **config service sdp shutdown** command before the **no far-end** command can be executed. Removing the far end IP address causes all *lsp-name* associations with the SDP to be removed.

Parameters

ip-address

Specifies the system address of the far-end device 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 or an unsigned 32-bit integer.

Values	a.b.c.d
	1 to 4294967295

metric

Syntax

metric *metric*

no metric

Context

config>service>sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the metric 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures 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. Users can configure one or two types of LSPs under the same SDP. Without this command, these commands are mutually exclusive.

Users can configure an RSVP LSP as a primary LSP type with an LDP LSP as a backup type. Users 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, 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 will activate it to forward service packets according to the following priority order:

- RSVP LSP type — one RSVP LSP can be configured per SDP.

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 the LDP LSP is not available, 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, 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. 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. See the *7210 SAS-K 2F6C4T, K 3SFP+ 8C MPLS Guide* for more information about this timer.

If the value of the **sdp-revert-time** timer is changes, 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 prefers 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.

LDP/BGP SDPs behave differently compared to RSVP/ LDP SDPs. For a /32 prefix, only a single route exists in the routing table: the IGP route or the BGP route. Either the LDP FEC or the BGP label route is active at any time. The tunnel table must be reprogrammed each time a route is deactivated and the other is activated. The SDP revert-time cannot be used because both LSP types cannot be active for the same /32 prefix.

The **no** form of this command disables the mixed-LSP mode of operation. The user must 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 this 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

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

ldp

Syntax

[no] ldp

Context

config>service>sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 **isp** 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 isp *isp-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-isp** option is enabled on the SDP.

Default

no ldp

isp

Syntax

isp *isp-name*

no isp *isp-name*

Context

config>service>sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command creates associations between one LSP and an MPLS SDP. This command is implemented only on MPLS-type encapsulated SDPs

In MPLS SDP configurations 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 shut down, 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 association from an SDP. If the *lsp-name* does not exist as an association or as a configured LSP, no error is returned. An *lsp-name* must be removed from all SDP associations before the *lsp-name* can be deleted from the system. The SDP must be administratively disabled (**shutdown**) before the last **lsp-name** associated 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* exists 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 be changed only while the administrative status of the SDP is down.

The SDP must be administratively shut down before the **signaling** command can be modified and re-enabled.

Default

tldp

Parameters

off

Keyword to specify 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

Keyword to specify that ingress and egress pseudowire signaling using T-LDP is enabled.

bgp

Keyword to specify that ingress and egress pseudowire signaling using BGP is enabled. This keyword 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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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**).

Default

no sr-isis

sr-ospf

Syntax

[no] **sr-ospf**

Context

config>service>sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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**).

Default

no sr-ospf

path-mtu

Syntax

path-mtu *bytes*

no path-mtu

Context

config>service>sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the Maximum Transmission Unit (MTU) in bytes that the 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, such as 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.

By default, the path-mtu defined on the system for the type of SDP is used.

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

Parameters

bytes

Specifies the number of bytes in the path MTU.

Values 576 to 9194

2.21.2.1.5 SDP keepalive commands

keep-alive

Syntax

keepalive

Context

config>service>sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 sent only 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.

hello-time

Syntax

hello-time *seconds*

no hello-time

Context

config>service>sdp>keep-alive

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 to the default value.

Default

hello-time 10

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

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

When an SDP keepalive response is received that indicates an error condition or the **max-drop-count** keepalive messages receive no reply, the *sdp-id* is immediately brought operationally down. If a keepalive response is received that indicates the error has cleared, the *sdp-id* becomes 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** value to the default.

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** will be enforced for the SDP ID.

Values 0 to 3600

Default 10

max-drop-count

Syntax

max-drop-count *count*

no max-drop-count

Context

config>service>sdp>keep-alive

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the number of consecutive SDP keepalive failed request attempts or remote replies that can be missed after which the SDP is operationally downed.

If the **max-drop-count** consecutive keepalive request messages cannot be sent or no replies are received, the SDP ID will be brought operationally down by the keepalive SDP monitoring.

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

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

Default 3

message-length

Syntax

message-length *octets*

no message-length

Context

config>service>sdp>keep-alive

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

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

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.

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.

Values 40 to 9198

Default 0

timeout

Syntax

timeout *timeout*

no timeout

Context

config>service>sdp>keep-alive

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

Parameters

timeout

Specifies the timeout in seconds.

Values 1 to 10

Default 5

2.21.2.1.6 Tools perform commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

tools

Syntax

tools

Context

root

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context enable useful tools for debugging purposes.

perform

Syntax

perform

Context

tools

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context enable tools to perform specific tasks.

service

Syntax

services

Context

tools>perform

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure tools for services.

id

Syntax

id *service-id*

Context

tools>perform>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies a service for which to configure tools.

Parameters

service-id

Specifies an existing service ID.

Values 1 to 2147483647

endpoint

Syntax

endpoint *endpoint-name*

Context

tools>perform>service>id

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures tools for a specific VLL service endpoint.

Parameters

endpoint-name

Specifies 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

Parameters

sdp-id

Specifies an existing spoke-SDP ID for the service.

Values 1 to 17407

vc-id

Specifies a virtual circuit ID.

Values 1 to 4294967295

spoke-sdp-fec spoke-sdp-fec-id

Specifies the spoke-SDP FEC ID for a FEC129 All Type 2 spoke-SDP. This parameter is mutually exclusive with the *sdp:vc-id* used for a FEC 128 spoke-SDP4.

Output

The following output is an example of service endpoint 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 *policy-id* [**allow-service-impact**]

Context

tools>perform>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command reevaluates the pseudowire template policy.

Parameters

policy-id

Specifies the pseudowire template policy.

Values 1 to 2147483647

allow-service-impact

Keyword to reevaluate the pseudowire template even if the service is impacted.

eval-expired-fec

Syntax

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

eval-expired-fec *all*

Context

tools>perform>service>pw-routing

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command resets the retry counter and retry timer for the specified spoke-SDP and attempts to reestablish the spoke-SDP.

spoke-sdp-fec-release

Syntax

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

Context

```
tools>perform>service
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command clears the MS-PW bindings associated with a particular SAll or TAll on an SPE.

2.21.2.2 Show, clear, debug commands

2.21.2.2.1 Services 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

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 that 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 12: 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 : 8
Contact     : Customer Service
Description : IES Customer
Phone      : (678) 555-1212

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

Customer-ID : 94043
Contact     : Test Engineer on Duty
Description : TEST Customer
Phone      : (789) 555-1212
-----
Total Customers : 8
-----
```

```

*A:ALA-12#
*A:ALA-12# show service customer 274
=====
Customer 274
=====
Customer-ID : 274
Contact      : Mssrs. Beaucoup
Description  : ABC Company
Phone       : 650 123-4567
-----
Multi Service Site
-----
Site        : west
Description : (Not Specified)
=====
*A:ALA-12#
    
```

Table 12: 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 or 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 multi-service site.
Assignment	The port ID, MDA, or card number, where the SAPs members of this multi- service site are defined.
I. Sched Pol	The ingress QoS scheduler policy assigned to this multi-service site.
E. Sched Pol	The egress QoS scheduler policy assigned to this multi-service site.
Service-ID	The ID that uniquely identifies a service.
SAP	Specifies the SAP assigned to the service.

fdb-mac

Syntax

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

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Keyword that shows the amount of time until the MAC address is aged out.

Output

The following output is an example of MAC address FDB entry information, and [Table 13: 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/    Last Change
          MAC                Source-Identifier  Expiry
-----
103      12:34:56:78:90:0f  sap:1/1/7:0       Static    02/02/2009 09:27:57
700      90:30:ff:ff:ff:8f  cpm                Host      02/02/2009 09:27:57
-----
No. of Entries: 2
=====
*A:ALA-48#
```

Table 13: 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/Age	Static — FDB entries created by management. Learned — Dynamic entries created by the learning process. OAM — Entries created by the OAM process. H — Host, the entry added by the system for a static configured subscriber host. D or DHCP — DHCP-installed MAC. Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease. P — Indicates the MAC is protected by the MAC protection feature.
Last Change	The time when the specific row entry was last changed.

sdp-using

Syntax

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

Context

show>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 for services using SDP address options, and [Table 14: Output fields: SDP-using](#) describes the output fields.

Sample output

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

Table 14: 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.
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] [mirror] [i-vpls] [m-vpls] [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 specific usage properties. If no optional parameters are specified, all services defined on the system are displayed.

Parameters

epipe

Displays matching Epipe services.

ies

Displays matching IES instances.

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 information for services matching specific usage properties, and [Table 15: 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           1
2           VPLS      Up   Up   1           2
3           VPLS      Up   Up   1           3
4           VPLS      Up   Up   1           4
2147483648  IES       Up   Down 1         _tmnx_InternalIesService
2147483649  intVpls   Up   Down 1         _tmnx_InternalVplsService
-----
Matching Services : 6
=====
```



```
*7210SAS>show>service#
```

Table 15: 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.

id

Syntax

`id service-id`

Context

`show>service`

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information about a specified service ID.

Parameters

service-id

Specifies the unique service ID number that identifies the service in the service domain.

Values *service-id*: 1 to 2147483647
svc-name: a string up to 64 characters

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

Description

This command displays the Ethernet ring information.

Parameters

ring-index

Specifies the ring index.

Values 1 to 128

status

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

hierarchy

Displays Ethernet ring hierarchical relationships.

path {a | b}

Displays information related to the configured Ethernet rings.

Output

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

- [Sample output](#) , [Table 16: Output fields: eth-ring status](#)
- [Sample output](#), [Table 17: Output fields: eth-ring](#)

Sample output

```
*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 16: Output fields: eth-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.

Sample output

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

=====
Ethernet Ring 10 Information
=====
Description      : (Not Specified)
Admin State      : Down           Oper State       : Down
Node ID          : 00:25:ba:03:48:04
Guard Time      : 5 deciseconds  RPL Node        : rplNone
Max Revert Time : 300 seconds    Time to Revert  : N/A
CCM Hold Down Time : 0 centiseconds CCM Hold Up Time : 20 deciseconds

APS Tx PDU       : N/A
Defect Status    :

-----
Ethernet Ring Path Summary
-----
Path Port   Raps-Tag   Admin/Oper   Type   Fwd State
-----
a -         -         -/-         -      -
b -         -         -/-         -      -
=====

*A:NS1015C0821>show#
```

Table 17: Output fields: eth-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.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays PW routing information at this 7210 node.

Parameters

local-prefix | static-route | paths | all

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

all-routes

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

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays information about PW templates.

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
-----
```

pw-template-using

Syntax

pw-template-using

Context

show>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays information about services that are using a PW template.

Output

The following output is an example of information about services that are using a PW template, and [Table 18: Output fields: PW template using](#) describes the output fields.

Sample output

```
*A:dut-a>show>service# pw-template-using
=====
Services using PW Template
=====
PW Template Id      SvcId
-----
5                    5
=====
*A:dut-a>show>service#
```

Table 18: Output fields: PW template using

Label	Description
PW Template Id	The PW template identifier.
Svc ID	The service identifier.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 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 for the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP + 8C, and [Table 19: Output fields: service SDP](#) describes the output fields.

Sample output

```
*A:Dut-D# show service id 1 sdp 17407:4294967294 detail
=====
Service Destination Point (Sdp Id : 17407:4294967294) Details
=====
-----
Sdp Id 17407:4294967294  -(not applicable)
-----
Description      : (Not Specified)
SDP Id           : 17407:4294967294          Type           : VplsPmsi
Split Horiz Grp  : (Not Specified)
VC Type          : Ether                    VC Tag         : n/a
Admin Path MTU   : 9194                      Oper Path MTU  : 9194
Delivery         : MPLS
Far End          : not applicable
Tunnel Far End   : n/a                      LSP Types      : None
Hash Label       : Disabled                  Hash Lbl Sig Cap : Disabled
Oper Hash Label  : Disabled
Admin State      : Up                       Oper State      : Up
Acct. Pol        : None                     Collect Stats   : Disabled
Ingress Label    : 0                       Egress Label   : 3
Ingr Mac Fltr-Id : n/a                      Egr Mac Fltr-Id : n/a
```

```
Ingr IP Fltr-Id      : n/a          Egr IP Fltr-Id      : n/a
Ingr IPv6 Fltr-Id   : n/a          Egr IPv6 Fltr-Id   : n/a
Admin ControlWord    : Not Preferred Oper ControlWord    : False
Last Status Change  : 12/14/2012 12:42:22 Signaling           : None
Last Mgmt Change    : 12/14/2012 12:42:19 Force Vlan-Vc      : Disabled
Endpoint            : N/A          Precedence         : 4
PW Status Sig       : Enabled
Class Fwding State  : Down
Flags               : None
Time to RetryReset  : never        Retries Left       : 3
Mac Move            : Blockable     Blockable Level    : Tertiary
Local Pw Bits       : None
Peer Pw Bits        : None
Peer Fault Ip       : None
Peer Vccv CV Bits   : None
Peer Vccv CC Bits   : None
Application Profile  : None
Max Nbr of MAC Addr: No Limit      Total MAC Addr     : 0
Learned MAC Addr    : 0            Static MAC Addr    : 0

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

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

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

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

Statistics          :
I. Fwd. Pkts.      : 0            I. Dro. Pkts.       : 0
I. Fwd. Octs.      : 0            I. Dro. Octs.       : 0
E. Fwd. Pkts.      : 2979761        E. Fwd. Octets      : 476761760

-----
Control Channel Status
-----
PW Status           : disabled      Refresh Timer       : <none>
Peer Status Expire  : false        Clear On Timeout    : true

MCAC Policy Name    :
MCAC Max Unconst BW: no limit      MCAC Max Mand BW   : no limit
MCAC In use Mand BW: 0            MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0            MCAC Avail Opnl BW: unlimited

-----
RSVP/Static LSPs
-----
Associated LSP List :
```

```

No LSPs Associated

-----
Class-based forwarding :
-----
Class forwarding      : Disabled          EnforceDSTELspFc   : Disabled
Default LSP          : Uknwn              Multicast LSP       : None

=====
FC Mapping Table
=====
FC Name              LSP Name
-----
No FC Mappings

-----
Stp Service Destination Point specifics
-----
Stp Admin State      : Down                Stp Oper State      : Down
Core Connectivity    : Down
Port Role            : N/A                 Port State          : Forwarding
Port Number          : 0                   Port Priority       : 128
Port Path Cost       : 10                  Auto Edge           : Enabled
Admin Edge           : Disabled            Oper Edge           : N/A
Link Type            : Pt-pt               BPDU Encap         : Dot1d
Root Guard           : Disabled            Active Protocol    : N/A
Last BPDU from       : N/A                 Designated Port Id : N/A

Fwd Transitions      : 0                   Bad BPDUs rcvd     : 0
Cfg BPDUs rcvd       : 0                   Cfg BPDUs tx       : 0
TCN BPDUs rcvd       : 0                   TCN BPDUs tx       : 0
TC bit BPDUs rcvd    : 0                   TC bit BPDUs tx    : 0
RST BPDUs rcvd       : 0                   RST BPDUs tx       : 0

-----
Number of SDPs : 1
-----

=====
*A:Dut-B# show service sdp 204 detail
=====
Service Destination Point (Sdp Id : 204) Details
=====
Sdp Id 204 -10.20.1.4
-----
Description          : (Not Specified)
SDP Id               : 204                 SDP Source          : manual
Admin Path MTU       : 0                   Oper Path MTU       : 1492
Delivery              : MPLS
Far End               : 10.20.1.4
Tunnel Far End       : n/a                 LSP Types           : RSVP

Admin State          : Up                   Oper State          : Up
Signaling             : TLDP                Metric              : 0
Acct. Pol             : None                 Collect Stats       : Disabled
Last Status Change   : 02/12/2013 22:10:43 Adv. MTU Over.     : No
Last Mgmt Change     : 02/12/2013 22:09:55 VLAN VC Etype      : 0x8100
Bw BookingFactor     : 100                 PBB Etype           : 0x88e7
    
```

```

Oper Max BW(Kbps)      : 0           Avail BW(Kbps)       : 0
Net-Domain            : default      Egr Interfaces      : Consistent
Flags                 : None

Mixed LSP Mode Information :
Mixed LSP Mode        : Disabled     Active LSP Type     : RSVP

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

RSVP/Static LSPs
-----
Associated LSP List :
Lsp Name            : lsp-b2d
Admin State          : Up           Oper State          : Up
Time Since Last Tran*: 00h17m33s
-----
Class-based forwarding :
-----
Class forwarding     : Disabled     EnforceDSTELspFc  : Disabled
Default LSP          : Uknwn       Multicast LSP      : None
=====
FC Mapping Table
=====
FC Name              LSP Name
-----
No FC Mappings
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-B#

*A:ALA-12# show service sdp
=====
Services: Service Destination Points
=====
SdpId   Adm MTU   Opr MTU   IP address   Adm  Opr       Deliver Signal
-----
10      4462     4462     10.20.1.3   Up   Dn NotReady MPLS  TLDP
40      4462     1534     10.20.1.20  Up   Up        MPLS  TLDP
60      4462     1514     10.20.1.21  Up   Up        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:Dut-A# show service sdp 1 detail
=====
Service Destination Point (Sdp Id : 1) Details
    
```

```

-----
Sdp Id 1 -(10.20.1.3)
-----
Description          : epipe sdp 1 for lspId 00:00:00:01:00:00:00:00
SDP Id               : 1                      SDP Source          : manual
Admin Path MTU      : 0                      Oper Path MTU       : 1492
Far End              : 10.20.1.3             Delivery             : MPLS
Admin State         : Up                    Oper State          : Up
Signaling           : TLDP                  Metric              : 0
Acct. Pol           : None                  Collect Stats       : Disabled
Last Status Change  : 12/08/2008 22:54:30   Adv. MTU Over.     : No
Last Mgmt Change    : 12/08/2008 22:54:01   VLAN VC Etype      : 0x8100
Bw BookingFactor    : 100                   PBB Etype           : 0x88e7
Oper Max BW(Kbps)   : 1000                  Avail BW(Kbps)     : 1000
Flags                : None

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

Associated LSP LIST :
Lsp Name             : tof1
Admin State          : Up                    Oper State          : Up
Time Since Last Tran*: 00h04m01s
-----
Class-based forwarding :
-----
Class forwarding     : disabled              EnforceDSTELspFc  : disabled
Default LSP          : Uknwn                 Multicast LSP      : None
=====
FC Mapping Table
=====
FC Name              LSP Name
-----
No FC Mappings
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-A>config>service#

*A:ALA-12# show service sdp 8
=====
Service Destination Point (Sdp Id : 8)
=====
SdpId   Adm MTU   Opr MTU   IP address   Adm  Opr        Deliver Signal
-----
8       4462     4462     10.10.10.104  Up   Dn NotReady MPLS   TLDP
=====
*A:ALA-12#

*A:ALA-12#
=====
Service Destination Point (Sdp Id : 8) Details
=====
Sdp Id 8 -(10.10.10.104)
-----
Description          : MPLS-10.10.10.104
SDP Id               : 8
Admin Path MTU      : 0                      Oper Path MTU       : 0
Far End              : 10.10.10.104             Delivery             : MPLS
Admin State         : Up                    Oper State          : Down
    
```

```

Flags                : SignalingSessDown TransportTunnDown
Signaling            : TLDP                VLAN VC Etype      : 0x8100
Last Status Change   : 02/01/2007 09:11:39  Adv. MTU Over.     : No
Last Mgmt Change     : 02/01/2007 09:11:46
KeepAlive Information :
Admin State          : Disabled            Oper State          : Disabled
Hello Time           : 10                  Hello Msg Len       : 0
Hello Timeout        : 5                    Unmatched Replies   : 0
Max Drop Count       : 3                    Hold Down Time      : 10
Tx Hello Msgs        : 0                    Rx Hello Msgs       : 0
Associated LSP LIST :
Lsp Name             : to-104
Admin State          : Up                  Oper State          : Down
Time Since Last Tran*: 01d07h36m
=====
* indicates that the corresponding row element may have been truncated.
*A:ALA-12#

*A:MV-SR12>config>service>vprn# show service sdp 10 detail
=====
Service Destination Point (Sdp Id : 10) Details
=====
Sdp Id 10  -(200.20.1.201)
-----
Description          : (Not Specified)
SDP Id               : 10                  SDP Source          : manual
Admin Path MTU       : 0                    Oper Path MTU       : 9182
Far End              : 200.20.1.201        Delivery            : MPLS/LDP
Admin State          : Up                  Oper State          : Up
Signaling            : TLDP                Metric              : 0
Acct. Pol            : None                Collect Stats       : Disabled
Last Status Change   : 02/12/2010 22:37:08  Adv. MTU Over.     : No
Last Mgmt Change     : 02/12/2010 22:37:03  VLAN VC Etype      : 0x8100
Bw BookingFactor     : 100                 PBB Etype           : 0x88e7
Oper Max BW(Kbps)    : 0                    Avail BW(Kbps)      : 0
Net-Domain           : default              Egr Interfaces      : Consistent
Mixed LSP Mode       : Enabled
Revert Time          : 0                    Revert Count Down   : n/a
Flags                : None

KeepAlive Information :
Admin State          : Disabled            Oper State          : Disabled
Hello Time           : 10                  Hello Msg Len       : 0
Hello Timeout        : 5                    Unmatched Replies   : 0
Max Drop Count       : 3                    Hold Down Time      : 10
Tx Hello Msgs        : 0                    Rx Hello Msgs       : 0
-----
LDP Information :
-----
LDP LSP Id          : 65539                LDP Active          : No
-----
RSVP/Static LSPs
-----
Associated LSP LIST :
Lsp Name            : To_7710
Admin State          : Up                  Oper State          : Up
Time Since Last Tran*: 01h20m56s
-----
Class-based forwarding :
-----
Class forwarding     : Disabled            EnforceDSTELspFc   : Disabled
Default LSP          : Uknwn                Multicast LSP       : None
=====
FC Mapping Table
    
```

```

=====
FC Name          LSP Name
-----
No FC Mappings
=====
* indicates that the corresponding row element may have been truncated.
*A:MV-SR12>config>service>vprn#
    
```

Table 19: Output fields: service 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 Signaling	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
Last Status Change	The time of the most recent operating status change to this SDP.
Adv. NTU Over	Specifies whether the advertised MTU of a VLL spoke-SDP bind includes the 14-byte Layer 2 header, so that it is backward compatible with pre-2.0 software.

Label	Description
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 since the keepalive was administratively enabled or the counter was cleared.
Rx Hello Msgs	The number of SDP echo request messages received since the keepalive was administratively enabled or the counter was cleared.
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field.
Lsp Name	Displays the LSP name.
Time Since Last Transaction	Displays the time of the last transaction.
Signaling	Specifies the signaling type.
Collect Stats	Specifies whether the agent collects accounting statistics for this SDP. When the value is true the agent collects accounting statistics on this SDP.
VLAN VC Etype	Displays the VLAN VC type.
BW Booking Factor	Specifies the value used to calculate the max SDP available bandwidth. The value specifies the percentage of the SDP max available bandwidth for VLL call admission. When the value of is

Label	Description
	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.
Revert Time	Specifies the time to wait before reverting back from LDP to the configured LSPs, after having failed over to LDP.
Revert Count Down	Indicates the timer countdown before reverting back from LDP on this SDP. The timer countdown begins after the first configured LSP becomes active.
Flags	Displays all the conditions that affect the operating status of this SDP.
Class Forwarding	Indicates the admin state of class-based forwarding on this SDP. When the value is true, class-based forwarding is enabled.
EnforceDSTELspFc	Specifies whether service manager must validate with RSVP the support of the FC by the LSP.
Default LSP	Specifies the LSP ID that is used as a default when class-based forwarding is enabled on this SDP. This object must be set when enabling class-based forwarding.
Multicast LSP	Displays the LSP ID that all multicast traffic will be forwarded on when class-based forwarding is enabled on this SDP. When this

Label	Description
	object has its default value, multicast traffic will be forwarded on an LSP according to its forwarding class mapping.
Number of SDPs	Displays the metric to be used within the Tunnel Table Manager for decision-making purposes. When multiple SDPs going to the same destination exist, this value is used as a tie-breaker by Tunnel Table Manager users, such as MP-BGP, to select the route with the lower value.

2.21.2.2.2 ETH-CFM show commands

eth-cfm

Syntax

eth-cfm

Context

show

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Description

This command displays Ethernet SFM association information.

Parameters

ma-index

Specifies the maintenance association (MA) index.

Values 1 to 4294967295

detail

Displays detailed information for the eth-cfm association.

Output

The following output is an example of Ethernet CFM association information, and [Table 20: Output fields: association](#) describes the output fields.

Sample output

```
*A:ALU-IPD# show eth-cfm association
=====
CFM Association Table
=====
Md-index   Ma-index   Name                               CCM-interval Bridge-id
-----
1           1          abcabcabcabc1                     1             1
1           2          abcabcabcabc2                     1             2
1           3          abcabcabcabc3                     1             3
1           4          abcabcabcabc4                     1             4
=====
*A:ALU-IPD#
```

Table 20: Output fields: 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 that 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 endpoint (MEP) identifier.

cfm-stack-table

Syntax

cfm-stack-table [{all-ports}] [level 0..7] [direction up | down]

Context

show>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays stack-table information. This stack-table is used to display the various management points MEPs and MIPs configured on the system. This can be Service based. The various options allow users to be specific. If no parameters are included, 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.

Values 0 to 7

direction

Displays the direction in which the MP faces on the bridge port.

Output

The following output is an example of Ethernet CFM stack-table information, and [Table 21: Output fields: CFM stack-table](#) describes the output fields.

Sample output

```
*A:ALU-IPD# show eth-cfm cfm-stack-table
=====
CFM SAP Stack Table
=====
Sap           Level Dir  Md-index  Ma-index  Mep-id  Mac-address
-----
lag-1:1.1     0     Down 2      1         10      00:f3:f0:98:97:1b
lag-1:1.1     6     Down 1      1         1       00:f3:f0:98:97:1b
lag-1:2.2     0     Down 2      2         20      00:f3:f0:98:97:1b
lag-1:2.2     6     Down 1      2         2       00:f3:f0:98:97:1b
=====
CFM Ethernet Tunnel Stack Table
=====
Eth-tunnel   Level Dir  Md-index  Ma-index  Mep-id  Mac-address
```

```

=====
CFM SDP Stack Table
=====
Sdp          Level Dir  Md-index  Ma-index  Mep-id Mac-address
-----
No Matching Entries
=====

CFM Virtual Stack Table
=====
Service      Level Dir  Md-index  Ma-index  Mep-id Mac-address
-----
No Matching Entries
=====
*A:ALU-IPD#
    
```

Table 21: Output fields: CFM stack-table

Label	Description
Sap	Displays associated SAP IDs.
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

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

The following output is an example of Ethernet CFM domain information, and [Table 22: Output fields: domain](#) describes the output fields.

Sample output

```
*A:7210-2# show eth-cfm domain
=====
CFM Domain Table
=====
Md-index   Level Name                                     Format
-----
1           6                                             none
2           0                                             none
=====
*A:7210-2#
```

Table 22: Output fields: domain

Label	Description
Md-index	Displays the Maintenance Domain (MD) index value.
Level	Displays an integer identifying the MD level. Higher numbers correspond to higher MDs, those with the greatest physical reach, with the highest values for customers CFM PDUs. Lower numbers correspond to lower MDs, those with more limited physical reach, with the lowest values for CFM PDUs protecting single bridges or physical links.
Name	Displays a generic MD name.
Format	Displays the type of the 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 mac-address]
```

Context

show>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document

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 Ethernet CFM MEP information, and [Table 23: Output fields: MEP](#) describes the output fields.

- [Sample output](#)
- [Sample output for two-way-slm-test](#)

Sample output

```
*A:7210-2# show eth-cfm mep 1 domain 1 association 1 linktrace
-----
Mep Information
-----
Md-index      : 1                Direction      : Down
Ma-index      : 1                Admin          : Enabled
MepId         : 1                CCM-Enable    : Enabled
IfIndex       : 1342177281      PrimaryVid    : 65537
FngState      : fngReset        ControlMep    : False
LowestDefectPri : macRemErrXcon  HighestDefect  : none
Defect Flags  : None
Mac Address   : 00:f3:f0:98:97:1b CcmLtmPriority : 7
CcmTx         : 531              CcmSequenceErr : 0
Eth-1Dm Threshold : 3(sec)
Eth-Ais      : Disabled
Eth-Tst      : Enabled          Eth-Tst Pattern: : allZerosNoCrc
Eth-Tst dataLength : 64          Eth-Tst Priority: : 7
Eth-Tst Threshold : 1(bitError)
CcmLastFailure Frame:
  None
XconCcmFailure Frame:
  None
-----
Mep Linktrace Message Information
-----
LtRxUnexplained : 0                LtNextSequence : 2
LtStatus        : False           LtResult        : False
TargIsMepId     : False           TargMepId      : 0
TargMac         : 00:00:00:00:00:00 TTL             : 64
EgressId        : 00:00:00:f3:f0:98:97:1b SequenceNum     : 1
LtFlags         : useFDBonly
-----
Mep Linktrace Replies
-----
SequenceNum     : 1                ReceiveOrder    : 1
Ttl             : 63                Forwarded       : False
LastEgressId    : 00:00:00:f3:f0:98:97:1b TerminalMep    : True
NextEgressId    : 00:00:00:e0:b1:99:cb:46 Relay           : n/a
ChassisIdSubType : unknown value (0)
ChassisId:
  None
ManAddressDomain:
  None
ManAddress:
  None
IngressMac      : 00:e0:b1:99:cb:46 Ingress Action  : ing0k
```

```

IngrPortIdSubType : unknown value (0)
IngressPortId:
  None
EgressMac          : 00:00:00:00:00:00          Egress Action    : egrNoTlv
EgrPortIdSubType  : unknown value (0)
EgressPortId:
  None
Org Specific TLV:
  None
-----
*A:7210-2#

*A:7210-2# show eth-cfm mep 1 domain 1 association 1 loopback
-----
Mep Information
-----
Md-index           : 1                          Direction        : Down
Ma-index           : 1                          Admin            : Enabled
MepId              : 1                          CCM-Enable      : Enabled
IfIndex            : 1342177281                 PrimaryVid      : 65537
FngState           : fngReset                   ControlMep      : False
LowestDefectPri    : macRemErrXcon              HighestDefect    : none
Defect Flags       : None
Mac Address        : 00:f3:f0:98:97:1b          CcmLtmPriority   : 7
CcmTx              : 566                        CcmSequenceErr  : 0
Eth-1Dm Threshold : 3(sec)
Eth-Ais            : Disabled
Eth-Tst            : Enabled                    Eth-Tst Pattern: : allZerosNoCrc
Eth-Tst dataLength: 64                          Eth-Tst Priority: : 7
Eth-Tst Threshold : 1(bitError)
CcmLastFailure Frame:
  None
XconCcmFailure Frame:
  None
-----
Mep Loopback Information
-----
LbRxReply          : 1                          LbRxBadOrder    : 0
LbRxBadMsdu        : 0                          LbTxReply       : 0
LbSequence         : 2                          LbNextSequence  : 2
LbStatus           : False                       LbResultOk      : True
DestIsMepId        : False                       DestMepId       : 0
DestMac            : 00:00:00:00:00:00          SendCount       : 0
VlanDropEnable     : True                        VlanPriority     : 7
Data TLV:
  None
-----
*A:7210-2#

*A:7210-2# show eth-cfm mep 1 domain 1 association 1 remote-mepid 10
=====
Eth-CFM Remote-Mep Table
=====
R-mepId Rx CC   Rx Rdi Port-Tlv If-Tlv Peer Mac Addr      CCM status since
-----
10      True  False Absent  Absent 00:e0:b1:99:cb:46 05/20/2010 12:59:55
=====
*A:7210-2#

*A:7210-2# show eth-cfm mep 1 domain 1 association 1 eth-test-results
=====
    
```

Eth CFM ETH-Test Result Table

```
=====
Peer Mac Addr      FrameCount      Current
                   ByteCount      ErrBits
                                                 CrcErrs
                   -----
00:e0:b1:99:cb:46 1                0
                   64              0
                                                 0
=====
```

*A:7210-2# show eth-cfm mep 1 domain 1 association 1 eth-test-results remote-peer 00:e0:b1:99:cb:46

Eth CFM ETH-Test Result Table

```
=====
Peer Mac Addr      FrameCount      Current
                   ByteCount      ErrBits
                   -----
00:e0:b1:99:cb:46 1                0
                   64              0
                                                 0
=====
```

*A:7210-2#

*A:7210-2# show eth-cfm mep 1 domain 1 association 1 one-way-delay-test

Eth CFM One-way Delay Test Result Table

```
=====
Peer Mac Addr      Delay (us)      Delay Variation (us)
                   -----
00:e0:b1:99:cb:46 10000           10000
=====
```

*A:7210-2# show eth-cfm mep 1 domain 1 association 1 one-way-delay-test remote-peer 00:e0:b1:99:cb:46

Eth CFM One-way Delay Test Result Table

```
=====
Peer Mac Addr      Delay (us)      Delay Variation (us)
                   -----
00:e0:b1:99:cb:46 10000           10000
=====
```

*A:7210-2#

*A:7210-2# show eth-cfm mep 1 domain 1 association 1 two-way-delay-test

Eth CFM Two-way Delay Test Result Table

```
=====
Peer Mac Addr      Delay (us)      Delay Variation (us)
                   -----
00:e0:b1:99:cb:46 10000           10000
=====
```

*A:7210-2#

*A:7210-2# # show eth-cfm mep 1 domain 1 association 1 two-way-delay-test remote-peer 00:e0:b1:99:cb:46

Eth CFM Two-way Delay Test Result Table

```

Peer Mac Addr      Delay (us)      Delay Variation (us)
-----
00:e0:b1:99:cb:46  10000          10000
=====
*A:7210-2#
A:dut-b# show eth-cfm mep 1 domain 1 association 1 linktrace
-----
Mep Information
-----
Md-index           : 1                Direction          : Down
Ma-index           : 1                Admin              : Enabled
MepId              : 1                CCM-Enable        : Enabled
IfIndex            : 35946496          PrimaryVid         : 1
FngState           : fngReset          ControlMep         : False
LowestDefectPri    : macRemErrXcon          HighestDefect      : none
Defect Flags       : None
Mac Address        : 00:25:ba:01:c3:6a          CcmLtmPriority     : 7
CcmTx              : 0                CcmSequenceErr    : 0
Eth-1Dm Threshold : 3(sec)
Eth-Ais            : Disabled
Eth-Tst            : Disabled
CcmLastFailure Frame:
  None
XconCcmFailure Frame:
  None
-----
Mep Linktrace Message Information
-----
LtRxUnexplained    : 0                LtNextSequence     : 2
LtStatus           : False           LtResult            : False
TargIsMepId        : False           TargMepId          : 0
TargMac            : 00:00:00:00:00:00          TTL                : 64
EgressId           : 00:00:00:25:ba:01:c3:6a          SequenceNum        : 1
LtFlags            : useFDBOnly
-----
Mep Linktrace Replies
-----
SequenceNum        : 1                ReceiveOrder       : 1
Ttl                : 63                Forwarded          : False
LastEgressId       : 00:00:00:25:ba:01:c3:6a          TerminalMep        : True
NextEgressId       : 00:00:00:25:ba:00:5e:bf          Relay              : rlyHit
ChassisIdSubType   : unknown value (0)
ChassisId:
  None
ManAddressDomain:
  None
ManAddress:
  None
IngressMac         : 00:25:ba:00:5e:bf          Ingress Action     : ingOk
IngrPortIdSubType : unknown value (0)
IngressPortId:
  None
EgressMac          : 00:00:00:00:00:00          Egress Action      : egrNoTlv
EgrPortIdSubType   : unknown value (0)
EgressPortId:
  None
Org Specific TLV:
  None
A:dut-b#
A:dut-b#
A:dut-b# show eth-cfm mep 1 domain 1 association 1 loopback
-----
Mep Information
    
```

```

-----
Md-index          : 1                Direction      : Down
Ma-index          : 1                Admin          : Enabled
MepId             : 1                CCM-Enable    : Enabled
IfIndex           : 35946496         PrimaryVid     : 1
FngState          : fngReset          ControlMep     : False
LowestDefectPri   : macRemErrXcon     HighestDefect  : none
Defect Flags      : None
Mac Address       : 00:25:ba:01:c3:6a  CcmLtmPriority : 7
CcmTx             : 0                CcmSequenceErr : 0
Eth-1Dm Threshold : 3(sec)
Eth-Ais           : Disabled
Eth-Tst           : Disabled
CcmLastFailure Frame:
None
XconCcmFailure Frame:
None
-----
    
```

Mep Loopback Information

```

-----
LbRxReply         : 1                LbRxBadOrder  : 0
LbRxBadMsdu       : 0                LbTxReply     : 0
LbSequence        : 2                LbNextSequence : 2
LbStatus          : False            LbResultOk    : True
DestIsMepId       : False            DestMepId     : 0
DestMac           : 00:00:00:00:00:00 SendCount      : 0
VlanDropEnable    : True              VlanPriority   : 7
Data TLV:
None
A:dut-b#
    
```

*A:dut-b# show eth-cfm mep 1 domain 4 association 4 two-way-delay-test remote-peer 00:25:ba:00:5e:bf

Eth CFM Two-way Delay Test Result Table

Peer Mac Addr	Delay (us)	Delay Variation (us)
00:25:ba:00:5e:bf	507	507

*A:dut-b#

*A:dut-b# show eth-cfm mep 1 domain 4 association 4 two-way-delay-test

Eth CFM Two-way Delay Test Result Table

Peer Mac Addr	Delay (us)	Delay Variation (us)
00:25:ba:00:5e:bf	507	507

*A:dut-b#

*A:dut-a# show eth-cfm mep 2 domain 4 association 4 eth-test-results remote-peer 00:25:ba:01:c3:6a

Eth CFM ETH-Test Result Table

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#
*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
                   ByteCount      CrcErrs         CrcErrs
-----
00:25:ba:01:c3:6a 6                0                0
                   384             0                0
=====
*A:dut-a# show eth-cfm mep 2 domain 4 association 4 one-way-delay-test remote-
peer 00:25:ba:01:c3:6a

=====
Eth CFM One-way Delay Test Result Table
=====
Peer Mac Addr      Delay (us)      Delay Variation (us)
-----
00:25:ba:01:c3:6a 402             402
=====
*A:dut-a#

*A:dut-a# show eth-cfm mep 2 domain 4 association 4 one-way-delay-test

=====
Eth CFM One-way Delay Test Result Table
=====
Peer Mac Addr      Delay (us)      Delay Variation (us)
-----
00:25:ba:01:c3:6a 402             402
=====
*A:dut-a#
    
```

Table 23: Output fields: MEP

Label	Description
Mep Information	
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.
MepId	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.

Label	Description
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 (LTMs) 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).
CcmSequenceErr	Displays the number of CCM errors.
Eth-1DM Threshold	Displays the one-way-delay threshold value.
Eth-Ais	Displays the state of the ETH-AIS test (enabled or disabled).
Eth-Tst	Displays the state of the ETH-Test (enabled or disabled).
CcmLastFailure Frame	Displays the frame that caused the last CCM failure.
XconCcmFailure Frame	Displays the frame that caused the XconCCMFailure.
Mep Linktrace Message Information	
LtRxUnexplained	Displays the number of unexplained linktrace messages (LTM) that have been received.
LtNextSequence	Displays the sequence number of the next linktrace message.
LtStatus	Displays the status of the linktrace.
LtResult	Displays the result of the linktrace.
TargIsMepId	Identifies whether the target interface has a MEP-ID (true or false).
TargMepId	Displays the MEP-ID of the target interface.
TargMac	Displays the MAC address of the target interface.

Label	Description
TTL	Displays the TTL value.
EgressId	Displays the egress ID of the linktrace message.
SequenceNum	Displays the sequence number of the linktrace message.
LtFlags	Displays the linktrace flags.
Mep Linktrace Replies	
SequenceNum	Displays the sequence number returned by a previous transmit linktrace message, indicating which linktrace message response will be returned.
ReceiveOrder	Displays the order in which the linktrace initiator received the linktrace replies.
Ttl	Displays the TTL field value for a returned linktrace reply.
Forwarded	Indicates whether the linktrace message was forwarded by the responding MEP.
LastEgressId	<p>Displays the last egress identifier returned in the linktrace reply egress identifier TLV of the linktrace reply.</p> <p>The last egress identifier identifies the MEP linktrace initiator that initiated, or the linktrace responder that forwarded, the linktrace message for which this linktrace reply is the response.</p> <p>This is the same value as the egress identifier TLV of that linktrace message.</p>
TerminalMep	Indicates whether the forwarded linktrace message reached a MEP enclosing its MA.
NextEgressId	<p>Displays the next egress identifier returned in the linktrace reply egress identifier TLV of the linktrace reply. The next egress identifier identifies the linktrace responder that transmitted this linktrace reply and can forward the linktrace message to the next hop. This is the same value as the egress identifier TLV of the forwarded linktrace message, if any.</p>
Relay	Displays the value returned in the Relay Action field.
IngressMac	Displays the MAC address returned in the ingress MAC address field.
Ingress Action	Displays the value returned in the Ingress Action field of the linktrace message.
IngressPortIdSubType	Displays the format of the ingress port ID.

Label	Description
IngressPortId	Displays the ingress port ID; the format is determined by the value of the IngressPortIdSubType.
EgressMac	Displays the MAC address returned in the egress MAC address field.
Egress Action	Displays the value returned in the Egress Action field of the linktrace message.
EgressPortIdSubType	Displays the format of the egress port ID.
EgressPortId	Displays the egress port ID; the format is determined by the value of the EgressPortIDSubType.
Org Specific TLV	Displays all organization-specific TLVs returned in the linktrace reply, if any. Includes all octets including and following the TLV length field of each TLV, concatenated.
Mep Loopback Information	
LbRxReply	Displays the number of received loopback (LB) replies.
LbRxBadOrder	Displays the number of received loopback messages that are in a bad order.
LbRxBadMsdu	Displays the number of loopback replies that have been received with the wrong destination MAC address (MSDU = MAC Service Data Unit).
LbTxReply	Displays the number of loopback replies transmitted out this MEP.
LbSequence	Displays the sequence number in the loopback message.
LbNextSequence	Displays the next loopback sequence.
LbStatus	Displays the loopback status as True or False: True — loopback is in progress False — no loopback is in progress
LbResultOk	Displays the result of the loopback test.
DestIsMepId	Identifies whether the destination interface has a MEP-ID (true or false).
DestMepId	Displays the MEP-ID of the destination interface.
DestMac	Displays the MAC address of the destination interface.

Sample 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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command displays connection profile information.

Parameters

conn-prof-id

Specifies the connection profile ID.

Values 1 to 1000

associations

Displays the SAP and the service ID that use this connection profile.

Output

The following outputs are examples of connection profile information, and [Table 24: Output fields: show connection profile](#) describes the output fields.

- [Sample output — standard](#)
- [Sample output — connection-profile associations, Table 24: Output fields: show connection profile](#)

Sample output — standard

```
*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#
    
```

Sample output — 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: show 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 SAS.
HasRange	Indicates whether VLAN range is configured.

3 VLL services

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

3.1 Epipe

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

3.2 Epipe service overview

An Epipe service is a Layer 2 point-to-point service where the customer data is encapsulated and transported across a service provider's network. An Epipe service is completely transparent to the subscriber's data and protocols. The 7210 SAS Epipe service does not perform any MAC learning. A local Epipe service consists of two SAPs on the same node, whereas a distributed Epipe service consists of two SAPs on different nodes.

The following support is available on different platforms:

- The 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T support only local Epipe services. The 7210 SAS-D provide null, QinQ SAPs and dot1q SAPs to provide point-point Layer 2 local services.
- The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C support both local and distributed Epipe services. The platforms provide null, QinQ SAPs and dot1q SAPs to provide a point-point Layer 2 local and distributed service. The platforms support the use of MPLS tunnels for distributed services.

Each SAP configuration includes a specific port on which service traffic enters the 7210 SAS router from the customer side (also called the access side). Each port is configured with an encapsulation type. If a port is configured with an IEEE 802.1Q (referred to as dot1q) encapsulation, a unique encapsulation value (ID) must be specified.

[Figure 17: Local Epipe/VLL service using access or access-uplink ports on the 7210 SAS](#) shows a local Epipe (VLL) service using access or access-uplink ports on a 7210 SAS. [Figure 18: Distributed Epipe/VLL service using network ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#) shows a distributed Epipe (VLL) service using network ports on a 7210 SAS-K 2F6C4T or 7210 SAS-K 3SFP+ 8C.

Figure 17: Local Epipe/VLL service using access or access-uplink ports on the 7210 SAS





Figure 18: Distributed Epipe/ VLL service using network ports on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C



3.3 Epipe oper state decoupling on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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 hand-off. The operator can also maintain performance and continuity measurement across their network regardless of the connectivity between the terminating node and the customer.

If the SAP between the operator and the customer enters a Oper Down state, the Epipe remains operationally up, so the results can continue to be collected uninterrupted. The operator receives applicable port or SAP alerts/alarms. This option is available only for the customer-facing SAP failures. If a network-facing SAP or spoke-SDP fails, the operational state of the Epipe service is set to down. 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.
- As the service remains operationally up, all bindings should remain operationally up and should be able to receive and transmit data. The PW status represents the failed SAP in the LDP status message, but this does not prevent the data from using the PW as a transport, in or out. This is the same as LDP status messaging.
- The SAP failure continues to trigger normal reactions, except the operational state of the service.
- ETH-CFM PM measurement tools (DMM/SLM) can be used with the Up MEP on the failed SAP to collect performance metric. Additionally, CFM troubleshooting tools and connectivity (LBM, LTM, AIS, CCM) can be used and will function regularly.
- 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.

3.4 Dynamic Multi-Segment Pseudowire routing



Note:

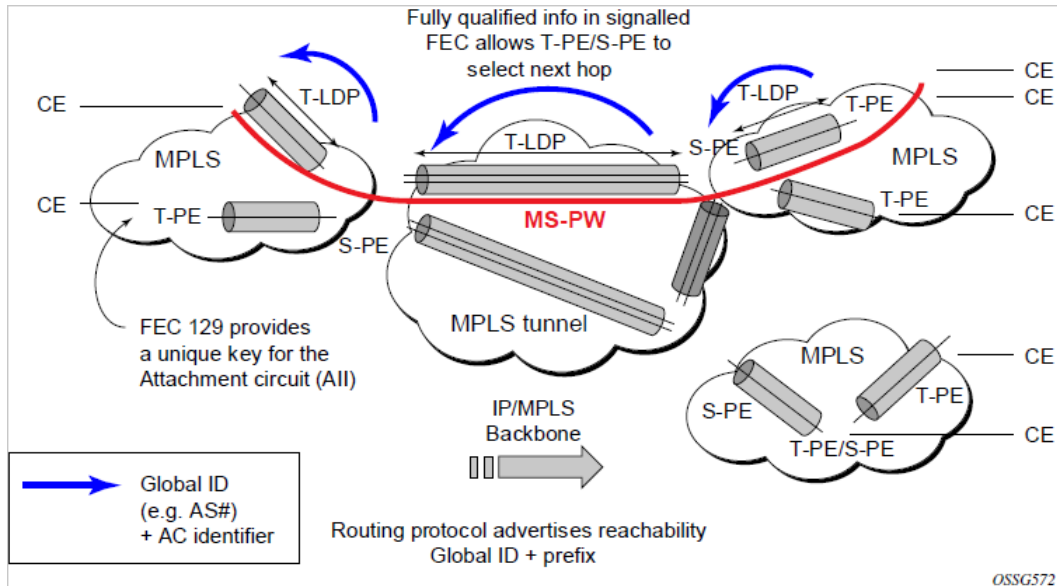
The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C can only be configured as a T-PE node and not as an S-PE node.

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

Dynamic Multi-Segment Pseudowire (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 19: Dynamic MS-PW operation



The FEC 129 All Type 2 structure depicted in the following figure is used to identify each individual pseudowire endpoint.

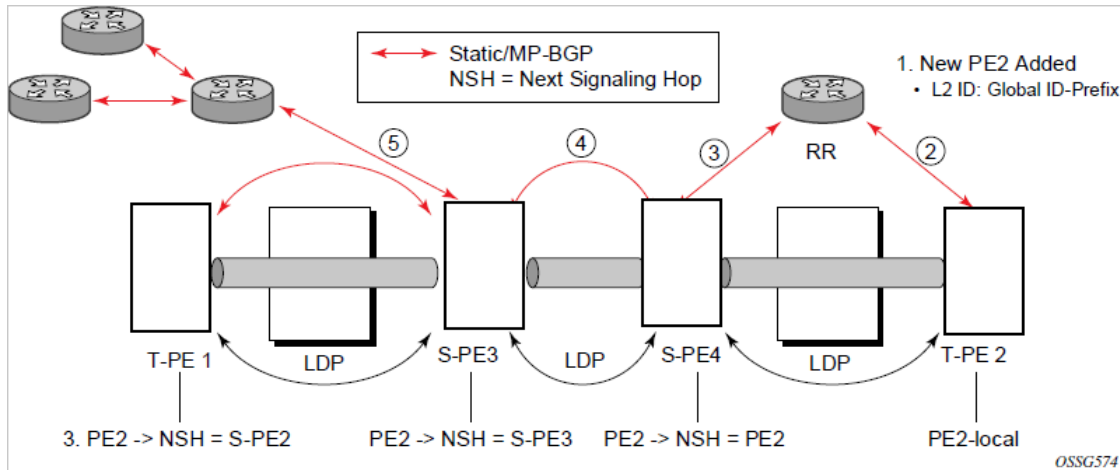
Figure 20: MS-PW addressing using FEC129 All Type 2



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

This new All type is mapped into the MS-PW BGP NLRI (a new BGP AFI of Layer 2 VPN, 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 21: Advertisement of PE addresses by PW routing



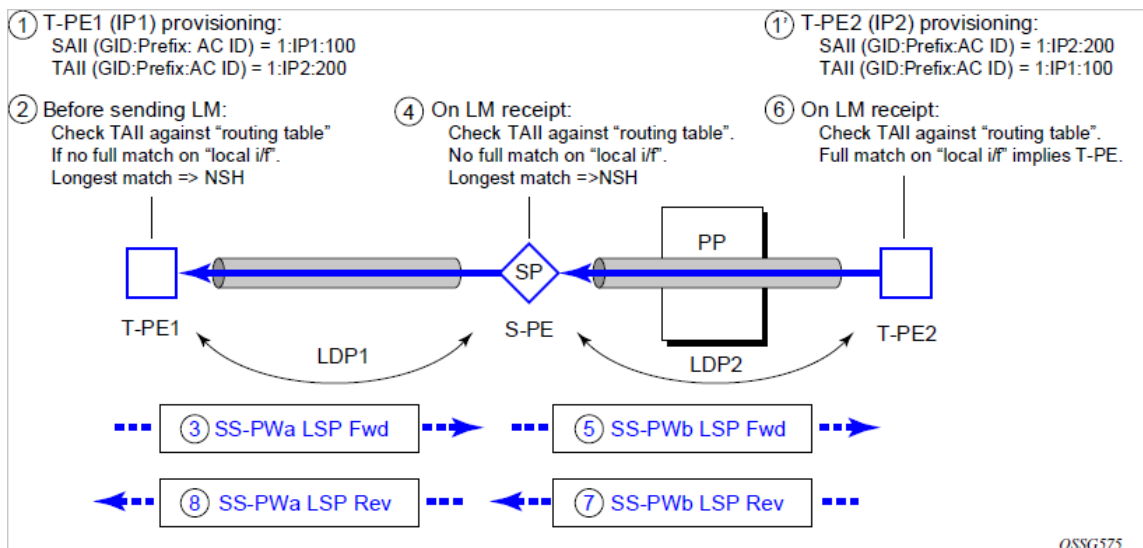
In step 1, a new T-PE (T-PE2) is configured with a local prefix.

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.

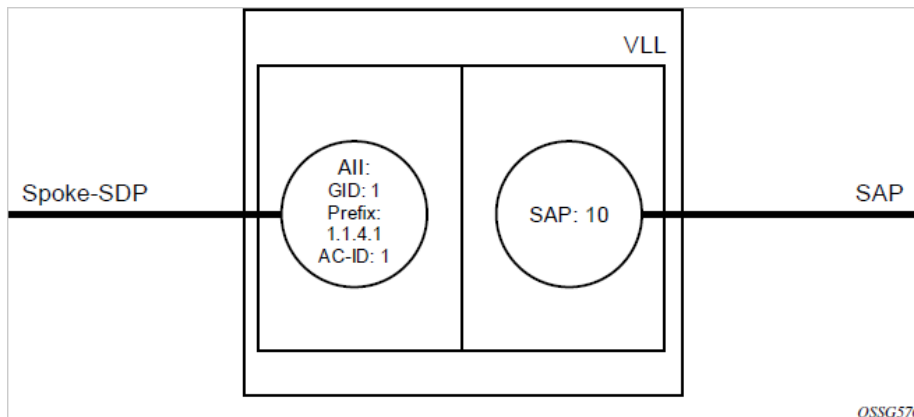
VLL services can then be established, as shown in the following figure.

Figure 22: 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 All (SAII), Target All (TAII). On the 7210 SAS, the AIIs are locally configured for each spoke-SDP, according to the model shown in Figure 23: Mapping of All to SAP. Therefore, the 7210 SAS provides for a flexible mapping of All to SAP. That is, the values used for the All are through local configuration, and it is the context of the spoke-SDP that binds it to a specific SAP.

Figure 23: 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.

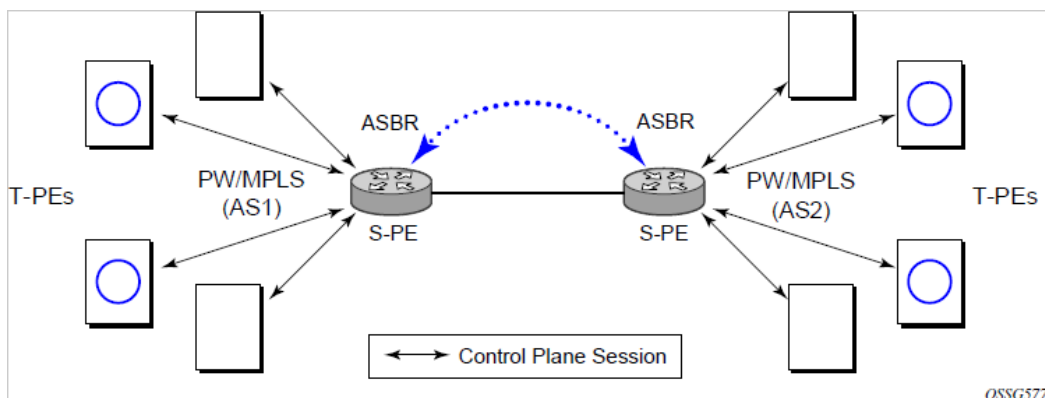
**Note:**

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 24: 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.

3.4.1 MS-PW routing — pseudowire routing

**Note:**

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C can only be configured as a T-PE node and not as an S-PE node.

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

Each S-PE and T-PE has a pseudowire routing table that contains a reference to the T-LDP session to use to signal to a set of next-hop S-PEs to reach a 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 8-byte RD, a 4-byte global ID, a 4-byte local prefix, and (optionally) a 4-byte AC ID. Support for the MS-PW address family is configured in CLI under **config>router>bgp>family ms-pw**.

MS-PW routing parameters are configured in the **config>service>pw-routing** context.

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-SDPs using dynamic MS-PWs at a T-PE unless an S-PE address has already been configured. The S-PE address is used as the address of a node used to populate the switching point TLV in the LDP label mapping message and the pseudowire status notification sent for faults at an S-PE.

Each T-PE is also be configured with the following parameters:

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 (RD) 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 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.

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. Nokia recommends enabling 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 parameter for iBGP and BGP communities can be configured under such a policy.

3.4.1.1 MS-PW routing — static routing

In addition to support for BGP routing, static MS-PW routes may also be configured using the **config>services>pw-routing>static-route** command. Each static route comprises the target T-PE global ID and prefix, and the IP address of the T-LDP session to the next-hop S-PE or T-PE that should be used.

If a static route is set to 0, 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.

3.4.1.2 MS-PW routing — explicit paths

A set of default explicit routes to a remote T-PE or S-PE prefix may be configured on a T-PE under **config>services>pw-routing** using the **path name** command. Explicit paths are used to populate the explicit route TLV used by MS-PW T-LDP signaling. Only strict (fully qualified) explicit paths are supported.



Note:

It is possible to configure explicit paths independently of the configuration of BGP or static routing.

3.4.2 MS-PW routing — configuring VLLs using dynamic MS-PWs

One or more spoke-SDPs may be configured for distributed Epipe VLL services. Dynamic MS-PWs use FEC129 (also known as the Generalized ID FEC) with Attachment Individual Identifier (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 Source Attachment Individual Identifier (SAII) and a Target Attachment Individual Identifier (TAII) to identify the end of a pseudowire at the T-PE. The SAII identifies the local end, while the TAII identifies the remote end. The SAII and TAII are each structured as follows:

- global ID - a 4-byte identifier that uniquely identifies an operator or the local network
- prefix - a 4-byte prefix, which should correspond to one of the local prefixes assigned under pw-routing
- AC ID - a 4-byte identifier for this end of the pseudowire. This should be locally unique within the scope of the global-id:prefix

3.4.2.1 MS-PW routing — active/passive T-PE selection

Dynamic MS-PWs use single-sided signaling procedures with double-sided configuration; a fully qualified FEC must be configured at both endpoints. That is, one T-PE (the source T-PE: ST-PE) of the MS-PW initiates signaling for the MS-PW, while the other end (the terminating T-PE: TT-PE) passively waits for the label mapping message from the far-end. 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 SAII with the TAII as unsigned integers. The T-PE with SAII>TAII assumes the active role. However, it is possible to override this behavior using the signaling **{master | auto}** command under the spoke-sdp-fec. If master is selected at a 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.

3.4.2.2 MS-PW routing — automatic endpoint configuration

Automatic endpoint configuration allows the configuration of an endpoint without specifying the TAII associated with that spoke-sdp-fec. It allows a single-sided provisioning model where an incoming label mapping message with a TAII that matches the SAII of that spoke-SDP can be automatically bound to that endpoint. This is useful in scenarios where a service provider needs to separate service configuration from the service activation phase.

Automatic endpoint configuration is supported required for Epipe VLL spoke-sdp-fec endpoints bound to a VLL SAP. It is configured using the **spoke-sdp-fec>auto-config** command, and excludes the TAIL from the configuration. When autoconfiguration is used, the node assumes passive behavior from a point of view of T-LDP signaling. Therefore, the far-end T-PE must be configured for signaling master for that spoke-sdp-fec.

3.4.2.3 MS-PW routing — selecting a path for an MS-PW

Path selection for signaling occurs in the outbound direction (ST-PE to TT-PE) for an MS-PW. In the TT-PE to ST-PE direction, a label mapping message 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* under 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.

3.4.2.4 MS-PW routing — pseudowire templates

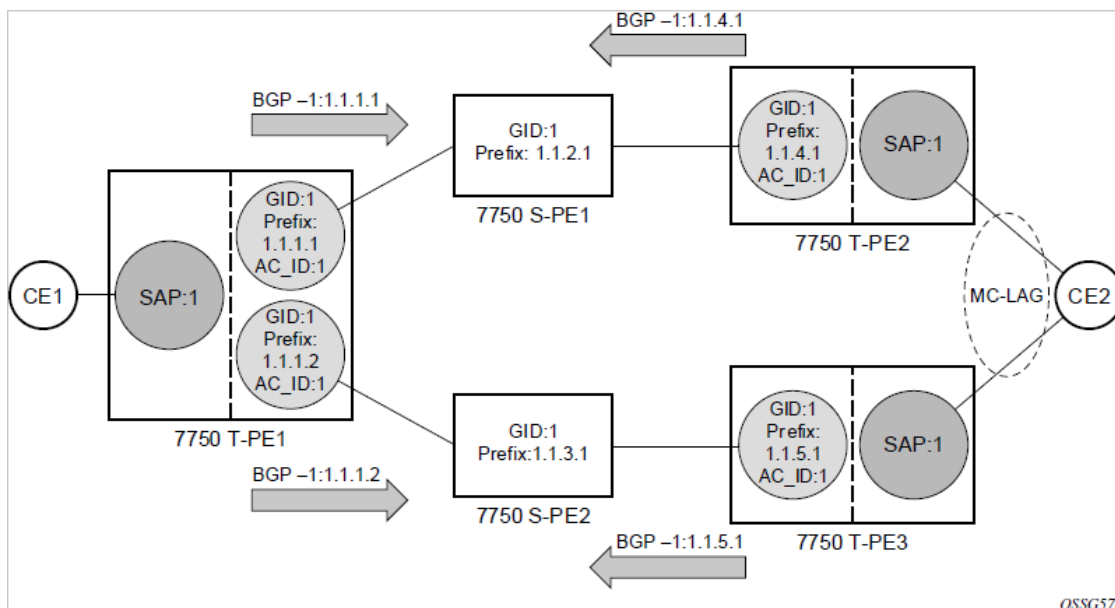
Dynamic MS-PWs support the use of the pseudowire template for specifying generic pseudowire parameters at the T-PE. The pseudowire template to use is configured in the **spoke-sdp-fec>pw-template-bind** *policy-id* context. Dynamic MS-PWs do not support the provisioned SDPs specified in the pseudowire template.

3.4.3 MS-PW routing — pseudowire redundancy

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

The following figure shows the use of pseudowire redundancy.

Figure 25: 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/TAIL 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 RD.

If the primary MS-PW fails, fail-over to a standby MS-PW, as per the normal pseudowire redundancy procedures. A configurable retry timer for the failed primary MS-PW is then started. When the timer expires, attempt to 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:

Since the SDP ID is determined dynamically at signaling time, it cannot be used as a tie breaker to choose the primary MS-PW between multiple MS-PWs of the same precedence. The user should therefore explicitly configure the precedence values to determine which MS-PW is active in the final selection.

3.4.4 MS-PW routing — 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.

3.4.5 MS-PW routing — 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*. 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.

If the SAIL:TAIL 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.

3.4.6 MS-PW routing — VCCV-trace on dynamic MS-PWs

The 7210 SAS supports the MS-PW path trace mode of operation for VCCV-trace, as per pseudowire switching, but using FEC129 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, because it is not configured for a specific MS-PW at S-PE nodes.

3.4.7 MS-PW routing — example dynamic MS-PW configuration

This section presents an example of how to configure dynamic MS-PWs for a VLL service between a set of 7210 SAS nodes. The network consists of two 7210 SAS T-PEs and two 7210 playing the role of S-PEs, as shown in the following figure. Each 7210 peers with its neighbor using LDP and BGP.

Figure 26: Dynamic MS-PW example



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

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.

```

config
router
  ldp
    targeted-session
      peer 10.20.1.5
    exit
  exit
  policy-options
    begin
    policy-statement "exportMsPw"
      entry 10
        from
          family ms-pw
        exit
        action accept
      exit
    exit
  exit
  commit
exit
bgp
  family ms-pw
  connect-retry 1
  min-route-advertisement 1
  export "exportMsPw"
  rapid-withdrawal
  group "ebgp"
  neighbor 10.20.1.5
    
```



```
        multihop 255
        peer-as 200
    exit
exit
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
  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
```

```
config
router
  ldp
  targeted-session
  peer 10.20.1.2
  exit
  exit
  ...
  policy-options
  begin
  policy-statement "exportMsPw"
  entry 10
  from
  family ms-pw
  exit
  action accept
  exit
  exit
  exit
  commit
exit

bgp
  family ms-pw
  connect-retry 1
  min-route-advertisement 1
```

```
export "exportMsPw"
  rapid-withdrawal
  group "ebgp"
    neighbor 10.20.1.2
      multihop 255
      peer-as 300
    exit
  exit
exit
config
  service
    pw-routing
      spe-address 6:10.20.1.6
      local-prefix 6:10.20.1.6 create
    exit
    path "path1_to_F" create
      hop 1 10.20.1.2
      hop 2 10.20.1.5
      no shutdown
    exit
  exit
  epipe 1 customer 1 vpn 1 create
    description "Default epipe
      description for service id 1"
  service-mtu 1400
    service-name "XYZ Epipe 1"
    sap 1/1/3:1 create
  exit
  spoke-sdp-fec 1 fec 129 aii-type 2 create
    retry-timer 10
    retry-count 10
    saii-type2 6:10.20.1.6:1
    taii-type2 3:10.20.1.3:1
    no shutdown
  exit
  no shutdown
exit
```

3.5 Master-slave operation

**Note:**

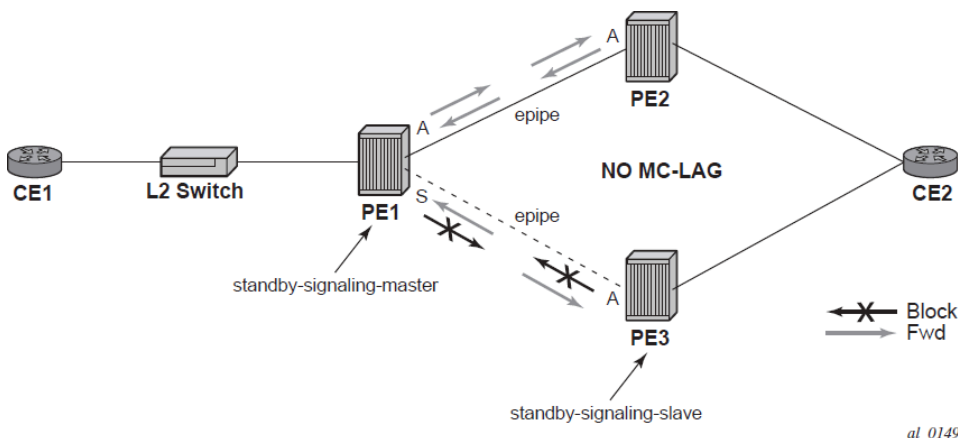
7210 SAS devices support only the **standby-signaling-master** option. The 7210 SAS does not support the CLI command **standby-signaling-slave**. In the following description, reference to the **standby-signaling-slave** command is only used to describe the solution. The 7210 SAS can be used only where **standby-signaling-master** is used in the following 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 signaled by the master endpoint. This satisfies a majority of deployments where bidirectional blocking of the forwarding on a standby spoke-SDP is required.

Figure 27: Master-slave pseudowire redundancy 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 27: 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. However, if standby-signaling-slave is enabled at the remote VLL endpoint then the Tx direction of the spoke-SDP will also be blocked, according to the rules in [PW redundancy — operation of master-slave pseudowire redundancy with existing scenarios](#).

**Note:**

Although master-slave operation provides bidirectional blocking of a standby spoke-SDP during steady-state conditions, it is possible that the Tx directions of more than one slave endpoint can be active for transient periods during a fail-over operation. This is due to slave endpoints transitioning a spoke-SDP from standby to active receiving or processing a pseudowire preferential forwarding status message before those transitioning a spoke-SDP to standby.

This transient condition is most likely when a forced switch-over is performed, or the relative preferences of the spoke-SDPs 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 due to common network faults that can occur during normal operation, a failure of connectivity on the path of the spoke-SDP or the SAP, would not result in such loops in the datapath.

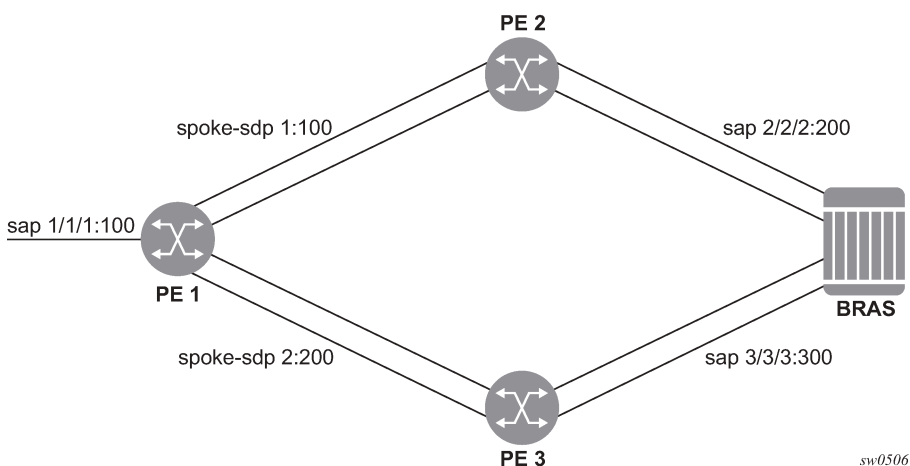
3.5.1 PW redundancy — operation of master-slave pseudowire redundancy with existing scenarios

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

3.5.1.1 PW redundancy — VLL resilience

The following figure shows a VLL resilience path example. An sample configuration follows.

Figure 28: VLL resilience path

**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.

```
PE1
configure service epipe 1
endpoint X
exit
endpoint Y
revert-time 0
standby-signaling-master
```

```
    exit
    sap 1/1/1:100 endpoint X
    spoke-sdp 1:100 endpoint Y
precedence primary
    spoke-sdp 2:200 endpoint Y
precedence 1
PE2
configure service epipe 1
    endpoint X
    exit
    sap 2/2/2:200 endpoint X
    spoke-sdp 1:100
        standby-signaling-slave
```

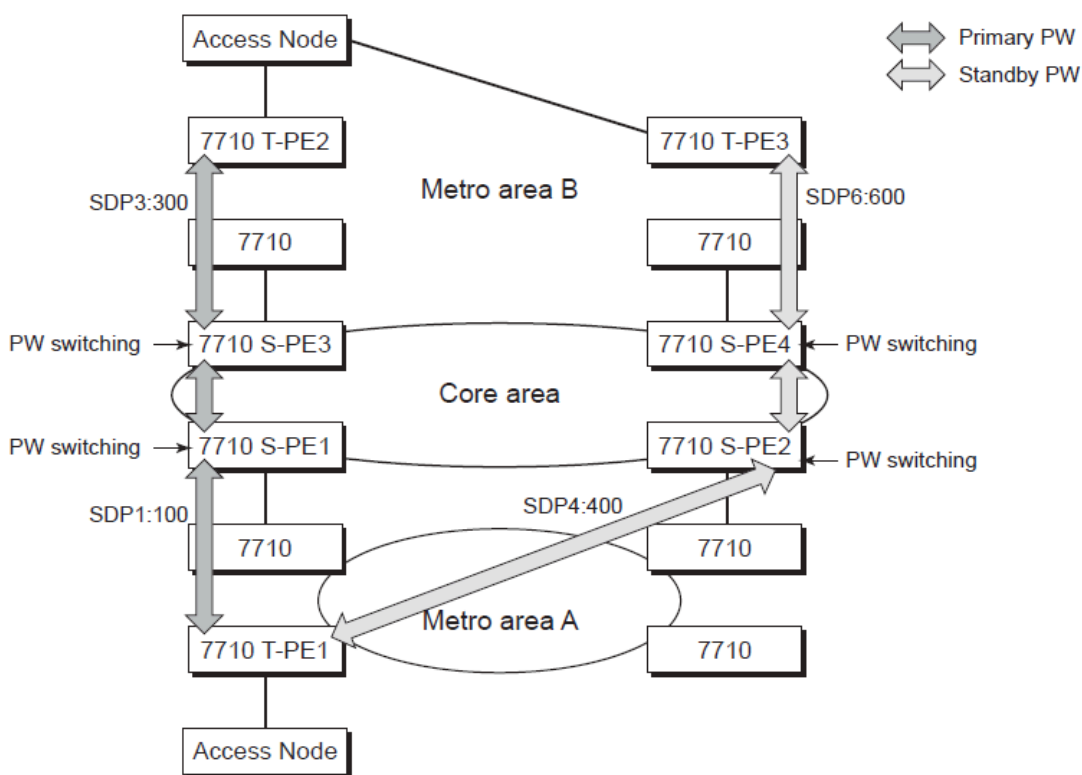
PE3

```
configure service epipe 1
    endpoint X
    exit
    sap 3/3/3:300 endpoint X
    spoke-sdp 2:200
        standby-signaling-slave
```

3.5.2 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 29: VLL resilience with PW redundancy and PW switching



OSSG114

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 with two destination PE nodes](#), 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 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 due to the differential delay between the paths. This will cause T-PE1 to switch the path of the VLL to the secondary standby pseudowire and remain there until the status notification is cleared. At that time, the VLL path is switched back to the primary pseudowire due to the revertive behavior operation. The path will not switch back to a secondary path when it comes up, even if it has a higher precedence than the currently active secondary path.

3.5.2.1 PW redundancy — VLL resilience for a switched PW path




Note:

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C can only be configured as a T-PE node and not as an S-PE node.

VC-switching is not supported on the 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

The following figure shows a VLL resilience for a switched pseudowire path example. A sample configuration follows.

Figure 30: VLL resilience with pseudowire switching



Configuration

```
T-PE1
  configure service epipe 1
  endpoint X
  exit
  endpoint Y
  revert-time 100
  standby-signaling-master
  exit
  sap 1/1/1:100 endpoint X
  spoke-sdp 1:100 endpoint Y
    precedence primary
  spoke-sdp 2:200 endpoint Y
    precedence 1
  spoke-sdp 3:300 endpoint Y
    precedence 1

T-PE2
  configure service epipe 1
  endpoint X
  exit
  endpoint Y
  revert-time 100
  standby-signaling-slave
  exit
```

```
sap 2/2/2:200 endpoint X
spoke-sdp 4:400 endpoint Y
  precedence primary
spoke-sdp 5:500 endpoint Y
  precedence 1
spoke-sdp 6:600 endpoint Y
  precedence 1
```

S-PE1

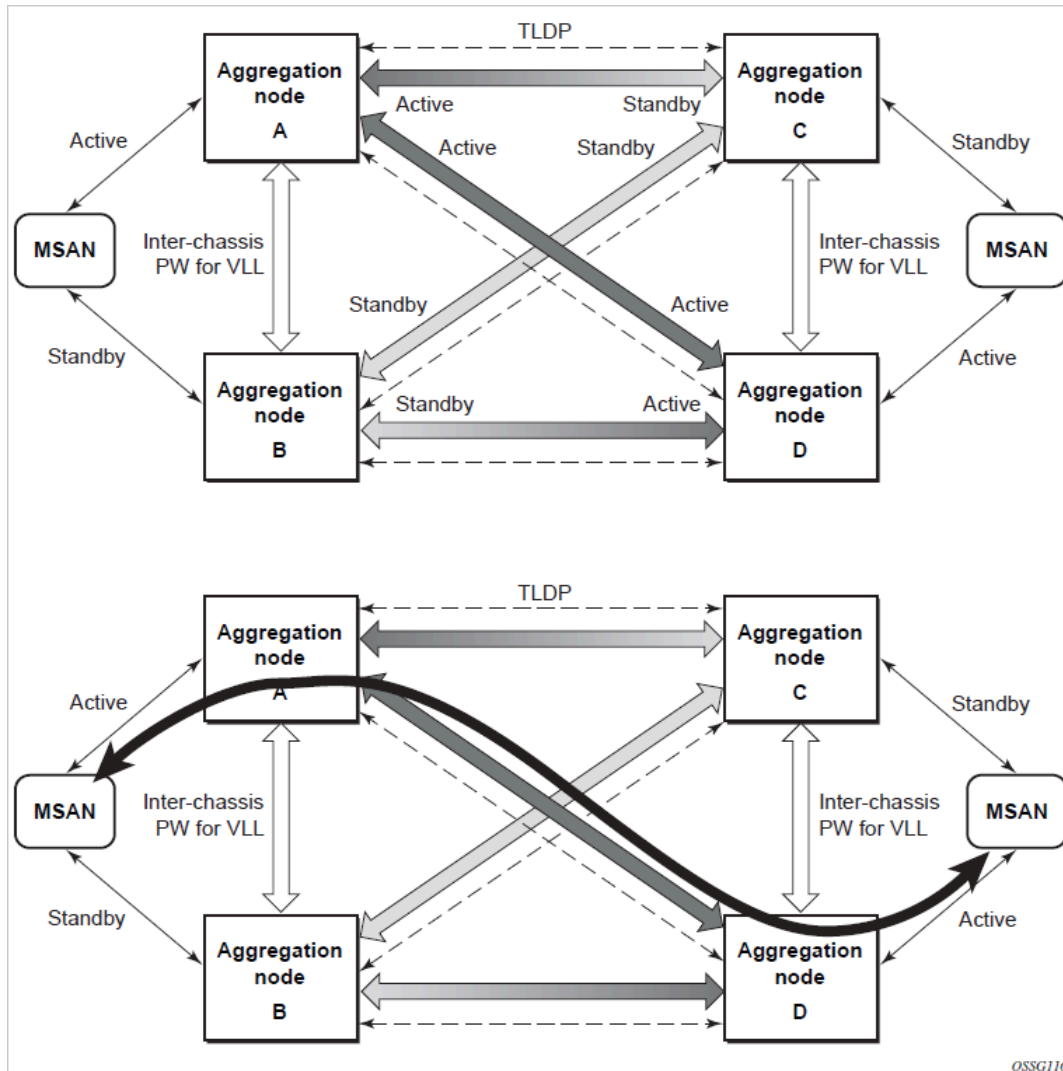
VC switching indicates a VC cross-connect so that the service manager does not signal the VC label mapping immediately, but will put this into passive mode.

```
configure service epipe 1 vc-switching
  spoke-sdp 1:100
  spoke-sdp 4:400
```

3.5.3 Access node 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 31: 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:

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 an MC-LAG instance.

3.5.4 Epipe pseudowire switching on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C



Note:

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C can only be configured as T-PE nodes and not as S-PE nodes. In this section, references to S-PE nodes refer to other products in the Nokia IP product family that support S-PE functionality.

The pseudowire switching feature provides the user with the ability to create a VLL service by cross-connecting two spoke-SDPs. This feature allows the scaling of VLL and VPLS services in a large network in which the otherwise full mesh of PE devices would require thousands of Targeted LDP (T-LDP) sessions per PE node.

Services with one SAP and one spoke-SDP are created on the PE; however, the target destination of the SDP is the pseudowire switching node instead of the remote PE. In addition, the user configures a VLL service on the pseudowire switching node using the two SDPs.

The pseudowire switching node (that is, S-PE node) acts in a passive role with respect to signaling of the pseudowires. It waits until one or both PEs (that is, T-PEs) sends the label mapping message before relaying it to the other PE. This is because it needs to pass the interface parameters of each PE to the other.

A 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 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 example, the user configures a regular Epipe VLL service PE1 and PE2 (acting as T-PE nodes). 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 SAS-K 2F6C4T PE1 [T-PE node] (Epipe)|---sdp 2:10---  
|7450 ESS, 7750 SR, and 7950 XRS PW SW [S-PE node] (Epipe)|---sdp 7:15---|7210 SAS-  
K12 PE2 [T-PE node] (Epipe)|
```

Configuration examples can be found in "Configuring Two VLL Paths Terminating on T-PE2".

3.5.4.1 Pseudowire switching with protection

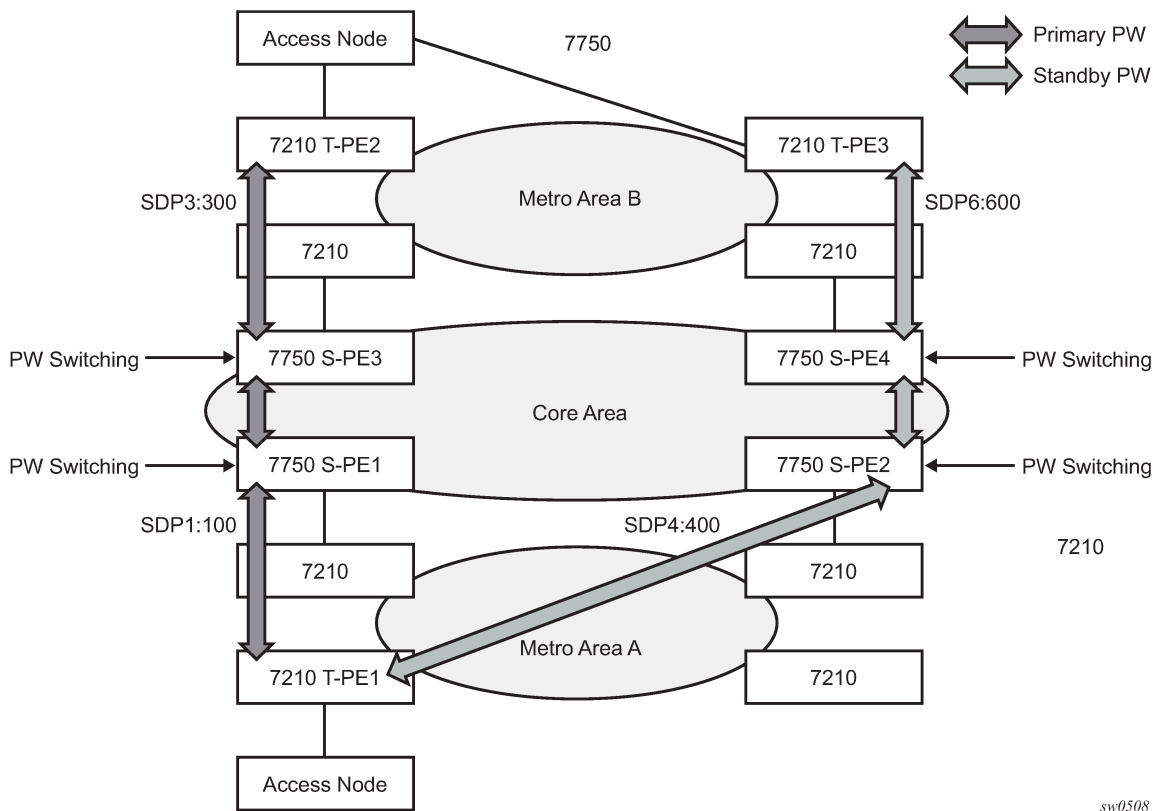
Pseudowire switching scales VLL and VPLS services over a multi-area network by removing the need for a full mesh of targeted LDP sessions between PE nodes. This illustrates the use of pseudowire redundancy to provide a scalable and resilient VLL service across multiple IGP areas in a provider network.

In the network in [Figure 32: 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 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 SPE2 will affect the spoke-SDP cross-connect only when both directions of the pseudowire have been signaled and matched.

Figure 32: VLL resilience with pseudowire redundancy and switching



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 and/or the S-PE nodes. Pseudowire status messages received by a switching node are processed, 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 TPE node and the S-PE should process and pass the message without changes except for the VCID value in the FEC TLV.

3.5.4.2 Pseudowire switching behavior

In the network in [Figure 32: 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 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. TPE1 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.

Pseudowire status notification messages can be generated by the T-PE nodes and/ or the S-PE nodes. Pseudowire status notification messages received by a switching node are processed, then passed on to the next hop. An S-PE node appends the optional pseudowire switching TLV, with its system address added to it, to the FEC in the pseudowire status notification message only if it originated the message or the message was received with the TLV in it. Otherwise, it means the message was originated by a T-PE node and the S-PE should process and pass the message without changes except for the VC ID value in the FEC TLV.

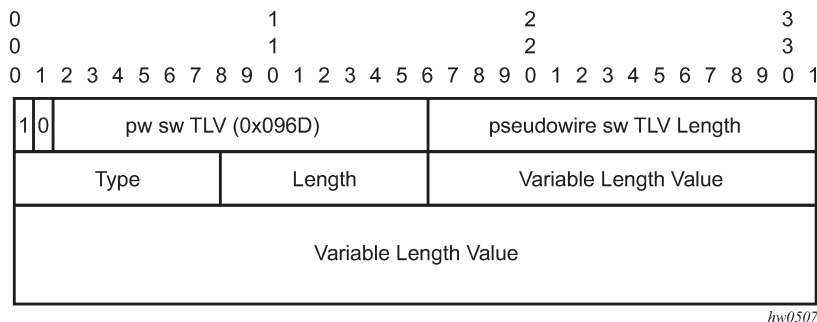
The merging of the received T-LDP status notification message and the local status for the spoke-SDPs from the service manager at a PE complies with the following rules:

- When the local status for both spoke-SDPs is up, the S-PE passes any received SAP or SDP-binding generated status notification message unchanged; for example, the status notification TLV is unchanged but the VC ID in the FEC TLV is set to the value of the pseudowire segment to the next hop.
- When the local operational status for any of the spoke-SDPs is down, the S-PE always sends SDP-binding down status bits, regardless of whether the received status bits from the remote node indicated SAP up/down or SDP-binding up/down.

3.5.4.2.1 Pseudowire switching TLV

The following figure shows the format of the pseudowire switching TLV.

Figure 33: Pseudowire switching TLV



PW sw TLV Length - Specifies the total length of all the following pseudowire switching point TLV fields in octets:

- Type - encodes how the value field is to be interpreted
- Length - specifies the length of the value field in octets
- Value - octet string of length octets that encodes information to be interpreted as specified by the type field

3.5.4.2.1.1 Pseudowire switching point sub-TLVs

The following are details specific to pseudowire switching point sub-TLVs:

- pseudowire ID of last pseudowire segment traversed - this sub-TLV type contains a pseudowire ID in the format of the pseudowire ID
- pseudowire switching point description string - an optional description string of text up to 80 characters
- IP address of pseudowire switching point

The IP V4 or V6 address of the pseudowire switching point. This is an optional sub TLV:

- MH VCCV capability indication

3.5.4.3 Pseudowire redundancy

Pseudowire redundancy provides the ability to protect a pseudowire with a preprovisioned secondary standby pseudowire and to switch traffic over to the secondary standby pseudowire in case of a SAP and/or network failure condition. Pseudowires are redundant due to the SDP redundancy mechanism. For example, if the SDP is an RSVP LSP and is protected by a secondary standby path and/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 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.


3.5.5 VLL service resilience with pseudowire on 7210 SAS-K 3SFP+ 8C

The following sections describe VLL resilience.

3.5.5.1 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 34: VLL resilience



OSSG115

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. However, in other applications such as those described in [Access node resilience using MC-LAG and pseudowire redundancy](#), it will be important to minimize service outage to end users.

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 back to the primary pseudowire.

The behavior of the pseudowire redundancy feature is the same if PE1 detects or is notified of a network failure that brought the spoke-SDP operational status to down. The following are the events that will cause PE1 to trigger a switchover to the secondary standby pseudowire:

- LDP peer (remote PE) node withdrew the pseudowire label.
- T-LDP peer signaled a FEC status indicating a pseudowire failure or a remote SAP failure.
- T-LDP session to the peer node times out.
- 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 SDP type for the primary and secondary pseudowires need not be the same. That is, the user can protect a RSVP-TE based spoke-SDP with a LDP one. This provides the ability to route the path of the two pseudowires over different areas of the network.

The 7210 SAS routers support the ability to configure multiple secondary standby pseudowire paths. For example, PE1 uses the value of the 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.

On the 7210 SAS, this application is supported with Epipe Service.


3.5.5.1.1 Master-slave operation

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. It adds the ability for the remote peer to react to the pseudowire standby status notification, even if only one spoke-SDP terminates on the VLL endpoint on the remote peer, by blocking the transmit (Tx) direction of a VLL spoke-SDP when the far-end PE signals standby. This solution enables the blocking of the Tx direction of a VLL spoke-SDP at both master and slave endpoints when standby is signaled by the master endpoint. This approach satisfies a majority of deployments where bidirectional blocking of the forwarding on a standby spoke-SDP is required.

Figure 35: Master-slave pseudowire redundancy 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 objectives of this feature is to ensure that only one pseudowire is used for forwarding in both directions by PE1, PE2 and PE3 in the absence of a native dual homing protocol between CE2 and PE2/PE3, such as MC-LAG. In normal operating conditions (the SAPs on PE2 and PE3 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 35: Master-slave pseudowire redundancy



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

This is achieved as follows. The **standby-signaling-master** state is activated on the VLL endpoint in PE1. In this case, a spoke-SDP is blocked in the transmit direction at this master endpoint if it is either in

operDown state, or it has lower precedence than the highest precedence spoke-SDP, or the 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)

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



Note:


Although master-slave operation provides bidirectional blocking of a standby spoke-SDP during steady-state conditions, it is possible that the Tx directions of more than one slave endpoint can be active for transient periods during a fail-over operation. This is due to slave endpoints transitioning a spoke-SDP from standby to active receiving or processing a pseudowire preferential forwarding status message before those transitioning a spoke-SDP to standby. This transient condition is most likely when a forced switch-over is performed, or the relative preferences of the spoke-SDPs is changed, or the active spoke-SDP is shutdown at the master endpoint. During this period, loops of unknown traffic may be observed. Fail-overs due to common network faults that can occur during normal operation, a failure of connectivity on the path of the spoke-SDP or the SAP, would not result in such loops in the datapath.

3.5.5.2 Operation of master-slave pseudowire redundancy with existing scenarios

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

The following figure shows a VLL resilience path example. An sample configuration follows.

Figure 36: VLL resilience



OSSG246



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.

```
PE1
configure service epipe 1
endpoint X
exit
endpoint Y
    revert-time 0
    standby-signaling-master
    exit
    sap 1/1/1:100 endpoint X
    spoke-sdp 1:100 endpoint Y
precedence primary
    spoke-sdp 2:200 endpoint Y
precedence 1
PE2
configure service epipe 1
endpoint X
exit
    sap 2/2/2:200 endpoint X
    spoke-sdp 1:100
        standby-signaling-slave
```

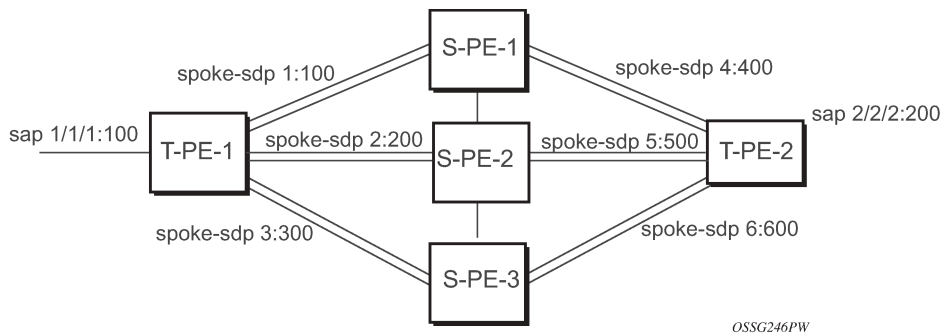
PE3

```
configure service epipe 1
endpoint X
exit
    sap 3/3/3:300 endpoint X
    spoke-sdp 2:200
        standby-signaling-slave
```

3.5.5.3 VLL resilience for a switched pseudowire path

The following figure displays VLL resilience for a switched pseudowire path example. A sample configuration follows.

Figure 37: VLL resilience with pseudowire switching



```
T-PE-1
configure service epipe 1
```

```
endpoint X
exit
endpoint Y
revert-time 100
standby-signaling-master
exit
sap 1/1/1:100 endpoint X
spoke-sdp 1:100 endpoint Y
  precedence primary
spoke-sdp 2:200 endpoint Y
  precedence 1
spoke-sdp 3:300 endpoint Y
  precedence 1
```

```
T-PE-2
configure service epipe 1
  endpoint X
  exit
  endpoint Y
  revert-time 100
  standby-signaling-slave
  exit
  sap 2/2/2:200 endpoint X
  spoke-sdp 4:400 endpoint Y
    precedence primary
  spoke-sdp 5:500 endpoint Y
    precedence 1
  spoke-sdp 6:600 endpoint Y
    precedence 1
```

```
S-PE-1
```

VC switching indicates a VC cross-connect so that the service manager does not signal the VC label mapping immediately but will put S-PE-1 into passive mode, as follows:

```
configure service epipe 1 vc-switching
  spoke-sdp 1:100
  spoke-sdp 4:400
```

3.5.5.4 Pseudowire redundancy service models


This section describes the various pseudowire redundancy scenarios as well as the algorithm used to select the active transmit object in a VLL endpoint.

The redundant VLL service model is described in the following section, [Pseudowire redundancy — redundant VLL service model](#).

3.5.5.4.1 Pseudowire redundancy — 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 38: 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 shown in [Figure 38: Redundant VLL endpoint objects](#).

[Figure 38: Redundant VLL endpoint objects](#) is just an example and the Y endpoint can also have a SAP 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 application. 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 any specific 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 an MC-LAG instance. A SAP that is not part of an 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

3.5.5.4.2 Pseudowire redundancy — T-LDP status notification handling rules

Using [Figure 38: 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:

Any allowed combination of objects as specified in [Pseudowire redundancy — redundant VLL service model](#) can be used on endpoints X and Y.

The following sections refer to the specific combination objects in [Figure 38: Redundant VLL endpoint objects](#) as an example to describe the more general rules.

3.5.5.5 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 an MC-LAG instance, the forwarding status of active is always advertised.

When the SAP in endpoint X is part of an 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 an MC-LAG instance, the forwarding status sent over all Y endpoint spoke-SDPs should always be set to zero (active by default).

3.5.5.6 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:

Ethernet SAP does not support SAP OAM protocol. All other SAP types cannot exist on the same endpoint as an ICB spoke-SDP since 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
- receive a SAP down notification 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 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.

3.5.6 MPLS entropy label and hash label

MPLS entropy label (RFC 6790) and the Flow Aware Transport label (known as the hash label) (RFC 6391) allow LSR nodes in a network to load-balance labeled packets in a much more granular fashion than allowed by hashing on the standard label stack. For more information, see the *7210 SAS-K 2F6C4T, K 3SFP+ 8C MPLS Guide*.

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C do not support MPLS entropy label or hash label.

3.6 VLL service considerations

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

3.6.1 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

3.6.2 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



Note:

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.

3.6.3 QoS policies

On the 7210 SAS-D and 7210 SAS-Dxp, when applied to 7210 SAS device Epipe services, service ingress QoS policies only create the unicast meters defined in the policy. The multi-point meters are not created on the service. Egress QoS policies function as with other services where the class-based queues per port are available.

On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, when applied to 7210 SAS device Epipe services, service ingress QoS policies only create the unicast queues defined in the policy. The multi-point queues are not created on the service. Service egress QoS policies function as with other services where the class-based queues are created as defined in the policy.



Note:

- Both Layer 2 or Layer 3 criteria can be used in the QoS policies for traffic classification in a service. For more information, see the *7210 SAS-D, Dxp Quality of Service Guide* and the *7210 SAS-K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Quality of Service Guide*.
- For the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, multicast queues are also created for the service ingress QoS policy, but only unicast queues are used for traffic.

3.6.4 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.

3.6.5 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 Epipe implementation does not perform any MAC learning on the service, so Epipe services do not consume any MAC hardware resources.

3.7 Configuring a VLL service with CLI

This section describes how to configure Virtual Leased Line (VLL) services using the command line interface.

3.7.1 Basic configurations

3.7.2 Common configuration tasks

This section provides a brief overview of the tasks that must be performed to configure the VLL services and provides the CLI commands:

- Associate the service with a customer ID.
- Define SAP parameters:
 - Optional - select ingress QoS policies (configured in the **config>qos** context)
 - Optional - select accounting policy (configured in the **config>log** context)
- Enable the service.

3.7.2.1 Creating an Epipe service for 7210 SAS-D

Use the following syntax to create an Epipe service.

```
config>service# epipe service-id [customer customer-id] [create] [svc-sap-type {null-star | any | dot1q-preserve}] [customer-vid vlan-id]
```

The following is a sample Epipe configuration output:

```
A:ALA-1>config>service# info
-----
...
    epipe 15 customer 40 svc-sap-type any create
        description "Local Epipe Service with ANY SVC_SAP_TYPE"
        no shutdown
    exit
-----
A:ALA-1>config>service#
```

3.7.2.2 Creating an Epipe service for 7210 SAS-Dxp

Use the following syntax to create an Epipe service.

```
config>service# 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] [uplink-type {l2 | mpls}]
```

The following is a sample Epipe configuration output:

```
A:ALA-1>config>service# info
-----
...
    epipe 15 customer 40 svc-sap-type any create
        description "Local Epipe Service with ANY SVC_SAP_TYPE"
        no shutdown
    exit
-----
A:ALA-1>config>service#
```

3.7.2.3 Creating an Epipe service with range SAPs

The following is a sample connection profile used to configure a range of SAPs and an Epipe configuration using the connection profile:

```
*A:7210SAS>config>connprof# info
-----
    ethernet
        ranges 0 2804-2805 2810-2811 2813 2832-2839
    exit
-----

*A:7210SAS>config>service>epipe# info
-----
    description "Default epipe description for service id 292"
    sap 1/1/4:292.* create
        description "Default sap description for service id 292"
```

```
        exit
    exit
    sap 1/1/9:cp-292 create
        description "Default sap description for service id 292"
    exit
    exit
    no shutdown
-----
```

3.7.2.4 Creating an Epipe service for 7210 SAS-K 2F1C2T

Use the following syntax to create an Epipe service.

```
config>service# epipe service-id [customer customer-id] [create] [svc-sap-type {any|dot1q-  
range}]
```

The following is a sample Epipe configuration output:

```
*A:SAH01-051>config>service>epipe$ info detail
-----
        shutdown
        no description
        service-mtu 1514
        eth-cfm
        exit
        pbb
        exit
-----
*A:SAH01-051>config>service>epipe$
-----
```

3.7.2.5 Creating an Epipe service for 7210 SAS-K 2F1C2T with range SAPs

The following is a sample connection profile used to configure a range of SAPs and an Epipe configuration using the connection profile:

```
*A:SAH01-051>config>connprof$ info detail
-----
        no description
        ethernet
            no ranges
        exit
-----
*A:SAH01-051>config>connprof$

*A:SAH01-051>config>service>epipe$ info detail
-----
        shutdown
        no description
        service-mtu 1514
        eth-cfm
        exit
        pbb
        exit
-----
```

```
*A:SAH01-051>config>service>epipe$
```

3.7.2.5.1 Configuring Epipe SAP parameters

A default QoS policy is applied to each ingress SAP, and a default access egress QoS policy is applied on the port where SAP is egressing. The access egress QoS policy is common to all SAPs on that port. Additional QoS policies can be configured in the **config>qos** context. Filter policies are configured in the **config>filter** context and explicitly applied to a SAP. There are no default filter policies.

```
- 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
```

3.7.2.5.1.1 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.

The following is a sample SAP configuration output for local Epipe service 500 on SAP 1/1/2 and SAP 1/1/3 on ALA-1:

```
A:ALA-1>config>service# epipe 500 customer 5 create
config>service>epipe$ description "Local epipe service
config>service>epipe# sap 1/1/2 create
config>service>epipe>sap? ingress
config>service>epipe>sap>ingress# qos 20
config>service>epipe>sap>ingress# filter ip 1
config>service>epipe>sap>ingress# exit
config>service>epipe>sap# egress
config>service>epipe# sap 1/1/3 create
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# qos 555
config>service>epipe>sap>ingress# filter ip 1
config>service>epipe>sap>ingress# exit
config>service>epipe>sap# no shutdown
config>service>epipe>sap# exit
```

```
A:ALA-1>config>service# info
```

```
.....
...
    epipe 500 customer 5 create
      description "Local epipe service"
```

```
    sap 1/1/2 create
      ingress
        qos 20
        filter ip 1
      exit
    exit
  sap 1/1/3 create
    ingress
      qos 555
      filter ip 1
    exit
  exit
no shutdown
exit
-----
A:ALA-1>config>service#
```

The following are sample SAP configurations output 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
exit

  epipe 5500 customer 5 vpn 5500 create
    description "Distributed epipe service to west coast"
    sap 441/1/4:550 create
      ingress
        qos 654
        filter ip 1020
      exit
    exit
  exit
exit
```

3.7.2.5.2 Distributed Epipe service

To configure a distributed Epipe service, you must configure service entities on the originating and far-end nodes. You should use the same service ID on both ends (for example, Epipe 5500 on ALA-1 and Epipe 5500 on ALA-2). The spoke-sdp sdp-id:vc-id must match on both sides. A distributed Epipe consists of two SAPs on different nodes.

By default, QoS policy ID 1 is applied to ingress service SAPs. On egress, QoS policies are associated with a port. Existing filter policies can be associated with service SAPs on ingress and egress.

Meters (defined in SAP-ingress policies) can be applied on ingress, 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](#).

The following example shows the command usage to configure a distributed service between ALA-1 and ALA-2:

```
A:ALA-1>epipe 5500 customer 5 create
config>service>epipe$ description "Distributed epipe service to east coast"
config>service>epipe# sap 221/1/3:21 create
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# qos 555
config>service>epipe>sap>ingress# filter ip 1
config>service>epipe>sap>ingress# exit
config>service>epipe>sap# no shutdown
config>service>epipe>sap# exit
config>service>epipe#

A:ALA-2>config>service# epipe 5500 customer 5 create
config>service>epipe$ description "Distributed epipe service to west coast"
config>service>epipe# sap 441/1/4:550 create
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# filter ip 1020
config>service>epipe>sap>ingress# exit
config>service>epipe>sap# egress
config>service>epipe>sap>egress# filter ip 6
config>service>epipe>sap>egress# exit
config>service>epipe>sap# no shutdown
config>service>epipe#
```

The following are sample SAP configuration outputs 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#
```

3.7.2.5.3 Configuring SDP bindings

Figure 39: SDPs — unidirectional tunnels displays an example of a distributed Epipe service configuration between two routers, identifying the service and customer IDs, and the unidirectional SDPs required to communicate to the far-end routers.

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

Figure 39: 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:

```
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#
```

3.7.2.5.4 Using spoke-SDP control words

The **control word** command provides the option to add a control word as part of the packet encapsulation for PW types for which the control word is optional. On 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. 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.

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#
```

Use the following syntax 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
```

3.7.2.5.5 Configuring ingress and egress SAP parameters

By default, QoS policy ID 1 is applied to ingress service SAPs. Existing filter policies or other existing QoS policies can be associated with service SAPs on ingress and egress ports.

Ingress and egress SAP parameters can be applied to local and distributed Epipe service SAPs.

The following are sample SAP ingress and egress parameters:

```
ALA-1>config>service# epipe 5500
config>service>epipe# sap 1/1/3:21
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# qos 555
config>service>epipe>sap>ingress# filter ip 1
config>service>epipe>sap>ingress# exit
config>service>epipe>sap#
```


The following is a sample Epipe SAP ingress and egress 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#
```

3.7.2.5.6 Configuring VLL resilience

The following figure displays an example to create VLL resilience. The zero revert-time value means that the VLL path will be switched back to the primary immediately after it comes back up.

Figure 40: VLL resilience




The following is a sample configuration output on PE1(a 7210 SAS-K 2F6C4T node):

```
*A:ALA-48>config>service>epipe# info
-----
endpoint "x" create
exit
endpoint "y" create
exit
spoke-sdp 1:100 endpoint "y" create
precedence primary
exit
spoke-sdp 2:200 endpoint "y" create
precedence 1
exit
no shutdown
-----
*A:ALA-48>config>service>epipe#
```

3.7.2.5.7 Configuring VLL resilience for a switched pseudowire path

The following figure shows an example of VLL resilience with pseudowire switching.

Figure 41: VLL resilience with pseudowire switching



OSSG247

T-PE1

The following is a sample configuration output on TPE1 (which can be a 7210 SAS-K 2F6C4T or 7210 SAS-K 3SFP+ 8C):

```
*A:ALA-48>config>service>epipe# info
-----
endpoint "x" create
exit
endpoint "y" create
exit
sap 1/1/1:100 endpoint "x" create
exit
spoke-sdp 1:100 endpoint "y" create
precedence primary
exit
spoke-sdp 2:200 endpoint "y" create
precedence 1
exit
spoke-sdp 3:300 endpoint "y" create
precedence 1
```

```
exit
no shutdown
-----
*A:ALA-48>config>service>epipe#
```

T-PE2

The following 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#
```

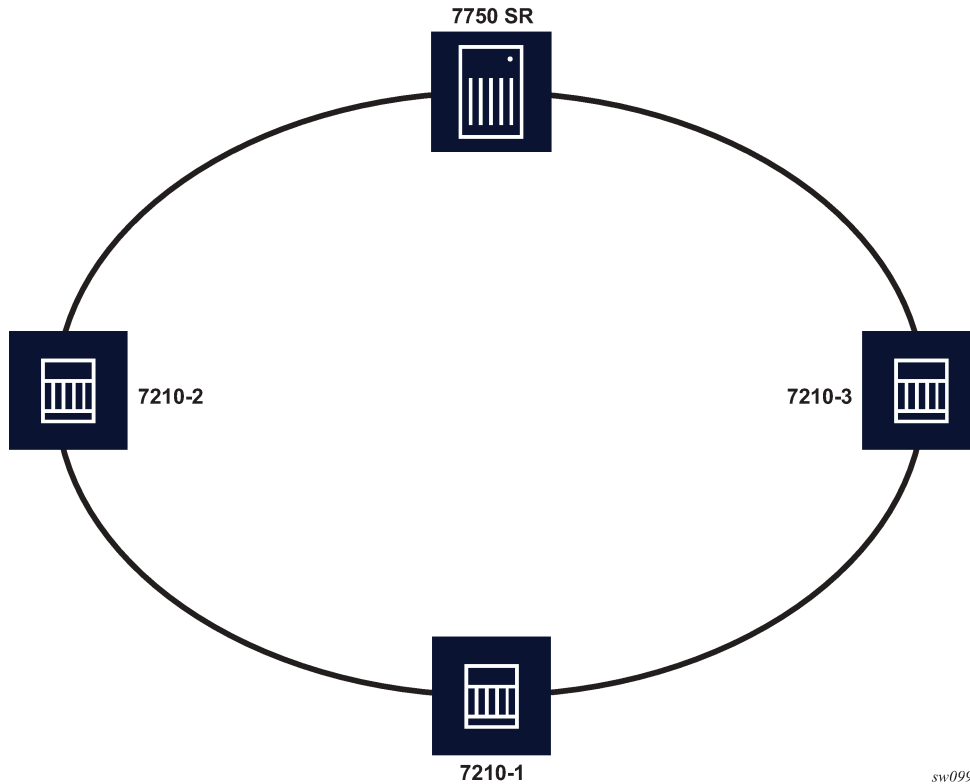
S-PE1

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#
```

3.7.2.6 Configuring default QinQ SAPs for Epipe transit traffic in a ring scenario in access-uplink mode

Figure 42: Default QinQ SAP for transit traffic in a ring scenario



In the preceding 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 (example – traffic from 7210-2 to 7210-3 OR from 7210-2 to 7750 –SR onto the core network).

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. The following are the sample configuration commands in this scenario:

```
ALA-1>config>service# epipe 8 customer 1 svc-sap-type null-star create
  sap 1/1/5:.* create
    statistics
    ingress
    exit
  exit
exit
  sap 1/1/6:.* create
    statistics
    ingress
    exit
  exit
exit
no shutdown
exit
```

3.7.3 Service management tasks

This section describes the Epipe service management tasks.

3.7.3.1 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
```

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#
```

3.7.3.2 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
```

3.7.3.3 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
```

3.7.3.4 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#
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
```

3.8 VLL services command reference

3.8.1 Command hierarchy

3.8.1.1 Epipe service configuration commands

- [Epipe service configuration commands](#)
 - [Epipe global commands for 7210 SAS-D](#)
 - [Epipe global commands for 7210 SAS-Dxp](#)
 - [Epipe global commands for 7210 SAS-K 2F1C2T](#)
 - [Epipe global commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)
 - [Epipe SAP configuration commands](#)
 - [Epipe SAP meter override commands](#)
 - [Epipe SAP statistics commands for 7210 SAS-D](#)
 - [Epipe SAP statistics commands for 7210 SAS-Dxp](#)

- Epipe SAP configuration — QoS and filter commands for 7210 SAS-D and 7210 SAS-Dxp
- Epipe SAP configuration — QoS and filter commands for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C
- Show commands
- Clear commands

3.8.1.1.1 Epipe global commands for 7210 SAS-D

```
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
- description description-string
- no description
- [no] endpoint endpoint-name [create]
- sap sap-id [create]
- no sap sap-id
- service-name service-name
- no service-name
- [no] shutdown
```

3.8.1.1.2 Epipe global commands for 7210 SAS-Dxp

```
config
- service
- 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] [uplink-type {l2 | mpls}]
- no epipe service-id
- description description-string
- no description
- sap sap-id [create]
- no sap sap-id
- service-name service-name
- no service-name
- [no] shutdown
```

3.8.1.1.3 Epipe global commands for 7210 SAS-K 2F1C2T

```
config
- service
- epipe service-id [customer customer-id] [create]
- epipe service-id [customer customer-id] [create] [svc-sap-type {any | dot1q-range}]
- epipe service-id [customer customer-id] [create] [svc-sap-type {any | dot1q-range}]
- no epipe service-id
- description description-string
- no description
- sap sap-id [create]
- no sap sap-id
- service-mtu octets
- no service-mtu
```

```
- service-name service-name  
- no service-name  
- [no] shutdown
```

3.8.1.1.4 Epipe global commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
config  
- service  
- epipe service-id [customer customer-id] [create]  
- epipe service-id [customer customer-id] [create] [svc-sap-type {any | dot1q-range}]  
- epipe service-id [customer customer-id] [create] [svc-sap-type {any | dot1q-range}]  
- no epipe service-id  
- description description-string  
- no description  
- [no] endpoint endpoint-name [create]  
- active-hold-delay active-endpoint-delay  
- no active-hold-delay  
- revert-time [revert-time | infinite]  
- no revert-time  
- standby-signaling-master  
- [no] standby-signaling-master  
- sap sap-id [create]  
- no sap sap-id  
- service-mtu octets  
- no service-mtu  
- service-name service-name  
- no service-name  
- [no] shutdown  
- [no] spoke-sdp
```

3.8.1.1.5 Epipe SAP configuration commands

```
config  
- service  
- epipe service-id  
- no epipe service-id  
- no sap sap-id  
- accounting-policy acct-policy-id  
- no accounting-policy acct-policy-id  
- [no] collect-stats  
- description description-string  
- no description  
- eth-cfm  
- [no] mep mep-id domain md-index association ma-index [direction {up |  
down}] primary-vlan-enable  
- [no] ais-enable  
- [no] client-meg-level [[level [level ...]]]  
- [no] interval {1 | 60}  
- [no] priority priority-value  
- no send-ais-on-port-down  
- send-ais-on-port-down  
- [no] ccm-enable  
- [no] ccm-ltm-priority priority  
- defect-oper-group name  
- no defect-oper-group  
- [no] description  
- [no] eth-test-enable  
- [no] bit-error-threshold bit-errors
```



```
xcon | noXcon}
    - [no] test-pattern {all-zeros | all-ones} [crc-enable]
    - [no] fault-propagation-enable {use-if-tlv | suspendccm}
    - low-priority-defect {allDef | macRemErrXcon | remErrXcon | errXcon |
- [no] mac-address mac-address
- [no] one-way-delay-threshold seconds
- [no] shutdown
- [no] squelch-ingress-levels [md-level [md-level]]
- mac-swap-enable
- no mac-swap-enable
- ethernet
  - [no] ais-enable
- [no] ignore-oper-down
- [no] shutdown
```

3.8.1.1.6 Epipe SAP meter override commands

```
config
- service
  - epipe service-id
  - no epipe service-id
    - no sap sap-id
    - ingress
      - meter-override
        - meter meter-id [create]
        - no meter meter-id
          - adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
          - cbs size [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
```

3.8.1.1.7 Epipe SAP statistics commands for 7210 SAS-D

```
config
- service
  - epipe service-id
  - no epipe service-id
    - no sap sap-id
    - statistics
      - egress
        - [no] forwarded-count
      - ingress
        - counter-mode {in-out-profile-count | forward-drop-count}
        - [no] received-count
```

3.8.1.1.8 Epipe SAP statistics commands for 7210 SAS-Dxp

```
config
- service
  - epipe service-id
```

```
- no epipe service-id
  - no sap sap-id
    - statistics
      - ingress
        - counter-mode {in-out-profile-count | forward-drop-count}
```

3.8.1.1.9 Epipe spoke-SDP configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
config
- service]
  - epipe service-id
  - no epipe service-id
    - spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] [no-endpoint]
    - spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] endpoint
    - no spoke-sdp sdp-id[:vc-id]
      - accounting-policy acct-policy-id
      - no accounting-policy
      - adv-service-mtu octets
      - no adv-service-mtu
      - [no] collect-stats
      - [no] control-word
      - [no] description
      - [no] egress
        - [no] vc-label egress-vc-label
      - eth-cfm
        - [no] ais-enable
        - [no] mep mep-id domain md-index association ma-index [direction {up |
down}]
          - [no] ais-enable
            - [no] client-meg-level [[level [level ...]]]
            - [no] interval {1 | 60}
            - [no] priority priority-value
          - [no] ccm-enable
          - [no] ccm-ltm-priority priority
          - [no] description
          - [no] eth-test-enable
            - [no] bit-error-threshold bit-errors
            - [no] test-pattern {all-zeros | all-ones} [crc-enable]
          - [no] fault-propagation-enable {use-if-tlv | 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 [signal-capability]
        - no hash-label
        - [no] ingress
          - [no] vc-label egress-vc-label
        - precedence [precedence-value] primary]
        - no precedence
        - no pw-status-signaling
        - pw-status-signaling
        - [no] shutdown
        - suppress-vc-vlan
        - no suppress-vc-vlan
        - vlan-vc-tag 0..4094
```

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

3.8.1.1.10 Epipe SAP configuration — QoS and filter commands for 7210 SAS-D and 7210 SAS-Dxp

```
config
- service
- epipe service-id [customer customer-id] [create] [svc-sap-type {null-star | dot1q-
preserve | any | dot1q-range}] [customer-vid vlan-id]
- epipe service-id [customer customer-id] [create] [svc-sap-type {null-star | dot1q-
preserve | dot1q}] [customer-vid vlan-id]
- no epipe service-id
- no sap sap-id [create]
- egress
- aggregate-meter-rate rate-in-kbps [burst burst-in-kbits] [enable-stats]
- no aggregate-meter-rate
- filter [ip ip-filter-id]
- filter [ipv6 ipv6 -filter-id]
- filter [mac mac-filter-id]
- no filter [ip ip-filter-id] [ipv6 ipv6 -filter-id] [mac mac-filter-id]
- ingress
- aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
- no aggregate-meter-rate
- filter [ip ip-filter-id]
- filter [ipv6 ipv6-filter-id]
- filter [mac mac-filter-id]
- no filter [ip ip-filter-id] [ ipv6 ipv6-filter-id] [mac mac-filter-id]
- qos policy-id
- no qos
```

3.8.1.1.11 Epipe SAP configuration — QoS and filter commands for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

```
config
- service
- epipe service-id [customer customer-id] [create] [svc-sap-type {any | dot1q-range}]
- no epipe service-id
- no sap sap-id [create]
- egress
- agg-shaper-rate agg-rate
- no agg-shaper-rate
- dot1p-inner dot1p-inner
- no dot1p-inner
- no dot1p-outer
- dot1p-outer dot1p-outer
- filter [ip ip-filter-id]
- filter [ipv6 ipv6 -filter-id]
- filter [mac mac-filter-id] (app
- no filter [ip ip-filter-id] [ipv6 ipv6 -filter-id] [mac mac-filter-id]
- qos policy-id
- no qos
- ingress
- agg-shaper-rate agg-rate
- no agg-shaper-rate
- aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
- no aggregate-meter-rate
- filter [ip ip-filter-id]
```

```
- filter [ ipv6 ipv6-filter-id]
- filter [mac mac-filter-id]
- no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-id]
- qos policy-id
- no qos
```

3.8.1.2 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)]
```

3.8.1.3 Show commands

```
show
- service
  - id service-id
    - all
    - base
    - sap
    - stp [sap-id] [detail]]
  - sap-using [sap sap-id]
  - sap-using [ingress | egress] filter filter-id
  - sap-using [ingress ] qos-policy qos-policy-id
  - service-using [epipe] [vpls] [mirror] [i-vpls] [m-vpls] [sdp sdp-id]
[customer customer-id]
show
- connection-profile [conn-prof-id] [associations]
```

3.8.1.4 Clear commands

```
clear
- service
  - id service-id
  - statistics
    - id service-id
  - counters
  - sap sap-id {all | counters | stp | l2pt}
```

3.8.2 Command descriptions

3.8.2.1 VLL service configuration commands

- [Generic commands](#)

- [VLL global commands](#)
- [VLL SAP commands](#)
- [Connection profile commands](#)
- [Service QoS and filter policy commands](#)
- [Spoke-SDP commands](#)

3.8.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

config>service>epipe

config>service>epipe>sap

config>service>epipe>sap>eth-cfm>mep

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.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up, then tries to enter the operationally up state.

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

description

Syntax

description *description-string*

no description

Context

config>service>epipe

config>service>epipe>sap

```
config>service>epipe>spoke-sdp  
config>connection-profile
```

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 **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

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.

eth-cfm

Syntax

```
eth-cfm
```

Context

```
config>service>epipe>sap
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Description

This command provisions the maintenance endpoint (MEP).

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

For more information about ETH-CFM support for different services, see the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C OAM and Diagnostics Guide.

Parameters

mep-id

Specifies the maintenance association endpoint 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}

Specifies the direction in which the maintenance association (MEP) faces on the bridge port.

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

down — Sends ETH-CFM messages away from the MAC relay entity.

primary-vlan-enable

Keyword that provides a method for linking with the primary VLAN configured under the bridge-identifier for the MA. This must be configured as part of the creation step and can only be changed by deleting the MEP and recreating it. Primary VLANs are only supported under Ethernet SAPs. This parameter is only supported on the 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

3.8.2.1.2 VLL global commands

epipe

Syntax

```
epipe service-id [customer customer-id] [create] [svc-sap-type {null-star|dot1q|dot1q-preserve}]  
    [customer-vid vlan-id]
```

```
epipe service-id [customer customer-id] [create] [svc-sap-type {null-star|dot1q-preserve|any|dot1q-range}] [customer-vid vlan-id]
```

```
epipe service-id [customer customer-id] [create] [svc-sap-type {any|dot1q-range}]
```

```
epipe service-id [customer customer-id] [create] [svc-sap-type {null-star|dot1q|dot1q-preserve|any|dot1q-range|qinq-inner-tag-preserve}]
```

```
epipe [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] [uplink-type {I2 | mpls}]
```

```
no epipe service-id
```

Context

config>service

Platforms

Supported on all 7210 SAS platforms as described in this document

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.



Note:

- 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T platforms only support local SAP to SAP service.
- 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C platforms support both local and distributed service.

No MAC learning or filtering is provided on an Epipe.

When a service is created, the **customer** keyword and *customer-id* must be specified and 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 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.

By default, no Epipe services exist until they are explicitly created with this command.

The tables in the following links list the allowed SAPs for a particular value of *svc-sap-type* on different 7210 platforms:

- [Table 7: SAP and service combinations for 7210 SAS-D and 7210 SAS-Dxp](#)
- [Table 8: SAP and service combinations for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C](#)

The **no** form of this command deletes the Epipe service instance with the specified *service-id*. The service cannot be deleted until the service has been shut down.

Parameters

service-id

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.

service-id-epipe

Values *service-id*: 1 to 2147483648
 svc-name: 64 characters maximum

customer *customer-id*

Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deletion.

epipe customer

Values 1 to 2147483647

svc-sap-type

Keyword that specifies the type of access SAPs and access-uplink SAPs allowed in the service.

vpn

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). This is supported only on 7210 SAS-D and 7210 SAS-Dxp.

dot1q — Specifies that the allowed SAPs in the service are dot1q SAPs and dot1q explicit null SAPs. This is supported only on 7210 SAS-Dxp.

dot1q-preserve — Specifies that the allowed SAPs in the service are dot1q. The dot1q ID is not stripped after packets match the SAP. This is supported only on 7210 SAS-D and 7210 SAS-Dxp.

dot1q-range — Specifies that the access SAP in the service can use VLAN ranges as the SAP tags. The VLAN ranges are configured using the **configure>connection-profile** CLI command. On ingress of the access dot1q SAP using VLAN ranges, the outermost tag is not removed before forwarding. This is supported only on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

any — Keyword that allows any SAP type. This is supported only on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

Default any

customer-vid *vlan-id*

Defines the dot1q VLAN ID to be specified while creating the local dot1q SAP for **svc-sap-type dot1q-preserve**. Applicable only for Access-uplink mode.

Values 1 to 4094

create

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

endpoint

Syntax

[no] **endpoint** *endpoint-name*

Context

config>service>epipe

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies that the node will delay sending the change in the T-LDP status bits for the VLL endpoint when the MC-LAG transitions the LAG subgroup that 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.

A value of zero (default) means that when the MC-LAG transitions the LAG subgroup that hosts the SAP for this VLL endpoint from **active** to **standby**, the node immediately sends new T-LDP status bits indicating the new value of **standby** over the spoke-SDPs that 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 that hosts the SAP from "standby" to "active", or when any object in the endpoint transitions to an operationally up state.

Default

0

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the time to wait before reverting back to the primary spoke-SDP defined on this service endpoint, after having failed over to a backup spoke-SDP.

Parameters

revert-time

Specify the time, in seconds, to wait before reverting to the primary SDP.

Values 0 to 600

infinite

Keyword that causes the endpoint to be non-revertive.

standby-signaling-master

Syntax

[no] standby-signaling-master

Context

config>service>epipe>endpoint

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

When this command is enabled, the pseudowire standby bit (value 0x00000020) is sent to T-LDP peer for each spoke-SDP of the endpoint that is selected as a standby.

Default

no standby-signaling-master

service-mtu

Syntax

service-mtu *octets*

no service-mtu

Context

config>service>epipe

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

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.

In the event that a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, all associated SAP and SDP binding operational states are automatically reevaluated.

The **no** form of this command reverts to the default **service-mtu** value for the indicated service type.

Default

epipe: 1514

The following table displays MTU values for specific VC types.

Table 25: MTU values

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

Description

This command configures an optional service name that 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 or administrator to identify and manage services.

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, up to 64 characters, to identify the service. Service names may not begin with an integer (0-9).

3.8.2.1.3 VLL SAP commands

sap

Syntax

sap *sap-id* [**create**]

no sap *sap-id*

Context

config>service>epipe

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters that identify the service access point on the interface and within the 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.

On the 7210 SAS-D and 7210 SAS-Dxp, 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 services. This restriction does not apply to the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, or 7210 SAS-K 3SFP+ 8C.

A SAP can be associated with only 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 also be defined on an 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, and QinQ.
- 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 be deleted. For IES, the IP interface must be shut down before the SAP on that interface may be removed.

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 only allowed 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 or disabled in the **environment>create** context.

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy [*acct-policy-id*]

Context

config>service>epipe>sap

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables 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

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>epipe>sap

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables accounting and statistical data collection for the SAP, network port, or IP interface. When applying accounting policies, by default the data is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued, the statistics are still accumulated by the cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued, the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default

no collect-stats

ethernet

Syntax

ethernet

Context

config>service>epipe>sap

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures Ethernet properties in this SAP.

ais-enable

Syntax

[no] **ais-enable**

Context

config>service>epipe>sap>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the generation and the reception of AIS messages.

client-meg-level

Syntax

client-meg-level [[*level* [*level* ...]]

no client-meg-level

Context

config>service>epipe>sap>eth-cfm>ais-enable

Platforms

Supported on all 7210 SAS platforms as described in this document

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

level

Specifies the client MEG level.

Values 1 to 7

Default 1

interval

Syntax

interval {1 | 60}

no interval

Context

config>service>epipe>sap>eth-cfm>ais-enable

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the transmission interval of AIS messages.

Parameters

1 | 60

Specifies the transmission interval of AIS messages, in seconds.

Default 1

priority

Syntax

priority *priority-value*

no priority

Context

config>service>epipe>sap>eth-cfm>ais-enable

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

Values 0 to 7

Default 1

ccm-enable

Syntax

[no] **ccm-enable**

Context

config>service>epipe>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Platforms

Supported on all 7210 SAS platforms as described in this document

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

defect-oper-group

Syntax

defect-oper-group *name*

no defect-oper-group

Context

config>service>epipe>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document, except 7210 SAS-D

Description

This command configures an operational group for fault propagation.

See the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C OAM and Diagnostics Guide for more information about fault propagation in an Epipe service.

The **no** form of this command removes the operational group.

Parameters

name

Specifies the operational group name, up to 32 characters.

llf

Syntax

[no] **llf**

Context

config>service>epipe>sap>ethernet

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables Link Loss Forwarding (LLF) on an Ethernet port or an ATM port. This command provides an end-to-end OAM fault notification for Ethernet VLL service. It brings down the Ethernet port (Ethernet LLF) toward the attached CE when there is a local fault on the Pseudowire or service, or a remote fault on the SAP or pseudowire, signaled with label withdrawal or T-LDP status bits. It ceases when the fault disappears.

The Ethernet port must be configured for null encapsulation.

ignore-oper-down

Syntax

[no] ignore-oper-down

Context

config>service>epipe>sap

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures the optional command for a specific SAP to ignore the transition of the operational state to down when a SAP fails. Only a single SAP in an Epipe may use this option.

Default

no ignore-oper-down

send-ais-on-port-down

Syntax

[no] send-ais-on-port-down

Context

config>service>epipe>sap>eth-cfm>mep>ais-enable

config>service>vpls>sap>eth-cfm>mep>ais-enable

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies that ETH-AIS should be generated for client MEPs immediately when port down event is detected on the port where the server MEP (and the associated SAP) resides. By default, the system generates an ETH-AIS message, if enabled, when CCM messages are not received within the configured time period. On a subsequent port up event, the AIS messages continue to be sent until valid CCMs are received. If there are no remote MEPs configured for the MEP, on a subsequent port up event, the AIS messages are not sent.

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

Default

no send-ais-on-port-down

eth-test-enable

Syntax

[no] eth-test-enable

Context

config>service>epipe>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the Ethernet test functionality on MEP. For ETH-test to work, users must configure ETH-test parameters on both sender and receiver nodes. The ETH-test can then be performed 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 performed 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>epipe>sap>eth-cfm>mep>eth-test-enable

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Keyword that specifies to use all zeros in the test pattern.

all-ones

Keyword that specifies to use all ones in the test pattern.

crc-enable

Keyword that generates a CRC checksum.

Default all-zeros

bit-error-threshold

Syntax

bit-error-threshold *errors*

no bit-error-threshold

Context

config>service>epipe>sap>eth-cfm>mep>eth-test-enable

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command is used to specify the threshold value of bit errors.

one-way-delay-threshold

Syntax

one-way-delay-threshold *seconds*

Context

config>service>epipe>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables eth-test functionality on MEP.

Parameters

seconds

Specifies the one-way delay threshold in seconds.

Values 0 to 600

Default 3

snquch-ingsress-levels

Syntax

snquch-ingsress-levels [*md-level* [*md-level*]]

no snquch-ingsress-levels

Context

config>service>epipe>sap>eth-cfm

Platforms

7210 SAS-Dxp

Description

This command configures the levels of the ETH-CFM packets that are silently discarded on ingress into the SAP from the wire that matches the service delineation of the SAP.

All ETH-CFM packets inbound to the SAP that match the configured levels are dropped without regard for any other ETH-CFM criteria. No statistical information or drop count is available for any ETH-CFM packet that is silently discarded by this option.

The user must configure a complete contiguous list of MD-levels up to the highest level that are to be dropped. The command must be retyped in complete form to modify a previous configuration if the user does not want to delete it first.

Entering the command with no associated MD-level information is equivalent to the **no** form of this command.

The **no** form of this command removes the silent discarding of previously matching ETH-CFM PDUs.

Default

no snquch-ingsress-levels

Parameters

md-level

Identifies the level

Values 0 to 7

mac-swap-enable

Syntax

[no] mac-swap-enable

Context

config>service>epipe>sap config>service>vpls>sap

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures that the source and destination MAC address for all the packets to be sent out of the SAP is swapped, then the packets are looped back to ingress of the SAP. SAP loopback is typically used with Y.1564 testhead functionality to loop back packets received on the end that is acting as a reflector, looping back packets received from the testhead generator.

This command is available for testing VLL services and VPLS services only. When enabled, no packets egress the service SAP, and packets received on ingress are not processed.



Note:

- Before enabling this command, turn off all Layer 2 and IP control protocols (such as LACP, EFM, 802.1x and so on) on the device and its peer to prevent errors such as protocol flaps because of timeout and so on.
- When SAP loopback with MAC-swap is enabled for broadcast and multicast packets (that is, packets received with multicast or broadcast destination address), the source MAC address is used as the destination MAC address and the system MAC address is the source MAC address.
- The user can enable MAC swap on multiple SAPs in the same service. The user can also configure static MAC on SAPs and traffic to prevent running into the MAC move, relearning issues and unpredictable device behavior.

Default

no mac-swap-enable

mip

Syntax

mip [**mac** *mac-address*]

mip default-mac

no mip

Context

config>service>epipe>sap>eth-cfm

config>service>epipe>spoke-sdp>eth-cfm (supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C only)

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command allows Maintenance Intermediate Points (MIPs) to be created if **mhf-creation** for the MA is configured using the default option.

The **no** form of this command deletes the MIP.

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

fault-propagation-enable

Syntax

fault-propagation-enable {**use-if-tlv** | **suspend-ccm**}

no fault-propagation-enable

Context

config>service>epipe>sap>eth-cfm>mep

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

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.

3.8.2.1.4 Connection profile commands

connection-profile

Syntax

connection-profile *conn-prof-id* [**create**]

no connection-profile *conn-prof-id*

Context

config

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command creates a profile to configure the list of VLAN values to be assigned to a dot1q SAP in an Epipe service.

A connection profile can only be applied to a dot1q SAP that 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 1000

create

Keyword to create a connection profile.

ethernet

Syntax

ethernet

Context

config>connprof

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

Specifies the list of VLAN ranges or individual VLAN IDs used for mapping the specific VLANs to the Epipe SAP.

The system validates that the values specified are valid VLAN IDs in the range 0 to 4094 (VLAN ID 4095 is reserved). Ranges are specified in the format "a-b", the expression (a < b) should be true. Up to 32 individual VLAN values or VLAN ranges can be specified. A maximum of 8 VLAN ranges are allowed per connection profile.

Parameters

vlan-ranges

Specifies the list of VLAN ranges or individual VLAN IDs 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 to 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
```

3.8.2.1.5 Service QoS and filter policy commands

egress

Syntax

egress

Context

```
config>service>epipe>sap
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command forces VC VLAN-type forwarding in the datapath for spokes that have either VC-type. This command is not allowed on VC VLAN-type SDPs.

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

ingress

Syntax

```
ingress
```

Context

```
config>service>epipe>sap
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure ingress SAP QoS policies.

If no SAP-ingress QoS policy is defined, the system default SAP-ingress QoS policy is used for ingress processing.

agg-shaper-rate

Syntax

agg-shaper-rate *agg-rate*

no agg-shaper-rate

Context

config>service>epipe>sap>ingress

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures the aggregate rate for the access SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues.

The **no** form of this command disables the use of the SAP aggregate rate shaper. That is, the SAP can use up the maximum bandwidth available.

Default

no agg-shaper-rate

Parameters

agg-rate

Specifies the rate in kilobits per second.

Values 50 to 3000000 | max (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
64 to 20000000 | max (7210 SAS-K 3SFP+ 8C)

Default max

agg-shaper-rate

Syntax

agg-shaper-rate *agg-rate*

no agg-shaper-rate

Context

config>service>epipe>sap>egress

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command configures the aggregate rate for the access SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use egress queues.

The **no** form of this command disables the use of the SAP aggregate rate shaper. That is, the SAP can use up the maximum bandwidth available.

Default

no agg-shaper-rate

Parameters

agg-rate

Specifies the rate in kilobits per second.

Values 50 to 1000000 | max (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
64 to 10000000 | max (7210 SAS-K 3SFP+ 8C)

Default max

aggregate-meter-rate

Syntax

aggregate-meter-rate *rate-in-kbps* [**burst** *burst-in-kbits*] [**enable-stats**]

no aggregate-meter-rate

Context

config>service>epipe>sap>egress

config>service>ies>if>sap>egress

config>service>vpls>sap>egress

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures the access SAP egress aggregate policer. The rate (PIR) of the SAP egress aggregate policer must be specified. The user can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the traffic sent out of the SAP and determines the final disposition of the packet, which is either forwarded or dropped.

The user can optionally associate a set of two counters to count total forwarded packets and octets and total dropped packets and octets. When use of this counter is enabled, the amount of resources required

increases by twice the amount of resources taken up when the counter is not used. If the **enable-stats** keyword is specified during the creation of the meter, the counter is allocated by the software, if available. To free up the counter and relinquish its use, use the **no aggregate-meter-rate** command, then recreate the meter using the **aggregate-meter rate** command.

If egress frame-based accounting is used, the SAP egress aggregate meter rate accounts for the Ethernet frame overhead. The system accounts for 12 bytes of IFG and 8 bytes of start delimiter. Frame-based accounting does not affect the count of octets maintained by the counter (if in use).



Note:

- Before enabling this command for a SAP, resources must be allocated to this feature from the egress-internal-tcam resource pool using the **config system resource-profile egress-internal-tcam egress-sap-aggregate-meter** command. See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide* for more information.
- The egress aggregate meter is not FC-aware. The forward and drop decisions are made based on the order in which 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 4000000, max (7210 SAS-D)
1 to 20000000, max (7210 SAS-Dxp)

Default max

burst-in-kbits

Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 to 16384, default (7210 SAS-D)
4 to 2146959, default (7210 SAS-Dxp)

Default 512

enable-stats

Keyword to specify whether the counter to count forwarded and dropped packets and octets is allocated. If this keyword is used while configuring the meter, the counter is allocated.

aggregate-meter-rate

Syntax

aggregate-meter-rate *rate-in-kbps* [**burst** *burst-in-kbits*]

no aggregate-meter-rate

Context

config>service>epipe>sap>ingress

config>service>ies>if>sap>ingress

config>service>vpls>sap>ingress

config>service>vprn>if>sap>ingress

Platforms

7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

IES contexts apply only to the 7210 SAS-Dxp, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

VPRN contexts apply only to the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command allows the user to configure the access SAP ingress aggregate policer. The rate of the SAP ingress aggregate policer must be specified by the user. The user can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the ingress traffic on different FCs using policers to rate-limit the flow and determines the final disposition of the packet. The packet is either forwarded to an identified profile or dropped.



Note:

- The sum of the CIR of the individual FCs configured under the SAP cannot exceed the PIR rate configured for the SAP. Although the 7210 SAS software does not block this configuration, Nokia does not recommend this configuration.
- The queued traffic flows are not limited by the aggregate meter. That is, only metered flows can use the aggregate meter. Queue flows can use only the aggregate shaper.

The following table lists 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	Final packet color
Within CIR	Green	Within PIR	Green or In-profile
Within CIR	Green	Above PIR	Red and Dropped

Per FC meter operating rate	Per FC assigned color	SAP aggregate meter operating rate	Final packet color
Above CIR, Within PIR	Yellow	Within PIR	Yellow or Out-of-Profile
Above CIR, Within PIR	Yellow	Above PIR	Red or Dropped
Above PIR	Red	Within PIR	Red or Dropped
Above PIR	Red	Above PIR	Red or Dropped

The SAP ingress meter counters increment the packet or octet counts based on the final disposition of the packet.

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.

Values 1 to 20000000, max (7210 SAS-Dxp)
 50 to 3000000, max (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T)
 64 to 20000000, max (7210 SAS-K 3SFP+ 8C)

Default max

burst-in-kbits

Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 to 2146959, default (7210 SAS-Dxp)
 1 to 16384, default (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T,
 7210 SAS-K 3SFP+ 8C)

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

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 will fail and an error message returned.

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



Note:

For filter support available on different 7210 SAS platforms, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C 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

dot1p-inner

Syntax

dot1p-inner [**use-rcvd-outer-dot1p** | **use-rcvd-inner-dot1p**]

no dot1p-inner

Context

config>service>epipe>sap>egress

config>service>vpls>sap>egress

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command defines the dot1p marking values used per SAP on egress for the inner tag when the SAP encapsulation is QinQ (that is, Q1.Q2 SAP).

This command takes effect only if remarking is enabled in the remark policy associated with this SAP (under the egress context). It overrides the marking values defined in the remark policy associated with this SAP, if any.

The following table describes the dot1p values extracted from the packet on SAP (ingress) when **dot1p-inner** and **dot1p-outer** CLI commands are configured.

Table 27: Dot1p values extracted from the packet on SAP (ingress) when dot1p-inner and dot1p-outer commands are configured

Ingress SAP type	Dot1p value extracted from the packet	
	rcvd-inner-dot1p is configured	use-rcvd-outer-dot1p is configured
Null SAP	None	None
Dot1q SAP	Dot1p from the outermost VLAN tag	Dot1p from the outermost VLAN tag

Ingress SAP type	Dot1p value extracted from the packet	
	rcvd-inner-dot1p is configured	use-rcvd-outer-dot1p is configured
Dot1q Default SAP (that is, * SAP)	None	None
Dot1q Explicit null SAP (that is, :0 SAP)	Dot1p from the outermost VLAN tag (if priority-tagged packet), else none.	Dot1p from the outermost VLAN tag (if priority-tagged packet), else none.
Dot1q range SAP	Dot1p from the outermost VLAN tag	Dot1p from the outermost VLAN tag
Q1.Q2 SAP	Dot1p of the inner tag	Dot1p from the outermost VLAN tag
Q1.* SAP	Dot1p from the outermost VLAN tag	Dot1p from the outermost VLAN tag
0.* SAP	Dot1p from the outermost VLAN tag	Dot1p from the outermost VLAN tag
Q1.0 SAP	Dot1p of the inner priority tag, if available, else from outermost VLAN tag	Dot1p from the outermost VLAN tag
QinQ Default SAP (that is, .*. SAP)	None	None

The following table describes the dot1p values marked in the packet on SAP egress when dot1p inner and dot1p-outer CLI commands are configured.

Table 28: Dot1p values marked in the packet on SAP egress when dot1p-inner and dot1p-outer commands configured

Egress SAP Type	Dot1p-inner dot1p value marked in the packet sent out from this SAP		Dot1p-outer dot1p value marked in the packet sent out from this SAP	
	use-rcvd-inner-dot1p is configured	use-rcvd-outer-dot1p is configured	use-rcvd-inner-dot1p is configured	use-rcvd-outer-dot1p is configured
Null SAP	NA	NA	NA	NA
Dot1q SAP	NA	NA	Dot1p bits in the outermost tag	Dot1p bits in the outermost tag
Dot1q Default SAP (that is, * SAP)	NA	NA	NA	NA

Egress SAP Type	Dot1p-inner dot1p value marked in the packet sent out from this SAP		Dot1p-outer dot1p value marked in the packet sent out from this SAP	
	use-rcvd-inner-dot1p is configured	use-rcvd-outer-dot1p is configured	use-rcvd-inner-dot1p is configured	use-rcvd-outer-dot1p is configured
Dot1q Explicit null SAP (that is, :0 SAP)	NA	NA	NA	NA
Dot1q range SAP	NA	NA	NA	NA
Q1.Q2 SAP	Dot1p bits from the inner tag	Dot1p bits from the outermost tag	Dot1p bits from the inner tag	Dot1p bits from the outermost tag
Q1.* SAP	NA	NA	Dot1p bits from the inner tag	Dot1p bits from the outermost tag
0.* SAP	NA	NA	NA	NA
Q1.0 SAP	NA	NA	Dot1p bits from the inner tag	Dot1p bits from the outermost tag
QinQ Default SAP (that is, *.* SAP)	NA	NA	NA	NA



Note:

The NA entry in the preceding table means egress encapsulation is not done, and neither the remark policy nor the **use-rcvd** command will be applicable at that level.

When the **no** form of this command is executed, the values defined in the remark policy associated with this SAP are used, if any. If no remark policy is associated with SAP egress, the default values are used.

Default

no dot1p-inner

Parameters

use-rcvd-inner-dot1p

For information about this option, see [Table 27: Dot1p values extracted from the packet on SAP \(ingress\) when dot1p-inner and dot1p-outer commands are configured](#) and [Table 28: Dot1p values marked in the packet on SAP egress when dot1p-inner and dot1p-outer commands configured](#).

use-rcvd-outer-dot1p

For information about this option, see [Table 27: Dot1p values extracted from the packet on SAP \(ingress\) when dot1p-inner and dot1p-outer commands are configured](#) and [Table 28: Dot1p values marked in the packet on SAP egress when dot1p-inner and dot1p-outer commands configured](#).

dot1p-outer

Syntax

```
dot1p-outer [use-rcvd-outer-dot1p | use-rcvd-inner-dot1p]
no dot1p-outer
```

Context

```
config>service>epipe>sap>egress
config>service>vpls>sap>egress
config>service>ies>if>sap>egress
config>service>vprn>if>sap>egress
```

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C; IES and VPRN contexts apply only to the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command defines the dot1p marking values to be used per SAP on egress for the outer tag when the SAP encapsulation is QinQ or dot1q. The command takes effect only if remarking is enabled in the remark policy associated with this SAP (under the egress context). It overrides the marking values defined in the remark policy associated with this SAP, if any. For more information, see [Table 27: Dot1p values extracted from the packet on SAP \(ingress\) when dot1p-inner and dot1p-outer commands are configured](#) and [Table 28: Dot1p values marked in the packet on SAP egress when dot1p-inner and dot1p-outer commands configured](#).

When the **no** form of this command is executed, the values defined in the remark policy associated with this SAP are used, if any. If no remark policy is associated with SAP egress, the default values are used.

Default

```
no dot1p-outer
```

Parameters

use-rcvd-inner-dot1p

For information about this option, see [Table 27: Dot1p values extracted from the packet on SAP \(ingress\) when dot1p-inner and dot1p-outer commands are configured](#) and [Table 28: Dot1p values marked in the packet on SAP egress when dot1p-inner and dot1p-outer commands configured](#).

user-rcvd-outer-dot1p

For information about this option, see [Table 27: Dot1p values extracted from the packet on SAP \(ingress\) when dot1p-inner and dot1p-outer commands are configured](#) and [Table](#)

28: Dot1p values marked in the packet on SAP egress when dot1p-inner and dot1p-outer commands configured.

meter-override

Syntax

[no] **meter-override**

Context

config>service>epipe>sap>ingress

config>service>vpls>sap>ingress

config>service>vprn>if>sap>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document.

Description



Note:

The **config>service>vprn>if>sap>ingress** context is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

This command, within the SAP ingress contexts, creates 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 removes 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>vprn>if>sap>ingress>meter-override

Platforms

Supported on all 7210 SAS platforms as described in this document.

Description



Note:

The **config>service>vprn>if>sap>ingress>meter-override** context is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

This command, within the SAP ingress contexts, creates 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 removes any existing overrides for the specified meter ID.

Parameters

meter-id

Required when executing the meter command within the meter-overrides context. The specified meter-id must exist within the sap-ingress QoS policy applied to the SAP. If the meter is not currently used by any forwarding class or forwarding type mappings, the meter will not actually exist on the SAP. This does not preclude creating an override context for the meter-id.

create

Keyword that is required when a **meter** 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>vprn>if>sap>ingress>meter-override>meter

Platforms

Supported on all 7210 SAS platforms as described in this document.

Description



Note:

The **config>service>vprn>if>sap>ingress>meter-override>meter** context is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

This command overrides 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

Keyword that specifies the constraints enforced when adapting the PIR rate defined within the **meter-override meter** 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 "trcm2", this parameter is interpreted as the EIR value. See the *7210 SAS-D, Dxp Quality of Service Guide* and *7210 SAS-K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Quality of Service Guide* for a description and relevant notes about meter modes.

cir

Keyword that specifies the constraints enforced when adapting the CIR rate defined within the **meter-override meter** 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** — Keyword that is mutually exclusive with the **min** and **closest** options. When **max** (maximum) is defined, the operational PIR for the meter is equal to or less than the administrative rate specified using the **meter-override** command.
 - min** — Keyword that is mutually exclusive with the **max** and **closest** options. When **min** (minimum) is defined, the operational PIR for the queue is equal to or greater than the administrative rate specified using the **meter-override** command.
 - closest** — Keyword that is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the meter is the rate closest that which is 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>vprn>if>sap>ingress>meter-override>meter
```

Platforms

Supported on all 7210 SAS platforms as described in this document.

Description



Note:

The **config>service>vprn>if>sap>ingress>meter-override>meter** context is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

This command overrides 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, the packets are marked as in-profile by the meter to indicate that the traffic is complying with meter-configured parameters.

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

Default

32 kbits

Parameters

size

Specifies the value in kbits, bytes, or kilobytes.

Values kbits:

- 4 to 16384, default (7210 SAS-D)
- 4 to 2146959, default (7210 SAS-Dxp)
- 1 to 16384, default (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)

bytes:

- 512 to 2097152, default (7210 SAS-D)
- 512 to 274810752, default (7210 SAS-Dxp)
- 64 to 2097152, default (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)

kbytes:

- 1 to 2048, default (7210 SAS-D, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)
- 1 to 268369, default (7210 SAS-Dxp)

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>vprn>if>sap>ingress>meter-override>meter

Platforms

Supported on all 7210 SAS platforms as described in this document.

Description



Note:

The **config>service>vprn>if>sap>ingress>meter-override>meter** context is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

This command overrides the default MBS for the meter. The maximum burst *size* parameter specifies the maximum burst size that can be transmitted by the source while still complying with the PIR. If the transmitted burst is lower than the MBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying with meter-configured parameters.

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

Default

32Kbits

Parameters

size

Specifies the value in kbits, bytes, or kilobytes.

Values

kbits:

- 4 to 16384, default (7210 SAS-D)
- 4 to 2146959, default (7210 SAS-Dxp)
- 1 to 16384, default (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)

bytes:

- 512 to 2097152, default (7210 SAS-D)
- 512 to 274810752, default (7210 SAS-Dxp)

- 64 to 2097152, default (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)

kbytes:

- 1 to 2048, default (7210 SAS-D, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)
- 1 to 268369, default (7210 SAS-Dxp)

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>vprn>if>sap>ingress>meter-override>meter

Platforms

Supported on all 7210 SAS platforms as described in this document.

Description



Note:

The **config>service>vprn>if>sap>ingress>meter-override>meter** context is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

This command overrides the SAP-ingress QoS policy configured mode parameters for the specified meter ID.

The **no** form of this command reverts the policy defined metering and profiling mode to a meter.

Parameters

mode

Specifies the rate mode of the meter-override.

Values trtcm1, trtcm2, srtcm (7210 SAS-D, 7210 SAS-Dxp)
trtcm2, srtcm (7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)

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>vprn>if>sap>ingress>meter-override>meter

Platforms

Supported on all 7210 SAS platforms as described in this document.

Description



Note:

The **config>service>vprn>if>sap>ingress>meter-override>meter** context is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

This command within the SAP ingress meter-overrides contexts overrides the SAP-ingress QoS policy configured rate parameters for the specified meter ID.

The **no** form of this command reverts 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. See the *7210 SAS-D, Dxp Quality of Service Guide* and *7210 SAS-K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Quality of Service Guide* for more information.

The actual PIR rate is dependent on the queue **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 0 to 4000000, max (7210 SAS-D)

1 to 3000000, max (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
1 to 20000000, max (7210 SAS-Dxp, 7210 SAS-K 3SFP+ 8C)

Default max

cir-rate

Specifies to override 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 to 4000000, max (7210 SAS-D)
0 to 3000000, max (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
0 to 20000000, max (7210 SAS-Dxp, 7210 SAS-K 3SFP+ 8C)

Default 0

qos

Syntax

qos *policy-id*

no qos

Context

config>service>epipe>sap>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command associates a QoS policy with an ingress 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 or IP interface. If the *policy-id* does not exist, an error is returned.

The **qos** command is used to associate ingress policies. The **qos** command only allows ingress policies to be associated on SAP ingress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a specific type returns an error.

By default, if no specific QoS policy is associated with the SAP for ingress, 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 policy ID to associate with SAP on ingress. The policy ID must already exist.

Values 1 to 65535

statistics

Syntax

statistics

Context

config>service>epipe>sap

Platforms

7210 SAS-D

Description

Commands in this context configure the counters associated with SAP ingress and egress.

egress

Syntax

egress

Context

config>service>epipe>sap>statistics

Platforms

7210 SAS-D

Description

Commands in this context configure the egress SAP statistics counter and set the mode of the counter. This counter counts the number of packets forwarded through the SAP.

ingress

Syntax

ingress

Context

config>service>epipe>sap>statistics

Platforms

7210 SAS-D and 7210 SAS-Dxp

Description

Commands in this context configure the ingress SAP statistics counter.

The ingress counters are not enabled by default for access-uplink SAPs. For access SAPs, if the ingress counter is enabled by default, it can be disabled.

The types of ingress SAP counters are the following:

- a counter that counts the total packets or octets received on the SAP. This counter is supported only on the 7210 SAS-D.
- a counter associated with meters defined in the QoS policy of the SAP. This counter counts the in-profile and out-of-profile packets or octets received on the SAP.

forwarded-count

Syntax

[no] **forwarded-count**

Context

config>service>epipe>sap>statistics>egress

Platforms

7210 SAS-D

Description

This command associates a counter with the SAP. The counter counts the number of packets forwarded through the SAP.

A limited number of such counters are available for use with access SAPs and access-uplink SAPs.

Use this command before enabling applicable accounting record collection on the SAP to associate a counter with the SAP.

The **no** form of this command disables the packet count.

counter-mode

Syntax

counter-mode {in-out-profile-count | forward-drop-count}

Context

config>service>epipe>sap>statistics>ingress

Platforms

7210 SAS-D and 7210 SAS-Dxp

Description

This command sets the counter mode for counters associated with SAP ingress meters. A pair of counters is available for each meter. These counters count different events based on the configured **counter-mode** value.



Note:

If an accounting policy is associated with a SAP, the counter mode can be changed. In this case, 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 sequence of commands on the specified SAP to ensure that the correct statistics are collected when the counter mode is changed.

1. Configure the **config>service>epipe>sap no collect-stats** command to disable writing of accounting records for the SAP.
2. Change the counter mode option by configuring the **config>service>epipe>sap counter-mode {in-out-profile-count | forward-drop-count}** command.
3. Configure the **config>service>epipe>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

counter-mode 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 the other counts the dropped packets. The forwarded count is the sum of in-profile and out-of-profile packets or octets received on SAP ingress. The dropped count is the count of packets or octets dropped by the policer. A packet is determined to be in-profile or out-of-profile based on the meter rate parameters. A packet is dropped by the policer if it exceeds the configured PIR. 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 the other 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. 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. The dropped counts are obtained by subtracting the sum of in-profile count and out-of-profile count from the total SAP ingress received count and displayed.

received-count

Syntax

[no] **received-count**

Context

config>service>epipe>sap>statistics>ingress

Platforms

7210 SAS-D

Description

This command associates a counter with the SAP. It counts the number of packets and octets received on the SAP (ingress).

A limited number of such counters are available for use with access-uplink SAPs.

Use this command before enabling applicable accounting record collection on the SAP.

The **no** form of this command disables counter.

3.8.2.1.6 Spoke-SDP commands

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 before it can be associated with an Epipe or VPL service. If the **sdp** command 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 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.

Special Cases

Epipe

Only one SDP ID can be bound to an Epipe service. Because an Epipe is a point-to-point service, it can have, at most, two endpoints. The two endpoints can be one SAP and one SDP or two SAPs.

Parameters

sdp-id

Specifies the SDP identifier. Allowed values are integers for existing SDPs.

Values 1 to 17407

vc-id

Specifies the virtual circuit identifier.

Values 1 to 4294967295

no endpoint

Keyword that removes the association of a spoke-SDP with an explicit endpoint name.

endpoint *endpoint-name*

Specifies the name of the service endpoint.

adv-service-mtu

Syntax

adv-service-mtu *octets*

no adv-service-mtu

Context

config>service>epipe>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the MTU value signaled in a targeted LDP for the spoke-SDP and the value used to validate the MTU value signaled by the far-end PE. If configured, this value is used instead of the service MTU. However, the configuration does not affect the locally enforced value, which is still based on the service MTU. This command for the MTU cannot be configured on a spoke-SDP that is bound to an SDP with the **config>service>sdp>adv-mtu-override** command.

When unconfigured, an adjusted service MTU is used. See the **config>service>epipe>service-mtu** command for more information.

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

Default

no adv-service-mtu

Parameters

octets

Specifies the MTU size, in octets, expressed as a decimal integer.

Values 0 to 9212

control-word

Syntax

[no] control-word

Context

config>service>epipe>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command adds 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, as 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 comes up only 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, according to Section 6.1 of RFC 4447. When the user also enables the control the remote peer, the remote peer withdraws its original label and sends a label mapping with the C-bit set to 1; the VLL service is then up in both nodes.

hash-label

Syntax

hash-label [signal-capability]

no hash-label

Context

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the hash label on VLL or VPLS services that are bound to RSVP SDP, 3107 BGP SDP, segment routing, or LDP SDP, using the auto-bind mode with the **ldp**, **rsvp-te**, or **mpls** options. When this command is enabled, the ingress datapath is modified so 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. The ingress datapath appends the hash label at the bottom of the stack (BoS) and sets the S-bit to one (1).



Note:

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

Packets generated in the CPM that are forwarded with a label 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 0.

Signaling of the hash label capability is enabled by adding the **signal-capability** option under the VLL spoke-SDP, VPLS spoke-SDP or mesh SDP interface, or PW template instance. In this case, the decision of the local PE to insert the hash label on the user and control plane packets is determined by the outcome of the signaling process and can override the local PE configuration. The following process flow applies when the **hash-label** and **signal-capability** options 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 the specific spoke-SDP or mesh SDP.
- If a 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 the user and control plane packets that it forwards on the spoke-SDP or mesh SDP. The local PE also drops user and control plane packets received from a remote PE if they include a hash label. The dropped packets may be caused by the following:
 - a remote 7210 SAS PE that does not support the **hash-label** command
 - a remote 7210 SAS PE that has the **hash-label** command enabled but does not support the **signal-capability** option
 - a remote 7210 SAS PE that supports the **hash-label** command and the **signal-capability** option, but the user did not enable them because of a misconfiguration
- If the remote PE sends a flow label sub-TLV in the PW ID FEC element with T=TRUE and R=TRUE, the local PE enables the hash label capability. Consequently, the local PE node inserts a hash label in the user and control plane packets that it forwards on the spoke-SDP or mesh SDP. The local PE node also accepts user and control plane packets from the remote PE with a hash label. The local PE node drops user and control plane packets from the remote PE without a hash label.

If the **hash-label** command is enabled on the local PE with the **signal-capability** option configured and on the remote PE without the **signal-capability** option configured on the spoke-SDP or mesh-SDP, the hash

label is included in the pseudowire packets received by the local PE. These packets must be dropped. To resolve this situation, you must disable the **signal-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 is not enabled on the local configuration of the spoke-SDP or mesh-SDP at the remote PE, the hash label is not included in the pseudowire received by the local PE.

If the **signal-capability** option is enabled or disabled in the CLI, 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 mLDP LSP in both the base router instance and in the multicast VPN (mVPN) instance.
- In 7750 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 to 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. 7210 SAS devices do not set the MSB in the hash label value for service traffic. Therefore, the user must ensure that both ends are correctly configured to either process hash labels or disable them. The MSB bit is set for MPLS/OAM traffic on 7210 SAS devices.
- The **cpe-ping**, **mac-ping**, and **svc-ping** commands are not supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C when the **hash-label** command is enabled.

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

Default

no hash-label

Parameters

signal-capability

Keyword that specifies to enable the signaling and negotiation of the use of the hash label between the local and remote PE nodes.

precedence

Syntax

precedence [*precedence-value* | **primary**]

no precedence

Context

no precedence >service>epipe>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 be assigned to only 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 reverts the precedence to the default value.

Default

4

Parameters

precedence-value

Specifies the spoke-SDP precedence.

Values 1 to 4

primary

Specifies to make this the primary spoke-SDP.

pw-status-signaling

Syntax

[no] **pw-status-signaling**

Context

config>service>epipe>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the egress VC label.

Parameters

vc-label

Specifies 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the ingress VC label.

Parameters

vc-label

Specifies a VC ingress value that indicates a specific connection.

Values 2048 to 18431

suppress-vc-vlan

Syntax

suppress-vc-vlan
no suppress-vc-vlan

Context

config>service>epipe>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command suppresses the processing of the **vlan-vc-tag** value in the dataplane. When this command is enabled, the dataplane does not add a VLAN tag with a value of 0 or the configured **vlan-vc-tag** value when forwarding the packet out of the PW. When this command is enabled, the dataplane does not remove the outermost VLAN tag in the encapsulated packet received on the PW.

The **no** form of this command disables the suppression.

Default

no suppress-vc-vlan

vlan-vc-tag

Syntax

vlan-vc-tag *vlan-id*
no vlan-vc-tag [*vlan-id*]

Context

config>service>epipe>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

vlan-id

Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

Values 0 to 4094

spoke-sdp-fec

Syntax

spoke-sdp-fec

spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**]

spoke-sdp-fec *spoke-sdp-fec-id* **no-endpoint**

spoke-sdp-fec *spoke-sdp-fec-id* [**fec** *fec-type*] [**aii-type** *aii-type*] [**create**] **endpoint** *name*

Context

config>service>epipe

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 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.

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 **config>service>sdp** context that reaches the first hop 7210 of the MS-PW. The 7210 associates an SDP with a service. If an SDP to that service is not already configured, an error message is generated. If the SDP ID exists, a binding between that SDP ID and the service is created.

This command differs from the **spoke-sdp** command because 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 identifying 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 (All) type used to identify the MS-PW endpoints.

Values 1 to 2

endpoint endpoint-name

Specifies the name of the service endpoint.

no endpoint

Keyword that adds or removes a spoke-SDP association.

auto-config

Syntax

[no] **auto-config**

Context

config>service>epipe>spoke-sdp-fec

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 contexts.

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 SAS initiates MS-PW signaling based on the

prefix values configured in the SAll and TAll of the spoke-SDP. If the SAll has the greater prefix value, the 7210 SAS initiates MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAll has the greater value prefix, the 7210 SAS 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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* value 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.

Parameters

name

Specifies the name of the explicit path to be used, as configured under the **config>service>pw-routing** context.

precedence

Syntax

precedence *prec-value*
precedence primary
no precedence

Context

config>service>epipe>spoke-sdp-fec

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 reverts the precedence to the default value.

Default

42

Parameters

prec-value

Specifies the spoke-SDP precedence.

Values 1 to 4

primary

Keyword that 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command binds 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This optional command specifies the number of attempts that should be made 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 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the source attachment individual identifier for the spoke-SDP. This is applicable only to FEC129 All type 2.

Parameters

global-id

Specifies the global ID of this 7210 SAS T-PE. This value must correspond to one of the `global_id` values configured for a local-prefix under the **config>service>pw-routing>local-prefix** context.

Values 1 to 4294967295

prefix

Specifies the prefix on this 7210 SAS T-PE that the spoke-SDP is associated with. This value must correspond to one of the prefixes configured under the **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

signaling

Syntax

signaling *signaling*

Context

config>service>epipe>spoke-sdp-fec

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures this 7210 SAS as the active or passive T-PE for signaling this MS-PW, or to automatically select whether this T-PE is active or passive based on the prefix.

In an active role, this endpoint initiates MS-PW signaling without waiting for a T-LDP label mapping message to arrive from the far end T-PE. In a passive role, the endpoint waits for the initial label mapping message from the far end before sending a label mapping for this end of the PW. In auto mode, if the SAll has the greater prefix value, the 7210 SAS initiates MS-PW signaling without waiting for a label mapping message from the far end. However, if the TAll has the greater value prefix, the 7210 SAS assumes that the far end T-PE initiates MS-PW signaling and will wait for that label mapping message before responding with a T-LDP label mapping message for the MS-PW in the reverse direction.

The **no** form of this command means that the 7210 T-PE automatically selects the 7210 SAS that will initiate MS-PW signaling based on the prefix values configured in the SAll and TAll of the spoke-SDP.

Default

auto

Parameters

signaling

Configures this 7210 SAS as the active T-PE for signaling this MS-PW.

Values auto, master

taii-type2

Syntax

taii-type2 *global-id:prefix:ac-id*

no taii-type2

Context

config>service>epipe>spoke-sdp-fec

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 the **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 the **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

4 Ethernet Virtual Private Networks

This chapter provides information about Ethernet Virtual Private Networks (EVPN) for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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.

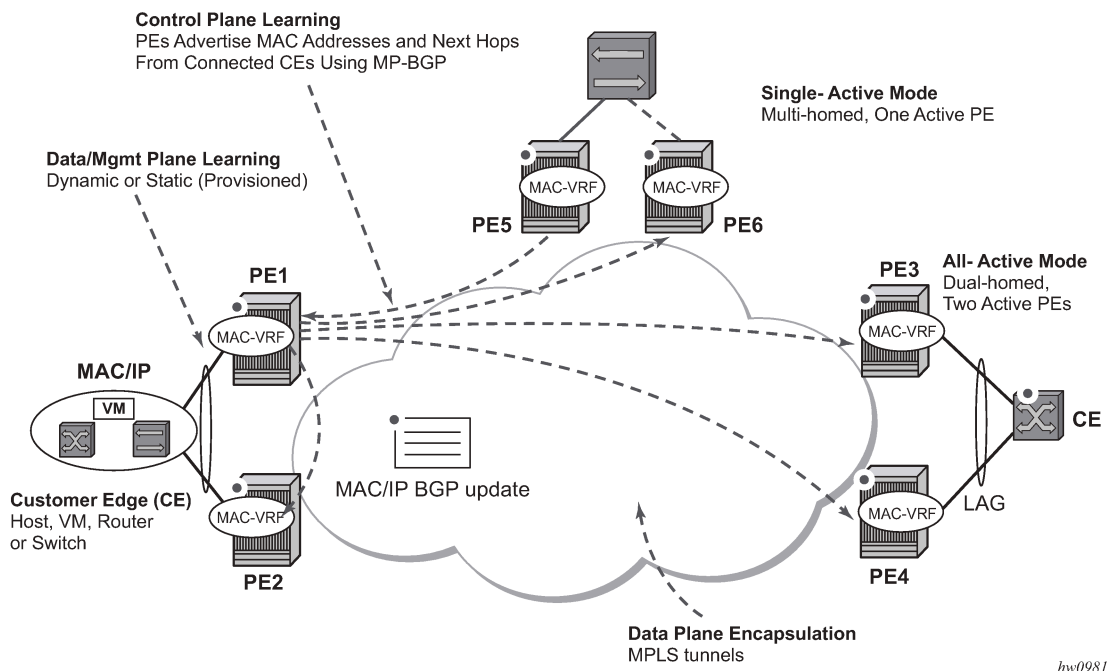
EVPN is designed to fill the gaps of traditional L2VPN technologies, such as VPLS. The main objective of EVPN is to build E-LAN services similar to IP-VPNs defined in RFC 4364, while supporting MAC learning in the control plane (distributed using multi-protocol BGP (MP-BGP)), efficient multi-destination traffic delivery, and single-active/active-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.

Figure 43: 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 all-active multi-homing (per-flow load-balancing multi-homing) and single-active multi-homing (per-service load-balancing multi-homing). Although VPLS already supports single-active multi-homing, EVPN single-active multi-homing is deemed 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 SR OS EVPN-MPLS implementation is compliant with RFC 7432.

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 EVPN routes supported in 7210 SAS SR OS and their usage in EVPN-MPLS.

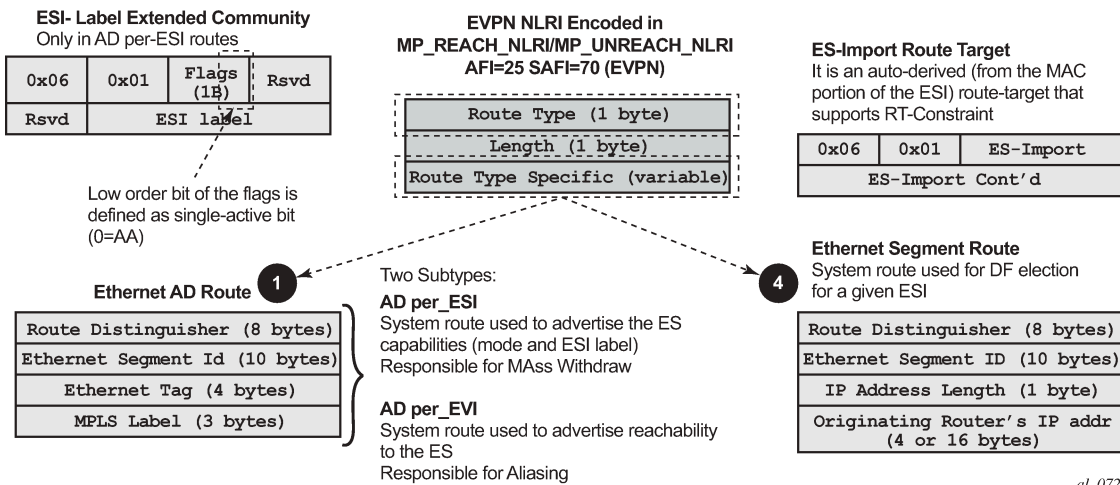
Table 29: EVPN routes and usage

EVPN route	Usage	EVPN-MPLS
Type 1 - Ethernet auto-discovery route (A-D)	Mass-withdraw, ESI labels, Aliasing	Yes
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.

Figure 44: EVPN routes type 1 and 4



4.2.1.1 EVPN route type 3 — inclusive multicast Ethernet tag route

Route type 3 is used 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 node 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.

The route type 2 generated by a router uses the following fields and values:

- **Route Distinguisher (RD)**
This is taken from the RD of the VPLS service within the BGP context. The RD can be configured or derived from the **bgp-evpn evi** value.
- **Ethernet Segment Identifier (ESI)**
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 for advertising 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 generated by a router 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.
- **Route-target extended community**
This is taken from the service level RT.

The system can send only a separate Ethernet AD per-ESI route per service.

The Ethernet AD per-EVI route generated by a router 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 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 a route-target 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. The following is a 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:**

The EVI and the system IP must be configured before executing the **configure>service/vpls>bgp-evpn>mpls>no shutdown** command.

The **evi** value is the EVPN instance (EVI) identifier and is unique in the system. It is used for the service-carving algorithm for multi-homing (if configured) and for auto-deriving route-target and route-distinguishers in EVPN-MPLS services.

If the **evi** value is not specified, the value is zero and no route-distinguisher or route-targets are auto-derived from it. If it is specified, and no other route-distinguisher or route-target are configured in the service, the following applies:

- the route-distinguisher is derived from: *system_ip:evi*
- the route-target is derived from: *autonomous-system:evi*

**Note:**

When vsi-import and vsi-export polices are configured, the route-target must be configured in the policies, and those values take precedence over the auto-derived route-targets. The operational route-target 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.

- **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.

- **ecmp**

When this command is set to a value greater than 1, aliasing is activated to the remote PEs that are defined in the same all-active multi-homing ES. See [EVPN all-active multi-homing](#) for more information.

- **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.

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.

```
*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	0	Yes	06/11/2015 19:59:03

```

192.0.2.254      ldp
                 262142      0      Yes      06/11/2015 19:59:03
                 bgp
-----
Number of entries : 7
=====

*A:Dut# show service id 1 fdb detail

=====
Forwarding Database, Service 1
=====
ServId    MAC                Source-Identifier      Type      Last Change
-----
1         00:ca:fe:ca:fe:69  eMpls:                EvpnS    06/11/15 21:53:48
                192.0.2.69:262118
1         00:ca:fe:ca:fe:70  eMpls:                EvpnS    06/11/15 19:59:57
                192.0.2.70:262140
1         00:ca:fe:ca:fe:72  eMpls:                EvpnS    06/11/15 19:59:57
                192.0.2.72:262141
-----
No. of MAC Entries: 3
-----
Legend:  L=Learned O=0am 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 allows EVPN-MPLS and VPLS to be integrated to the same network within the same service. Because EVPN is not deployed in greenfield networks, 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) and mesh-SDPs.
 - 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):
 - The following CLI command is added under the **bgp-evpn>mpls** context so that the EVPN-MPLS endpoints can be added to a split-horizon group: **bgp-evpn mpls [no] split-horizon-group group-name**.

- The **bgp-evpn mpls split-horizon-group** command must reference a user-configured split-horizon group. User-configured split-horizon groups 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.
- If the **bgp-evpn mpls split-horizon-group** command is not used, the default split-horizon group (that contains all the EVPN endpoints) is still used but cannot be referred to using SAPs/spoke-SDPs.
- The system disables the advertisement of MACs learned on spoke-SDPs or SAPs that are part of an EVPN split-horizon group:
 - When the SAPs or spoke-SDPs (manual) are configured within the same split-horizon group as the EVPN endpoints, MAC addresses are still learned on them, but they are not advertised in EVPN.
 - The preceding statement is also true if proxy-ARP/proxy-ND is enabled and an IP-to-MAC pair is learned on a SAP or SDP-binding that belongs to the EVPN split-horizon group.
 - 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 45: EVPN-VPLS integration



The following is a sample configuration 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
```

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

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 sections describes procedures implemented in SR OS to provide all-active multi-homing for a specified ES.

**Note:**

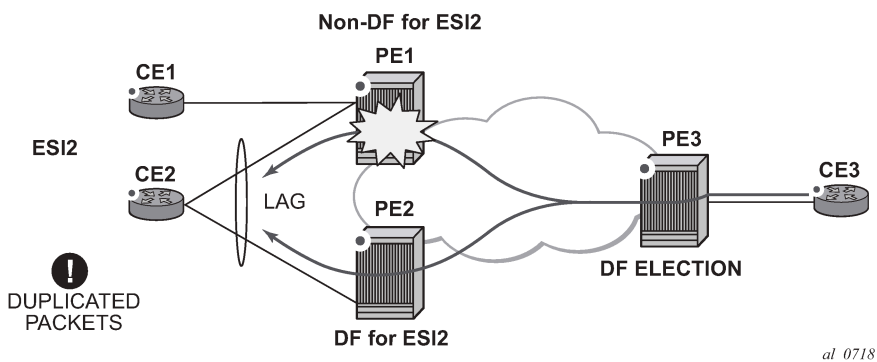
The 7210 SAS only supports two PEs per ES for all-active multi-homing.

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 46: DF election



al_0718



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 [Figure 47: Split-horizon](#)), 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 47: 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 48: Aliasing



4.2.4.2 All-active multi-homing service model

The following is a sample output of the PE1 configuration that provides all-active multi-homing to the CE2 shown in [Figure 48: Aliasing](#).

```
*A:PE1>config>lag(1)# info
-----
mode access
encap-type dot1q
port 1/1/2
lacp active administrative-key 1 system-id 00:00:00:00:00:22
no shutdown

*A:PE1>config>service>system>bgp-evpn# info
-----
route-distinguisher 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
```

In the same way, PE2 is configured as follows:

```
*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 "ESI12" 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 49: ES discovery and DF election



al_0719

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 as follows:

- 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 49: 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 49: 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      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:              EvpnS    06/11/15 00:09:06
                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: 3
-----
Legend:  L=Learned 0=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      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:              EvpnS    06/11/15 00:09:40
                192.0.2.69:262141
1        00:ca:fe:ca:fe:70  eMpls:              EvpnS    06/11/15 00:09:40
                192.0.2.70:262140
    
```

```
-----
No. of MAC Entries: 3
-----
Legend: L=Learned O=Oam 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
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  : 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 49: 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
-----
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 50: All-active multi-homing ES failure



al_0715

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 and ES routes, 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 step 2, 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 51: 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 (**StandbyforMHprotocol**). 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 due to MAC route delays

The following figure shows scenarios that may cause potential transient issues in the network.

Figure 52: 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 might 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 [Figure 52: Transient issues caused by "slow" MAC learning](#), 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 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**

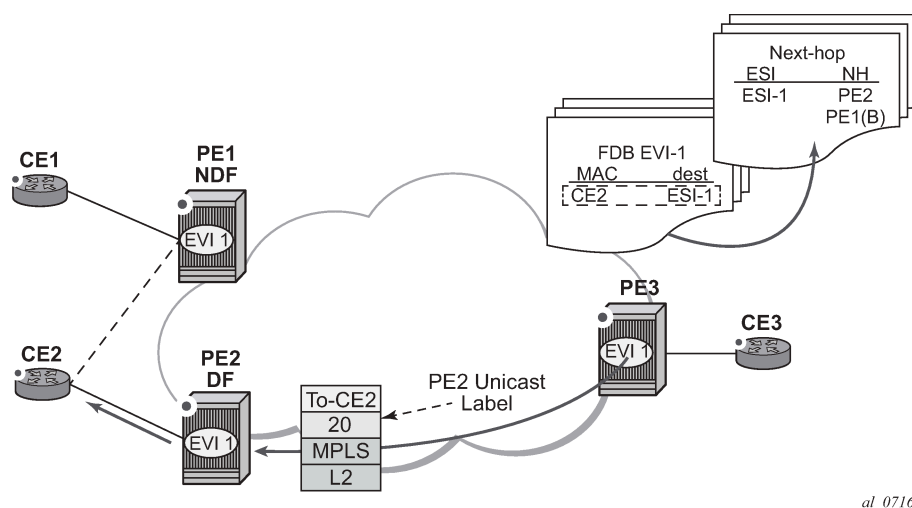
The DF election in single-active multi-homing determines the forwarding for BUM traffic from the EVPN network to the ES CE. DF election also determines the forwarding of any traffic (unicast or BUM) in any direction (to or from the CE).

- **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 53: Backup PE



al_0716

4.2.5.1 Single-active multi-homing service model

Example

The following example shows a PE1 configuration that provides single-active multi-homing to CE2, as shown in [Figure 53: Backup PE](#).

```
*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
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
```

Example

The PE2 example configuration for this scenario is as follows.

```
*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
  multi-homing single-active
  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 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, the ESI label is also used for single-active ESs.
- For single-active multi-homing, the ES can be associated with a **port** or **lag-id**, as shown in [Figure 53: 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-D# /show service system bgp-evpn ethernet-segment name "eslag1"
=====
Service Ethernet Segment
=====
Name                : eslag1
Admin State         : Enabled          Oper State           : Up
ESI                 : 00:de:01:00:00:00:00:00:01
Multi-homing        : singleActive      Oper Multi-homing    : singleActive
Lag Id              : 6
ES Activation Timer  : 3 secs (default)
Exp/Imp Route-Target : target:de:01:00:00:00:00
Svc Carving         : auto
ES SHG Label        : 130988
=====
*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 53: 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:                  Evpn  06/12/15 00:33:39
          01:74:13:00:74:13:00:00:74:13
1        00:ca:fe:ca:fe:69  eMpls:                EvpnS 06/11/15 21:53:47
          192.0.2.69:262118
1        00:ca:fe:ca:fe:70  eMpls:                EvpnS 06/11/15 19:59:57
          192.0.2.70:262140
1        00:ca:fe:ca:fe:72  eMpls:                EvpnS 06/11/15 19:59:57
          192.0.2.72:262141
-----
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      Num. MACs  Mcast    Last Change
              Transport
-----
10.20.1.2     131061         0          Yes      03/13/2000 11:26:29
              rsvp
10.20.1.2     131062         1          No       03/13/2000 12:10:04
              rsvp
10.20.1.5     131061         0          Yes      03/13/2000 11:18:23
              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  453            03/13/2000 12:10:14
-----
Number of entries: 1
-----
=====
```

```
*A:Dut-D# /show service id 1 evpn-mpls esi 00:de:01:00:00:00:00:00:01
=====
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                               Last Change
Transport
-----
```

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 54: 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.

4.3 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 that can be enabled or disabled independently per service. When enabled (proxy-ARP or proxy-ND **no shutdown**), the system populates the corresponding proxy-ARP or proxy-ND table with IP-to-MAC entries learned from the following sources:

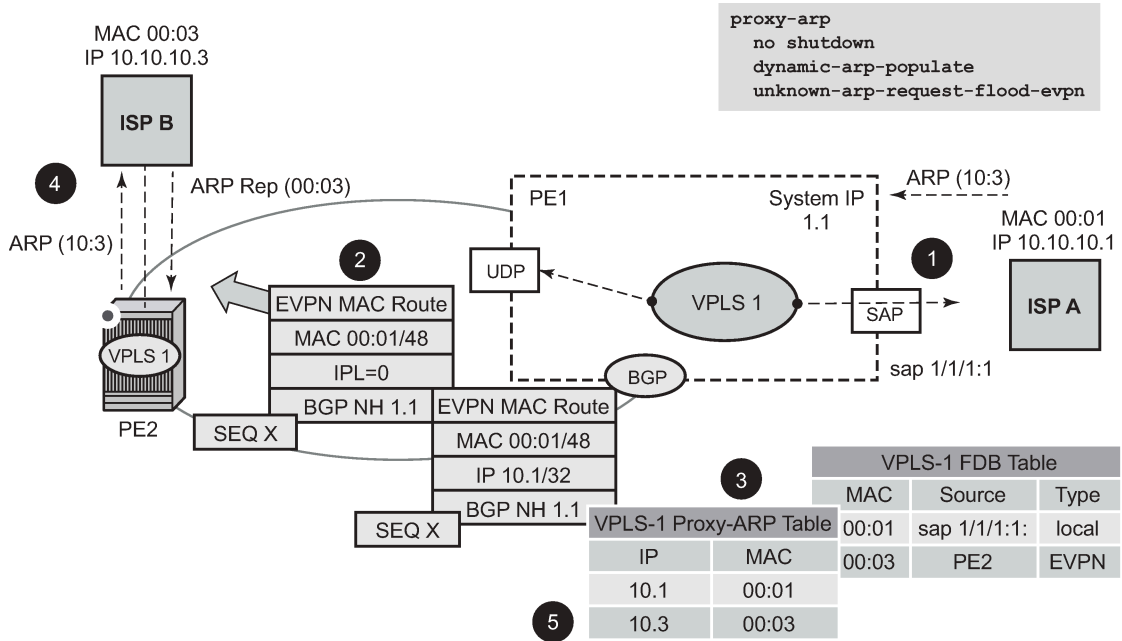
- EVPN-received IP-to-MAC entries
- user-configured static IP-to-MAC entries
- snooped dynamic IP-to-MAC entries (learned from ARP, GARP, or NA messages received on local SAPs or SDP-bindings)

In addition, any ingress ARP or ND frame on a SAP or SDP-binding are intercepted and processed. The system answers ARP requests and Neighbor Solicitation messages if the requested IP address is present in the proxy table.

Example

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 55: Proxy-ARP example usage in an EVPN network



al_0626

In the preceding figure, PE1 is configured as follows:

```
*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
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#
```

Figure 55: Proxy-ARP example usage in an EVPN network 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 = inActv** 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 30: 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 all-active 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 [Table 30: Proxy-ARP entry combinations](#).

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.

4.3.2 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: SVC MGR #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: MAC duplication parameter configuration

```
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: Conditional static MAC configuration

```
*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
Peer Router Id : 10.20.1.2
Flags        : Used Valid Best IGP
Route Source  : Internal
AS-Path      : No As-Path
EVPN type    : MAC
ESI          : 00:bc:01:00:00:00:00:00:00:01
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 56: 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 [Figure 56: Network topology of an access ring](#), 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 [Figure 56: Network topology of an access ring](#).

```

#-----
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
        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
        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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

4.4.1 EVPN-MPLS configuration examples

This section provides EVPN-MPLS configuration examples.

4.4.1.1 EVPN all-active multi-homing example

This section shows a configuration example, specified by the following assumptions:

- PE-1 and PE-2 are multi-homed to CE-12 that uses a LAG to get connected to the network. CE-12 is connected to LAG SAPs configured in an all-active multi-homing ES.
- PE-3 is a remote PE that performs aliasing for traffic destined for the CE-12

Example

The following configuration example applies to a VPLS-1 on PE-1 and PE-2, as well as the corresponding **ethernet-segment** and **lag** commands.

```
*A:Dut-C>config>lag# info
-----
mode access
encap-type dot1q
port 1/4/6
lacp active administrative-key 1 system-id 00:00:00:00:69:72
no shutdown
-----
*A:Dut-C>config>lag# /configure service system bgp-evpn
*A:Dut-D>config>service>system>bgp-evpn>eth-seg# info
-----
esi 00:de:01:00:00:00:00:00:01
service-carving
mode auto
exit
multi-homing all-active
lag 6
no shutdown
-----
description "Default tls description
split-horizon-group "vpls1" create
description "Default description for SHG vpls1"
exit
bgp-evpn
evi 1
mpls
```

```
        control-word
        split-horizon-group "vpls1"
        ingress-replication-bum-label
        auto-bind-tunnel
            resolution any
        exit
        no shutdown
    exit
exit
stp
    shutdown
exit
sap lag-6:1 create
    description "Default sap description for service id 1"
    egress
    exit
exit
static-mac
    mac 00:81:00:00:00:02 create sap lag-6:1 monitor fwd-status
exit
no shutdown
*A:Dut-C>config>service>vpls#
A:Dut-B>config>lag# info
*A:Dut-E>config>service>system>bgp-evpn>eth-seg# info
-----
        esi 00:de:01:00:00:00:00:00:01
        service-carving
            mode auto
        exit
        multi-homing all-active
        lag 6
        no shutdown
-----
*A:Dut-B>config>service>system>bgp-evpn# /configure service vpls 1
*A:Dut-E>config>service>vpls# info
-----
        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
            split-horizon-group "vpls1"
            ingress-replication-bum-label
            auto-bind-tunnel
                resolution any
            exit
            no shutdown
        exit
    exit
    stp
        shutdown
    exit
    sap lag-6:1 create
        description "Default sap description for service id 1"
        egress
        exit
    exit
    no shutdown
```

Example

The configuration on the remote PE (for example, PE-3), which supports aliasing to PE-1 and PE-2, is shown below. PE-3 does not have any ES configured and to perform aliasing requires only the VPLS-1 configuration and `ecmp>1`.

```
*A:PE3>config>service>vpls# info
-----
      bgp
      exit
      bgp-evpn
      evi 1
      mpls
      ingress-replication-bum-label
      ecmp 4
      auto-bind-tunnel
      resolution any
      exit
      no shutdown
    exit
  exit
  stp
  shutdown
  exit
  sap 1/1/1:1 create
  exit
  spoke-sdp 4:13 create
  no shutdown
  exit
  no shutdown
-----
```

4.4.1.2 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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

4.5.1 Command hierarchies

- [EVPN configuration commands](#)
- [EVPN show commands](#)
- [EVPN clear commands](#)
- [EVPN tools commands](#)

4.5.1.1 EVPN configuration commands

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls]
- no vpls service-id
- [no] bgp [bgp-instance]
- route-distinguisher [ip-addr:comm-val | as-number:ext-comm-val ]
- no route-distinguisher
- route-target ext-community
- route-target export ext-community [import ext-community]
- route-target import ext-community
- no route-target

```

```

- vsi-export policy-name [policy-name...(up to 5 max)]
- no vsi-export
- vsi-import policy-name [policy-name...(up to 5 max)]
- no vsi-import
- 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
- ecmp max-ecmp-routes
- [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] [virtual]
- es-activation-timer seconds
    
```

```
- no es-activation-timer
- esi esi
- no esi
- lag lag-id
- no lag
- multi-homing single-active [no-esi-label]
- multi-homing all-active
- no multi-homing
- port port-id
- no port
- 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
```

```
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
 - usage
- proxy-nd
 - usage

4.5.2 Command descriptions

4.5.2.1 EVPN configuration commands

bgp-evpn

Syntax

bgp-evpn
no bgp-evpn

Context

config>service>vpls

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables BGP-EVPN in the base instance.

The **no** form of this command disables BGP-EVPN.



Note:

CFM is not supported with 7210 SAS EVPN VPLS services.

evi

Syntax

evi value
no evi

Context

config>service>vpls>bgp-evpn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies a 2-byte EVPN instance that is unique in the system. It is used by the service-carving algorithm for multihoming and auto-deriving route target and route distinguishers.

If not specified, the value is zero and no route distinguisher or route targets are auto-derived from it. If the **evi value** is specified and no other route distinguisher or route target is configured in the service, the following rules apply:

- the route distinguisher is derived from **<system_ip>:evi**
- the route-target is derived from **<autonomous-system>:evi**



Note:

If VSI import and export policies are configured, the route target must be configured in the policies, and those values take precedence over the auto-derived route targets. The operational route target 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

mpls

Syntax

mpls

Context

config>service>vpls>bgp-evpn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Specifies to disable the automatic binding of a BGP-EVPN MPLS service to tunnels to MP-BGP peers.

any

Specifies to enable the binding to any supported tunnel type in a BGP-EVPN MPLS context following TTM preference.

filter

Specifies to enable 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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the BGP tunnel type.

BGP EVPN will 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 will not be 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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the LDP tunnel type.

BGP will search 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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the RSVP-TE tunnel type.

BGP will search for the best metric RSVP LSP to the address of the BGP next hop. This address can correspond to the system interface or to another loopback used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id.

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the Segment Routing (SR) tunnel type programmed by an IS-IS 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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies 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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 like 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. 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

Syntax

bgp *bgp-instance*

Context

config>service>vpls

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure the BGP related parameters for BGP EVPN.

route-distinguisher

Syntax

route-distinguisher [*ip-addr:comm-val* | *as-number:ext-comm-val*]

no route-distinguisher

Context

config>service>vpls>bgp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the Route Distinguisher (RD) component that will be signaled in the MP-BGP NLRI for L2VPN and EVPN families. This value is used for BGP-AD 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 BGP node - RD=VPLS-ID
- 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) will be automatically generated.

The **no** form of this command removes the RD component.

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 AS number.

Values *as-number*: 1 to 65535
ext-comm-val: 0 to 4294967295

route-target

Syntax

route-target *ext-community*
route-target export *ext-community*
route-target import *ext-community*
no route-target

Context

config>service>vpls>bgp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the route target (RT) component that will be signaled in the related MP-BGP attribute to be used for BGP auto-discovery and EVPN, if these features are configured in this VPLS service.

If this command is not used, the RT is built automatically using the VPLS ID. The *ext-comm* value 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 services, the route target can also be auto-derived from the **evi** value (**config service vpls bgp-evpn evi**), if this command is not configured. See the command description for more information.

The **no** form of this command removes the RT component.

Parameters

export *ext-community*

Specifies communities allowed to be sent to remote PE neighbors.

import *ext-community*

Specifies communities allowed to be accepted from remote PE neighbors.

vsi-export

Syntax

vsi-export *policy-name* [*policy-name* ... (up to 5 max)]

no vsi-export

Context

config>service>vpls>bgp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the name of the VSI export policies to be used for BGP auto-discovery, if it is configured in this VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

The **no** form of this command removes the VSI export policy.

Parameters

policy-name

Specifies a VSI export policy, 32 characters maximum.

vsi-import

Syntax

vsi-import *policy-name* [*policy-name* ... (up to 5 max)]

no vsi-import

Context

config>service>vpls>bgp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the name of the VSI import policies to be used for BGP auto-discovery, if it is configured in this VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The policy name list is handled by the SNMP agent as a single entity.

The **no** form of this command removes the VSI import policy.

Parameters

policy-name

Specifies a VSI import policy, 32 characters maximum.

mac-advertisement

Syntax

[no] **mac-advertisement**

Context

config>service>vpls>bgp-evpn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the timer after which the MAC in hold-down state is automatically flushed and the MAC duplication process starts again. This value is expected to be equal to 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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

ecmp

Syntax

ecmp *max-ecmp-routes*

no ecmp

Context

```
config>service>vpls>bgp-evpn>mpls
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

When configured in a VPLS service, this command controls the number of paths to reach a specified MAC address when that MAC in the FDB is associated with a remote all-active multi-homed ES.

The configuration of two or more ECMP paths to a specified MAC enables the aliasing function described in RFC 7432.

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

Default

no ecmp

Parameters

max-ecmp-routes

Specifies the maximum number of tunnels that may be used as ECMP next hops for the service.

Values 0 to 4

Default 0

force-vlan-vc-forwarding

Syntax

```
[no] force-vlan-vc-forwarding
```

Context

```
config>service>vpls>bgp-evpn>mpls
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command allows the system to preserve the VLAN ID and 802.1p bits of the service-delimiting qtag in a new tag added in the customer frame before sending it to the EVPN-MPLS destinations.

This command may be used in conjunction with the **sap ingress vlan-translation** command. In this case, the configured translated VLAN ID is the VLAN ID sent to the EVPN-MPLS destinations as opposed to the service-delimiting tag VLAN ID. If the ingress SAP/SDP binding is null-encapsulated, the output VLAN ID and pbits are zero.

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

Default

no force-vlan-vc-forwarding

ingress-replication-bum-label

Syntax

[no] **ingress-replication-bum-label**

Context

config>service>vpls>bgp-evpn>mpls

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 will advertise two labels per EVPN VPLS instance, one for unicast and one for BUM traffic. The ingress PE will use 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 specific 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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the proxy-ARP in a VPLS service.

The **no** form of this command removes the proxy-ARP context.

Default

no proxy-arp

proxy-nd

Syntax

[no] proxy-nd

Context

config>service>vpls

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the proxy-ND in a VPLS service.

The **no** form of this command removes the proxy-ND context.

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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/proxy-ND table. The proxy-ARP/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/ND frames and the dynamic/EVPN dup proxy-ARP/proxy-ND table entries are flushed. All the static entries are kept in the table as "inactive", regardless of their previous "Status".

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 multi-homing) or added to the default-multicast-list (in case of all-active multi-homing).

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 value
```

```
no esi
```

Context

```
config>service>system>bgp-evpn>ethernet-segment
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

value

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 800

multi-homing

Syntax

multi-homing single-active [no-esi-label]

multi-homing all-active

no multi-homing

Context

```
config>service>system>bgp-evpn>ethernet-segment
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the multi-homing mode for the specified **ethernet-segment** as **single-active** or **all-active** multi-homing, as defined in RFC7432.

By default, the use of **esi-label** is enabled for **all-active** and **single-active** as defined in RFC7432 (for **single-active** multi-homing, the ESI label is used to avoid transient loops).

When **single-active no-esi-label** is specified, the system will not allocate an ESI label and hence advertise ESI label 0 to peers. Even if the ESI is configured to not send the ESI label, upon reception of an ESI label from a peer, the PE will always send traffic to that peer using the received ESI label.

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.

all-active

Specifies all-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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a port ID associated with the ES. If the **ethernet-segment** is configured as **all-active**, only a LAG can be associated to 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 **ethernet-segment**. 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 Ethernet segment 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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 route distinguisher 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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 31: Output fields: EVPN MPLS tunnel endpoints](#) describes the output fields.

Sample output

```
*A:Dut-B# show service evpn-mpls
=====
EVPN MPLS Tunnel Endpoints
=====
EvpnMplsTEP 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 31: 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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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, the 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 32: 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 32: 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
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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays the existing EVPN-MPLS destinations for a specified service and all related information. The command allows filtering based on **esi** (for EVPN multi-homing) to display the EVPN-MPLS destinations associated to 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 33: 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 33: 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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

The command allows for specific IP addresses to be displayed. Dynamic IP entries associated to 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 detailed information.

Output

The following output is an example of proxy-ARP information for a specified service, and [Table 34: 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 34: 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

Label	Description
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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 to 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 (eight 16-bit pieces)
 x:x:x:x:x:d.d.d.d
 where:
 x - [0 to FFFF]H
 d - [0 to 255]D

detail

Displays detailed information.

Output

The following output is an example of proxy-ND information for a specified service, and [Table 35: 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          Send Refresh    : disabled
Table Size      : 250              Total           : 1
Static Count    : 0                EVPN Count      : 0
Dynamic Count   : 1                Duplicate Count : 0
Dup Detect
-----
Detect Window   : 3 mins           Num Moves       : 5
Hold down      : 9 mins
Anti Spoof MAC : None
EVPN
-----
Unknown NS Flood : enabled          ND Advertise    : Router
Rtr Unsol NA Flood: disabled      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 35: 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
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

Label	Description
Last Update	Displays the date and time of the last update

system

Syntax

system

Context

show>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure the **system** BGP EVPN show command.

bgp-evpn

Syntax

bgp-evpn

Context

show>service>system

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command shows system BGP EVPN information.

Output

The following output is an example of system BGP EVPN information, and [Table 36: 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 36: 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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 37: 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 37: 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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures tools to display service dump information.

usage

Syntax

usage

Context

tools>dump>service>proxy-arp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command provides information about the usage and limit of the system-wide proxy-ARP table for all the services. The command also shows if the limit has been exceeded and a trap raised.

Output

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
```

usage

Syntax

usage

Context

tools>dump>service>proxy-nd

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 which simplifies the LAN/WAN boundary and allows for rapid and flexible service provisioning.

VPLS provides a balance between point-to-point Frame Relay service and outsourced routed services. VPLS enables each customer to maintain control of their own routing strategies. All customer routers in the VPLS service are part of the same subnet (LAN) which simplifies the IP addressing plan, especially when compared to a mesh constructed from many separate point-to-point connections. The VPLS service management is simplified 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 (considered a local service) or more (which is considered a distributed service) service routers. The connection appears to be a bridged domain to the customer sites so protocols, including routing protocols, can traverse the VPLS service.

Other VPLS advantages include:

- VPLS is a transparent, protocol-independent service.
- There is no Layer 2 protocol conversion between LAN and WAN technologies.
- There is no need to design, manage, configure, and maintain separate WAN access equipment, therefore, eliminating the need to train personnel on WAN technologies such as Frame Relay.

The 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T do not support IP/MPLS uplinks and support only local VPLS service. These devices support provisioning of access SAPs or access-uplink SAPs to connect to the provider edge IP/MPLS routers, allowing users to provision an end-to-end VPLS VPN service using the concept of hierarchical VPLS.


The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C support IP/MPLS uplinks and support both local VPLS service and distributed VPLS service. This provides an option to use VPLS service configured on the node with spoke-SDPs to connect to the provider edge IP/MPLS routers or use access SAPs or access-uplink SAPs to connect to the provider edge IP/MPLS routers, allowing users to provision an end-to-end VPLS VPN service using the concept of hierarchical VPLS. The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C also support mesh SDPs for use in smaller network configurations.

5.1.1 VPLS packet walk-through for a local service

This section provides an example of VPLS processing for a local VPLS service using QinQ access uplink SAPs to connect to the PE router. It is applicable to the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K

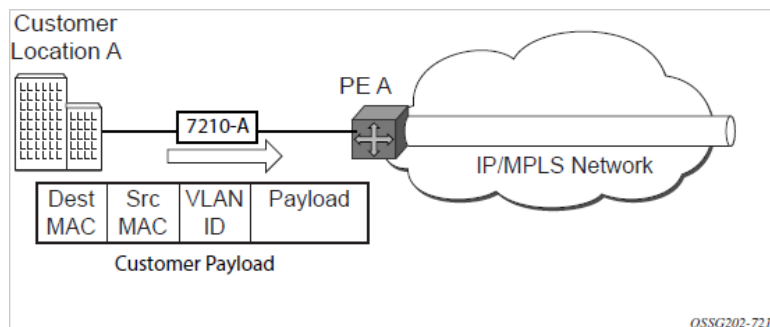
2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C when using access uplink SAPs). It describes VPLS processing of a customer packet sent across the network from site-A connected to PE-Router-A through a 7210 SAS to site-C connected through 7210 SAS to PE-Router-C (Figure 57: VPLS service architecture). This section does not describe the processing on the PE routers, but only on 7210 SAS routers.

Figure 57: VPLS service architecture



1. 7210-A (Figure 58: Access port ingress packet format and lookup)
 - 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.

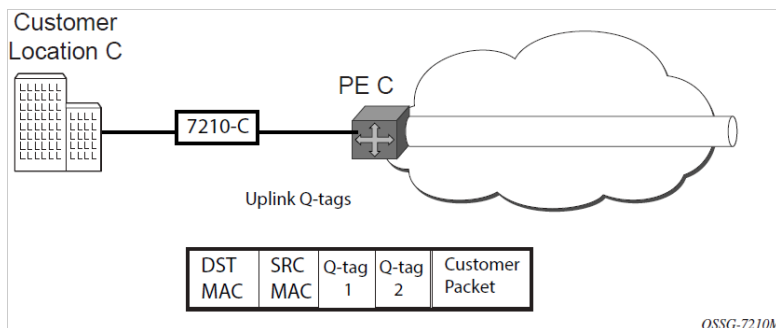
Figure 58: Access port ingress packet format and lookup



- 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).

Figure 59: Network port egress packet format and flooding



- i. For a Known MAC Address (Figure 59: Network port egress packet format and flooding):
 - If the destination MAC address has already been learned by 7210 an existing entry in the FIB table identifies destination uplink QinQ SAP used for sending the packet toward the PE-Router-A.
 - 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.
- ii. For an Unknown MAC Address (Figure 59: Network port egress packet format and flooding):
 - If the destination MAC address has not been learned, 7210 will flood the packet to all the uplink SAPs participating in the service.

2. Core Router Switching

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 58: 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).
 - i. 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.
 - ii. If the destination MAC address has not been learned, 7210 will flood the packet to all the access SAPs participating in the service.

5.1.2 VPLS packet walk-through for a distributed service


This section provides an example of VPLS processing for a local VPLS service using spoke-SDP/mesh-SDP as uplinks to connect to the PE routers. It is applicable to the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C when using spoke-SDP/mesh-SDP as uplinks to connect to the PE routers. It describes VPLS processing of a customer packet sent across the network from site-A connected to PE-Router-A through a 7210 SAS to site-C connected through 7210 SAS to PE-Router-C (Figure 60: VPLS service architecture). This section does not describe the processing on the PE routers, but only on 7210 SAS routers.

Figure 60: VPLS service architecture



1. 7210-A (Figure 61: Access port ingress packet format and lookup)
 - 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.

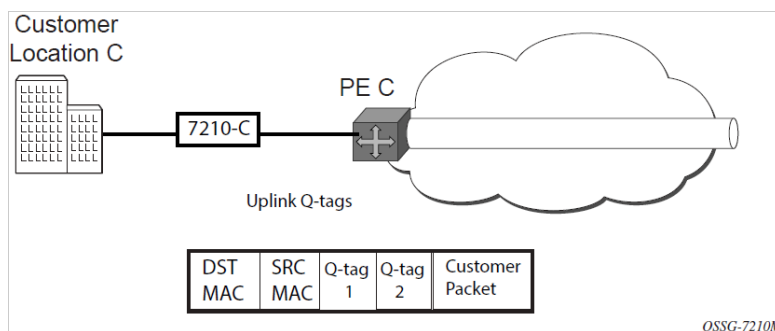
Figure 61: Access port ingress packet format and lookup



- 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 in-stance. 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).
 - i. For a Known MAC Address (Figure 61: Access port ingress packet format and lookup):
 If the destination MAC address has already been learned by 7210, an existing entry in the FIB table identifies the far-end PE-Router and the service VC-label (inner label) used before sending the packet to PE-Router-A.

 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.
 - ii. For a Unknown MAC Address (Figure 61: Access port ingress packet format and lookup):
 If the destination MAC address has not been learned, 7210 will flood the packet to all the spoke-SDPs participating in the service, as shown in the following figure.

Figure 62: Network port egress packet format and flooding



2. Core Router Switching

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 61: 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 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 in-stance. 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).
 - i. 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.
 - ii. If the destination MAC address has not been learned, 7210 will flood the packet to all the access SAPs and possible other spoke-SDPs participating in the service.

5.2 VPLS features

5.2.1 VPLS enhancements on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

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
 - aging timers for learned MAC addresses
- Implementation of Spanning Tree Protocol (STP) parameters on a per VPLS and per SAP basis.
- IGMP snooping on a per-SAP basis.

5.2.2 VPLS enhancements on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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
 - aging timers for learned MAC addresses

- Implementation of Spanning Tree Protocol (STP) parameters on a per VPLS and per SAP basis.
- IGMP snooping on a per-SAP and SDP basis.

5.2.3 VPLS over MPLS - supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The VPLS architecture proposed in draft-ietf-ppvpn-vpls-ldp-0x.txt specifies the use of provider equipment (PE) capable of learning, bridging, and replication on a per-VPLS basis. The PE routers that participate in the service are connected using MPLS Label Switched Path (LSP) tunnels in a full-mesh composed of mesh SDPs or based on an LSP hierarchy (Hierarchical VPLS (H-VPLS)) composed of mesh SDPs and spoke-SDPs. The 7210 SAS 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 provider edge (PE) routers through a mesh/spoke-SDP. The PE routers are connected using a full mesh of LSPs.
- Negotiating per-service VC labels using draft-Martini encapsulation.
- Replicating unknown and broadcast traffic in a service domain.
- Enabling MAC learning over tunnel and access ports (see [VPLS MAC learning and packet forwarding](#)).
- Using a separate forwarding information base (FIB) per VPLS service.

5.2.4 VPLS over SAPs

All 7210 SAS platforms supported in this document support QinQ SAPs or Dot1q SAPs that allows them to connect to upstream PE nodes providing IP/MPLS transport.

VPLS is provided over QinQ/Dot1q SAPs by:

- Connecting bridging-capable 7210 SAS devices.
- Replicating unknown and broadcast traffic in a service domain.
- Enabling MAC learning over QinQ/Dot1q SAPs and access ports (see [VPLS MAC learning and packet forwarding](#)).
- Using a separate forwarding information base (FIB) per VPLS service.

5.2.5 VPLS MAC learning and packet forwarding

The 7210 SAS edge devices perform the packet replication required for broadcast and multicast traffic across the bridged domain. MAC address learning is performed by the 7210 SAS device to reduce the amount of unknown destination MAC address flooding.

On the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the user can configure Layer 2 uplinks (such as Dot1q SAPs or QinQ SAPs) for uplink connectivity and access SAPs for customer connectivity as logical bridge ports in a VPLS service.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the user can configure access SAPs for customer connectivity and mesh-sdp/spoke-sdp for uplink connectivity as logical bridge ports in a VPLS service.

Each 7210 SAS maintains a Forwarding Information Base (FIB) for each VPLS service instance and learned MAC addresses are populated in the FIB table of the service. All traffic is switched based on MAC addresses using QinQ SAPs created on access uplink ports. Unknown destination packets (for example, the destination MAC address has not been learned) are forwarded on all SAPs for that service until the target station responds and the MAC address is learned by the 7210 SAS associated with that service.

5.2.5.1 IGMP snooping on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

In Layer 2 switches, multicast traffic is treated like an unknown MAC address or broadcast frame, which causes the incoming frame to be flooded out (broadcast) on every port within a VLAN. Although this is acceptable behavior for unknowns and broadcast frames, this flooded multicast traffic may result in wasted bandwidth on network segments and end stations, as IP multicast hosts can join and be interested in only specific multicast groups.

IGMP snooping entails using information in Layer 3 protocol headers of multicast control messages to determine the processing at Layer 2. By doing so, an IGMP snooping switch provides the benefit of conserving bandwidth on those segments of the network in which no node has expressed interest in receiving packets addressed to the group address.

IGMP snooping can be enabled in the context of VPLS services. The IGMP snooping feature allows for optimization of the multicast data flow to only the SAPs members of the group. The system builds a database of group members per service by listening to IGMP queries and reports from each SAP:

- When the switch receives an IGMP report from a host for a multicast group, the switch adds the host port number to the forwarding table entry.
- When it receives an IGMP leave message from a host, it removes the host port from the table entry, if no other group members are present. It also deletes entries if it does not receive periodic IGMP membership reports from the multicast clients.

The following are IGMP snooping features:

- IGMP v1, v2, and v3 are supported (RFC 1112, *Host Extensions for IP Multicasting*, and RFC 2236, *Internet Group Management Protocol, Version 2*).
- IGMP snooping can be enabled and disabled on individual VPLS service instances.
- IGMP snooping can be configured on individual SAPs that are part of a VPLS service. When IGMP snooping is enabled on a VPLS service, all its contained SAPs automatically have snooping enabled.
- Fast leave terminates the multicast session immediately, instead of using the standard group-specific query to check if other group members are present on the network.
- SAPs can be statically configured as multicast router ports. This allows the operator to control the set of ports to which IGMP membership reports are forwarded.
- Static multicast group membership on a per SAP basis can be configured.
- The maximum number of multicast groups (static and dynamic) that a SAP can join can be configured. An event is generated when the limit is reached.
- The maximum number of multicast groups (static and dynamic) that a VPLS instance simultaneously supports can be configured.

- Proxy summarization of IGMP messages reduces the number of IGMP messages processed by upstream devices in the network.
- IGMP filtering allows a subscriber to a service or the provider to block, receive, or transmit permission (or both) to individual hosts or a range of hosts.

The following types of filters can be defined:

- Filter group membership that report from a host or range of hosts. This filtering is performed by importing an appropriately-defined routing policy into the SAP.
- Filters that prevent a host from transmitting multicast streams into the network. The operator can define a data-plane filter (ACL) that drops all multicast traffic, and apply this filter to a SAP.

5.2.5.2 IGMP snooping on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

In Layer 2 switches, multicast traffic is treated like an unknown MAC address or broadcast frame, which causes the incoming frame to be flooded out (broadcast) on every port within a VLAN. Although this is acceptable behavior for unknowns and broadcast frames, this flooded multicast traffic may result in wasted bandwidth on network segments and end stations, as IP multicast hosts can join and be interested in only specific multicast groups.

IGMP snooping entails using information in Layer 3 protocol headers of multicast control messages to determine the processing at Layer 2. By doing so, an IGMP snooping switch provides the benefit of conserving bandwidth on those segments of the network in which no node has expressed interest in receiving packets addressed to the group address.

IGMP snooping can be enabled in the context of VPLS services. The IGMP snooping feature allows for optimization of the multicast data flow to only those SAPs and SDPs that are members of the group. The system builds a database of group members per service by listening to IGMP queries and reports from each SAP and SDP:

- When the switch receives an IGMP report from a host for a multicast group, the switch adds the host port number to the forwarding table entry.
- When it receives an IGMP leave message from a host, it removes the host port from the table entry, if no other group members are present. It also deletes entries if it does not receive periodic IGMP membership reports from the multicast clients.

The following are IGMP snooping features:

- IGMP v1, v2, and v3 are supported (RFC 1112, Host Extensions for IP Multicasting, and RFC 2236, Internet Group Management Protocol, Version 2). On 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, IGMP v1, v2, v3 is supported with both with L2 uplinks and MPLS uplinks.
- IGMP snooping can be enabled and disabled on individual VPLS service instances.
- IGMP snooping can be configured on individual SAPs and SDPs that are part of a VPLS service. When IGMP snooping is enabled on a VPLS service, all its contained SAPs and SDPs automatically have snooping enabled.
- Fast leave terminates the multicast session immediately, instead of using the standard group-specific query to check if other group members are present on the network.
- SDPs and SAPs can be statically configured as multicast router ports. This allows the operator to control the set of ports to which IGMP membership reports are forwarded.
- Static multicast group membership on a per SAP and SDP basis can be configured.

- The maximum number of multicast groups (static and dynamic) that a SAP and SDP can join can be configured. An event is generated when the limit is reached.
- The maximum number of multicast groups (static and dynamic) that a VPLS instance simultaneously supports can be configured.
- Proxy summarization of IGMP messages reduces the number of IGMP messages processed by upstream devices in the network.
- IGMP filtering allows a subscriber to a service or the provider to block, receive, or transmit permission (or both) to individual hosts or a range of hosts. The following types of filters can be defined:
 - Filter group membership that report from a particular host or range of hosts. This filtering is performed by importing an appropriately-defined routing policy into the SAP and SDP.
 - Filters that prevent a host from transmitting multicast streams into the network. The operator can define a data-plane filter (ACL) that drops all multicast traffic, and apply this filter to a SAP and SDP.

5.2.5.3 Multicast VLAN Registration (MVR) support

Multicast VPLS Registration (MVR) is a bandwidth optimization method for multicast in a broadband services network. MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on one or more network-wide multicast VPLS instances.

MVR assumes that subscribers join and leave multicast streams by sending IGMP join and leave messages. The IGMP leave and join message are sent inside the VPLS to which the subscriber port is assigned. The multicast VPLS is shared in the network while the subscribers remain in separate VPLS services. Using MVR, users on different VPLS cannot exchange any information between them, but still multicast services are provided.

On the MVR VPLS, IGMP snooping must be enabled. On the user VPLS, IGMP snooping and MVR work independently. If IGMP snooping and MVR are both enabled, MVR reacts only to join and leave messages from multicast groups configured under MVR. Join and leave messages from all other multicast groups are managed by IGMP snooping in the local VPLS. This way, potentially several MVR VPLS instances could be configured, each with its own set of multicast channels:

- **MVR by proxy**

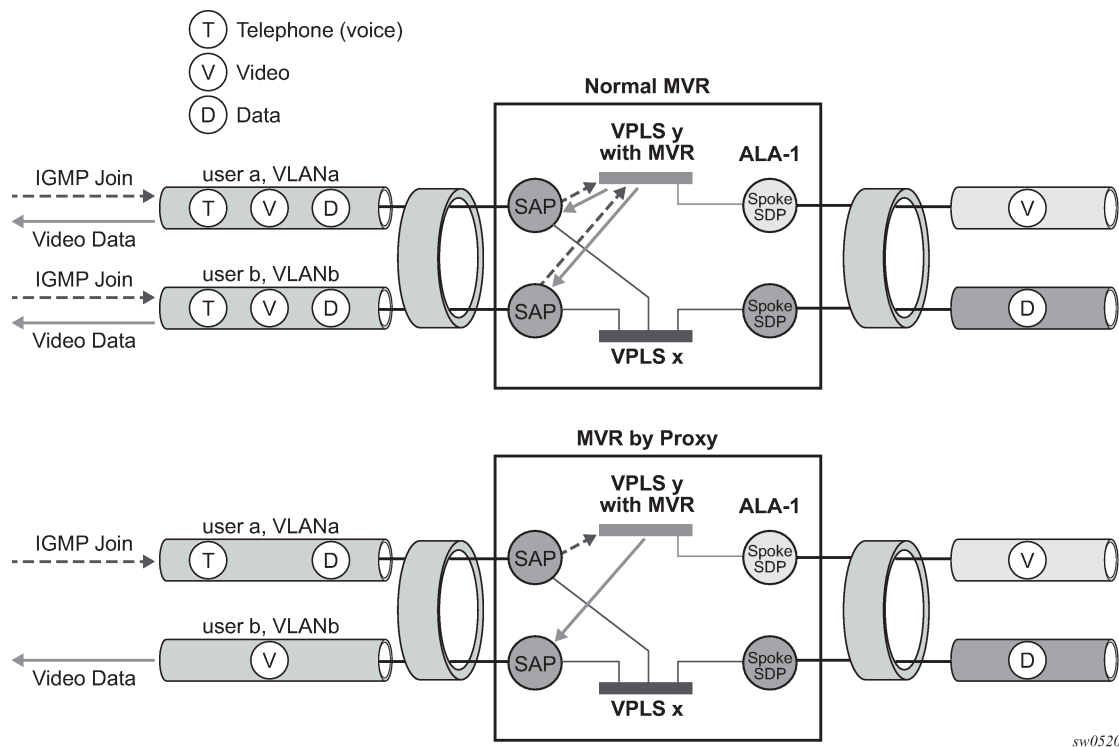
In some situations, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behavior) but to another SAP. This is called MVR by proxy.



Note:

In the following figure, the reference to SDP is only applicable to the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

Figure 63: MVR and MVR by proxy



5.2.5.3.1 Configuration guidelines for MVR

In a MVR configuration, the svc-sap-type of the VPLS service that is the source, known as 'mvr vpls service' and the svc-sap-type of the VPLS service that is the sink, known as 'user vpls service' should match.

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 L2 forwarding table management

The following sections describe VPLS features related to management of the Forwarding Information Base (FIB).

5.2.7.1 FIB size on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

The following MAC table management features are required for each instance of a SAP within a particular VPLS service instance:

- **MAC FIB size limits**

Allows users to specify the maximum number of MAC FIB entries learned locally for a SAP. If the configured limit is reached, then no new addresses will be learned from the SAP until at least one FIB entry is aged out or cleared:

- When the limit is reached on a SAP, packets with unknown source MAC addresses are still forwarded (this default behavior can be changed by configuration). By default, if the destination MAC address is known, it is forwarded based on the FIB, and if the destination MAC address is unknown, it will be flooded. Alternatively, if discard unknown is enabled at the VPLS service level, unknown destination MAC addresses are discarded.
- The log event SAP MAC limit reached is generated when the limit is reached. When the condition is cleared, the log event SAP MAC Limit Reached Condition Cleared is generated.
- Disable learning at the VPLS service level allows users to disable the dynamic learning function on the service. Disable learning is not supported at the SAP level.
- Disable aging allows users to turn off aging for learned MAC addresses on a SAP of a VPLS service instance.

5.2.7.2 FIB size on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The following MAC table management features are required for each instance of a SAP or spoke-SDP within a VPLS service instance:

- **MAC FIB size limits** - Allows users to specify the maximum number of MAC FIB entries learned locally for a SAP or remotely for a spoke-SDP. If the configured limit is reached, then no new addresses will be learned from the SAP and spoke-SDP until at least one FIB entry is aged out or cleared:
 - When the limit is reached on a SAP and spoke-SDP, packets with unknown source MAC addresses are still forwarded (this default behavior can be changed by configuration). By default, if the destination MAC address is known, it is forwarded based on the FIB, and if the destination MAC address is unknown, it will be flooded. Alternatively, if discard unknown is enabled at the VPLS service level, unknown destination MAC addresses are discarded.
 - The log event SAP/SDP MAC limit reached is generated when the limit is reached. When the condition is cleared, the log event SAP/SDP MAC Limit Reached Condition Cleared is generated.
 - Disable learning at the VPLS service level allows users to disable the dynamic learning function on the service. Disable learning is not supported at the SAP level and spoke-SDP level.
 - Disable aging allows users to turn off aging for learned MAC addresses on a SAP and spoke-SDP of a VPLS service instance.

5.2.7.3 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.

5.2.7.4 Local aging timers on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

Like a Layer 2 switch, learned MACs within a VPLS instance can be aged out if no packets are sourced from the MAC address for a specified period of time (the aging time). The age timer for the VPLS instance specifies the aging time for locally learned MAC addresses.

The aging mechanism is considered a low priority process. In most situations, the aging out of MAC addresses can happen in within tens of seconds beyond the age time. To minimize overhead, local MAC addresses on a LAG port, in some circumstances, can take up to two times their respective age timer to be aged out.

5.2.7.5 Local and remote aging timers on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Like a Layer 2 switch, learned MACs within a VPLS instance can be aged out if no packets are sourced from the MAC address for a specified time-period (the aging time). In each VPLS service instance, there are independent aging timers for locally learned MAC and remotely learned MAC entries in the forwarding database (FIB). A local MAC address is a MAC address associated with a SAP because it ingress on a SAP. A remote MAC address is a MAC address received by an SDP from another router for the VPLS instance. The local-age timer for the VPLS instance specifies the aging time for locally learned MAC addresses, and the remote-age timer specifies the aging time for remotely learned MAC addresses.

In general, the remote-age timer is set to a longer period than the local-age timer to reduce the amount of flooding required for destination unknown MAC addresses. The aging mechanism is considered a low priority process. In most situations, the aging out of MAC addresses can happen in within tens of seconds beyond the age time. To minimize overhead, local MAC addresses on a LAG port and remote MAC addresses, in some circumstances, can take up to two times their respective age timer to be aged out.

5.2.7.6 Disable MAC aging

The MAC aging timers can be disabled which will prevent any learned MAC entries from being aged out of the FIB. When aging is disabled, it is still possible to manually delete or flush learned MAC entries. Aging can be disabled for learned MAC addresses on a SAP of a VPLS service instance.

5.2.7.7 Disable MAC learning

When MAC learning is disabled for a service, new source MAC addresses are not entered in the VPLS FIB. MAC learning can be disabled for services.

5.2.7.8 Unknown MAC discard

Unknown MAC discard is a feature which discards all packets ingressing the service where the destination MAC address is not in the FIB. The normal behavior is to flood these packets to all end points in the service.

Unknown MAC discard can be used with the disable MAC learning and disable MAC aging options to create a fixed set of MAC addresses allowed to ingress and traverse the service.

5.2.7.9 VPLS and rate limiting

Traffic that is usually flooded throughout the VPLS can be rate limited on SAP ingress using service ingress QoS policies. In a service ingress QoS policy, individual meters can be defined per forwarding class to provide rate-limiting/policing of broadcast traffic, MAC multicast traffic and unknown destination MAC traffic.



Note:

On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C, the option is available to use either service ingress queues with rate shapers or meters to limit different traffic types.

5.2.7.10 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-D, 7210 SAS-Dxp, or 7210 SAS-K 2F1C2T shut down the SAP and create an alarm event when the threshold is exceeded.

When MAC move is enabled, the 7210 SAS-K 2F6C4T or 7210 SAS-K 3SFP+ 8C shut down the SAP or spoke-SDP and create an alarm event when the threshold is exceeded.

5.2.7.11 Split horizon SAP groups



Note:

- Split Horizon group per service is only supported on the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.
- Port-based split horizon groups are only supported on the 7210 SAS-D and 7210 SAS-Dxp.

In many applications, the split-horizon group concept involving a group of SAPs is useful to prevent direct customer-to-customer traffic exchange (without the traffic being sent to the head-end service nodes). This extension is referred to as a split horizon SAP group. Traffic arriving on a SAP or a spoke-SDP within a split horizon group will not be forwarded to other SAPs configured in the same split horizon group, but will be forwarded to other SAPs, which are not part of the split horizon group.

5.2.7.11.1 Split horizon SAP groups and split horizon spoke-SDP groups on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Within the context of VPLS services, a loop-free topology inside a fully meshed VPLS core is achieved by applying a split-horizon forwarding concept. The packets received from a mesh SDP are never forwarded to other mesh SDPs within the same service. The advantage of this approach is that no protocol is required to detect loops within the VPLS core network.

In applications, such as residential aggregation, it is useful to extend this split-horizon concept also to groups of SAPs or spoke-SDPs. This extension is referred to as a split horizon SAP group/spoke-SDP group. Traffic arriving on a SAP or a spoke-SDP within a split horizon group will not be forwarded to other SAPs and spoke-SDPs configured in the same split horizon group, but will be forwarded to other SAPs/spoke-SDPs, which are not part of the split horizon group.

5.2.8 VPLS and Spanning Tree Protocol

The 7210 SAS VPLS service provides a bridged or switched Ethernet Layer 2 network. Equipment connected to SAPs forward Ethernet packets into the VPLS service. The 7210 SAS participating in the service learns where the customer MAC addresses reside, on ingress SAPs.

Unknown destinations, broadcasts, and multicasts are flooded to all other SAPs in the service. If SAPs are connected, either through misconfiguration or for redundancy purposes, loops can form and flooded packets can keep flowing through the network. 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 parameters can affect the behavior. For information about command usage, descriptions, and CLI syntax, see [Configuring a VPLS service with CLI](#).

5.2.8.1 Spanning tree operating modes

Per VPLS instance, a preferred STP variant can be configured. The STP variants supported are:

- **rstp**
rapid Spanning Tree Protocol (RSTP) compliant with IEEE 802.1D-2004 - default mode
- **dot1w**
compliant with IEEE 802.1w
- **comp-dot1w**
operation as in RSTP but backwards compatible with IEEE 802.1w (this mode allows interoperability with some MTU types)
- **mstp**
compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q-REV/D5.0-09/2005. This mode of operation is only supported in an mVPLS

While the 7210 SAS initially uses the mode configured for the VPLS, it will dynamically fall back (on a per-SAP basis) to STP (IEEE 802.1D-1998) based on the detection of a BPDU of a different format. A trap or log entry is generated for every change in spanning tree variant.

Some older 802.1W compliant RSTP implementations may have problems with some of the features added in the 802.1D-2004 standard. Inter-working with these older systems is improved with the comp-dot1w mode. The differences between the RSTP mode and the comp-dot1w mode are:

- The RSTP mode implements the improved convergence over shared media feature, for example, RSTP will transition from discarding to forwarding in 4 seconds when operating over shared media. The comp-dot1w mode does not implement this 802.1D-2004 improvement and transitions conform to 802.1w in 30 seconds (both modes implement fast convergence over point-to-point links).
- In the RSTP mode, the transmitted BPDUs contain the port's designated priority vector (DPV) (conforms to 802.1D-2004). Older implementations may be confused by the DPV in a BPDU and may fail to recognize an agreement BPDU correctly. This would result in a slow transition to a forwarding state (30 seconds). For this reason, in the comp-dot1w mode, these BPDUs contain the port's port priority vector (conforms to 802.1w).

The 7210 SAS supports one BPDU encapsulation formats, and can dynamically switch between the following supported formats (on a per-SAP basis):

- IEEE 802.1D STP
- Cisco PVST

5.2.8.2 Multiple Spanning Tree Protocol

The Multiple Spanning Tree Protocol (MSTP) extends the concept of the IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) by allowing grouping and associating VLANs to Multiple Spanning Tree Instances (MSTI). Each MSTI can have its own topology, which provides architecture enabling load balancing by providing multiple forwarding paths. At the same time, the number of STP instances running in the network is significantly reduced as compared to Per VLAN STP (PVST) mode of operation. Network fault tolerance is also improved because a failure in one instance (forwarding path) does not affect other instances.

The 7210 SAS implementation of Management VPLS (mVPLS) is used to group different VPLS instances under single RSTP instance. Introducing MSTP into the mVPLS allows the following:

- Inter-operation with traditional Layer 2 switches in access network.
- Provides an effective solution for dual homing of many business Layer 2 VPNs into a provider network.

5.2.8.2.1 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 [Figure 64: Access resiliency](#) 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 64: Access resiliency



The following mechanisms are supported for the VPLS:

- **STP/RSTP** can be used for all access types
- **M-VPLS with MSTP** can be used as is just for access Type A. MSTP is required for access type B and C.
- **LAG and MC-LAG** can be used for access Type A and B.

5.2.8.3 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.

5.2.8.3.1 MSTP general principles

MSTP represents modification of RSTP that allows the grouping of different VLANs into multiple MSTIs. To enable different devices to participate in MSTIs, they must be consistently configured. A collection of interconnected devices that have the same MST configuration (region-name, revision and VLAN-to-instance assignment) comprises an MST region.

There is no limit to the number of regions in the network, but every region can support a maximum of 16 MSTIs. Instance 0 is a special instance for a region, known as the Internal Spanning Tree (IST) instance. All other instances are numbered from 1 to 4094. IST is the only spanning-tree instance that sends and receives BPDUs (typically BPDUs are untagged). All other spanning-tree instance information is included

in MSTP records (M-records), which are encapsulated within MSTP BPDUs. This means that single BPDU carries information for multiple MSTI which reduces overhead of the protocol.

Any 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 a SR-Series PE can be considered as the part of the same MST region only if the VLAN assignment to IST and MSTIs is identical to the one of Layer 2 switches in access network.

5.2.8.3.2 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.

5.2.8.4 Enhancements to the Spanning Tree Protocol in MPLS-based VPLS service (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C Only)

To interconnect 7210 SAS devices (PE devices) across the backbone, service tunnels (SDPs) are used. These service tunnels are shared among multiple VPLS instances. 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 introduced are fully compliant with the 802.1D-2004 STP specification.

When running MSTP, spoke-SDPs cannot be configured. Also, ensure that all bridges connected by mesh SDPs are in the same region. If not, the mesh will be prevented from becoming active (trap is generated).

To achieve this, all mesh SDPs are dynamically configured as either root ports or designated ports. The PE devices participating in each VPLS mesh determine (using the root path cost learned as part of the normal protocol exchange) which of the 7210 SAS devices is closest to the root of the network. This PE device is internally designated as the primary bridge for the VPLS mesh. Because of this, all network ports on the primary bridges are assigned the designated port role and therefore remain in the forwarding state.

The second part of the solution ensures that the remaining PE devices participating in the STP instance see the SDP ports as a lower cost path to the root 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 preceding features ensure that network ports are never blocked and maintain interoperability with bridges external to the mesh running STP instances.

5.2.8.4.1 L2PT termination

L2PT is used to transparently transport protocol data units (PDUs) of Layer 2 protocols such as STP, CDP, DTP, VTP, PAGP, and UDLD. This allows running these protocols between customer CPEs without involving backbone infrastructure.

The 7210 SAS routers allow transparent tunneling of PDUs across the VPLS core. However, in some network designs, the VPLS PE is connected to CPEs through a legacy Layer 2 network, instead of having direct connections. In such environments termination of tunnels through such infrastructure is required.

L2PT tunnels protocol PDUs by overwriting MAC destination addresses at the ingress of the tunnel to a proprietary MAC address such as 01-00-0c-cd-cd-d0. At the egress of the tunnel, this MAC address is then overwritten back to MAC address of the respective Layer 2 protocol.

The 7210 SAS nodes support L2PT termination for STP BPDUs. More specifically:

- At ingress of every SAP which is configured as L2PT termination, all PDUs with a MAC destination address, 01-00-0c-cd-cd-d0 will be intercepted and their MAC destination address will be overwritten to MAC destination address used for the corresponding protocol. The type of protocol can be derived from LLC and SNAP encapsulation.
- In egress direction, PDUs of the corresponding protocol received on all VPLS ports will be intercepted and L2PT encapsulation will be performed for SAPs configured as L2PT termination points. Because of the implementation reasons, PDU interception and redirection to CPM can be performed only at ingress. Therefore, to comply with the preceding requirement, as soon as at least 1 port of a specific VPLS service is configured as L2PT termination port, redirection of PDUs to CPM will be set on all other ports (SAPs) of the VPLS service.

L2PT termination can be enabled only if STP is disabled in a context of the specific VPLS service.

5.2.8.4.2 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, 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 adding or removing all the required VLAN information in a BPDU payload).

This feature can be enabled on a SAP only if STP is disabled in the context of the specific VPLS service.

5.2.8.4.3 L2PT and BPDU translation

Cisco Discovery Protocol (CDP), Digital Trunking Protocol (DTP), Port Aggregation Protocol (PAgP), Uni-directional Link Detection (UDLD), Virtual Trunk Protocol (VTP), STP (Spanning Tree Protocol) and PVST

(per-VLAN Spanning Tree protocol) are supported on the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

The existing L2PT limitations apply:

- The protocols apply only to VPLS.
- The protocols are mutually exclusive with running STP on the same VPLS as soon as one SAP has L2PT/BPDU translation enabled.
- Forwarding occurs on the CPM and uses CPU processing cycles.

5.2.9 VPLS access redundancy



Note:

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the mechanisms discussed in this section are applicable when access ports or access-uplink ports are used as uplinks.

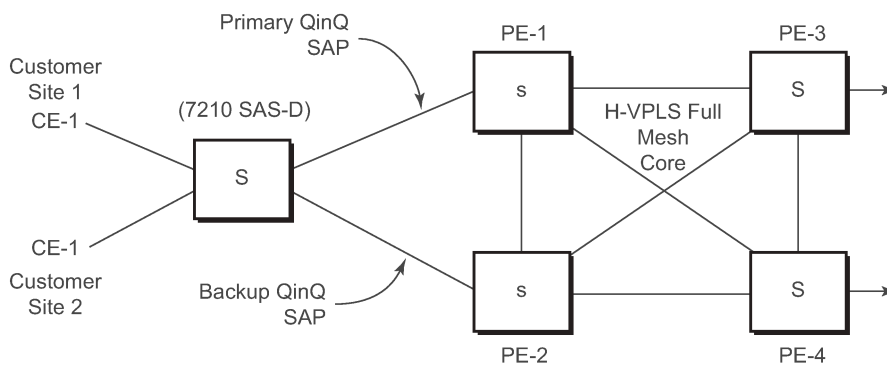
One application of hierarchical VPLS is to use 7210 SAS nodes with access ports or access-uplink ports as uplinks, which serve as Ethernet links to the closest PE node.

The following are several mechanisms 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
- Non-STP-based access using mechanisms such as active/standby links and MC-LAG on the PE

5.2.9.1 STP-based redundant access to VPLS

Figure 65: Dual homed 7210 SAS-D acting as MTU-s in two-tier hierarchy H-VPLS



sw1042

In configuration shown in the preceding figure, STP is activated on the MTU and two PEs to resolve a potential loop.

In this configuration, the scope of STP domain is limited to MTU and PEs, while any topology change needs to be propagated in the whole VPLS domain. When TCN (topology change notification) is received

with in the VPLS domain all MACs learned on SAPs are flushed except the SAP on which TCN was received.

5.2.10 VPLS redundancy on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C using MPLS uplinks

The VPLS standard (RFC 4762, Virtual Private LAN Services Using LDP Signaling) includes provisions for hierarchical VPLS, using point-to-point spoke-SDPs. Two applications have been identified for spoke-SDPs:

- to connect to Multi-Tenant Units (MTUs) to PEs in a metro area network
- to interconnect the VPLS nodes of two networks

In both applications, the spoke-SDPs serve to improve the scalability of VPLS. While node redundancy is implicit in non-hierarchical VPLS services (using a full mesh of SDPs between PEs), node redundancy for spoke-SDPs needs to be provided separately. In VPLS services, only two spoke-SDPs are allowed in an endpoint.

The 7210 SAS routers have implemented special features for improving the resilience of hierarchical VPLS instances, in both MTU and inter-metro applications.


5.2.10.1 Spoke-SDP redundancy on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 VPLS in this topology is not efficient, the node includes functionality which can associate several VPLSes to a single STP instance running over the redundant-SDPs. Node redundancy is 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



5.2.10.2 Spoke-SDP-based redundant access on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP + 8C

This feature provides the ability to have a node deployed as MTUs (Multi-Tenant Unit Switches) multi-homed for VPLS to multiple routers deployed as PEs without requiring the use of mVPLS.

In the configuration example displayed in [Figure 66: H-VPLS with spoke redundancy](#), the MTUs have spoke-SDPs to two PEs devices. One is designated as the primary and one as the secondary spoke-SDP. This is based on a precedence value associated with each spoke. If the primary and secondary spoke-SDPs have the same precedence value, the spoke-SDP with lower ID functions as the primary SDP.

The secondary spoke is in a blocking state (both on receive and transmit) if the primary spoke is available. When the primary spoke becomes unavailable (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.

5.2.10.3 Inter-domain VPLS resiliency using multi-chassis endpoints on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C



Note:


MC-EP is not supported on the 7210 SAS. This section provides an example of how 7210 SAS devices can be used as MTU devices in an MC-EP solution. In this solution, the 7750 SR routers provide the MC-EP functionality.

Inter-domain VPLS refers to a VPLS deployment where sites may be in different domains. An example of inter-domain deployment can be where different Metro domains are interconnected over a Wide Area Network (Metro1-WAN-Metro2) or where sites are in different autonomous systems (AS1-ASBRs-AS2).

Multi-chassis endpoint (MC-EP) provides an alternate solution that does not require RSTP at the gateway VPLS PEs while still using pseudo wires to interconnect the VPLS instances located in the two domains.

MC-EP expands the single chassis endpoint based on active-standby pseudo wires for VPLS shown in the following figure. In the solution depicted by the following figure, 7210 devices are used as MTUs.

Figure 67: H-VPLS resiliency based on AS pseudowires



The active-standby Pseudowires solution is appropriate for the scenario when only one VPLS PE (MTU-s) needs to be dual-homed to two core PEs (PE1 and PE2).

5.2.11 VPLS access redundancy when using MPLS uplinks on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

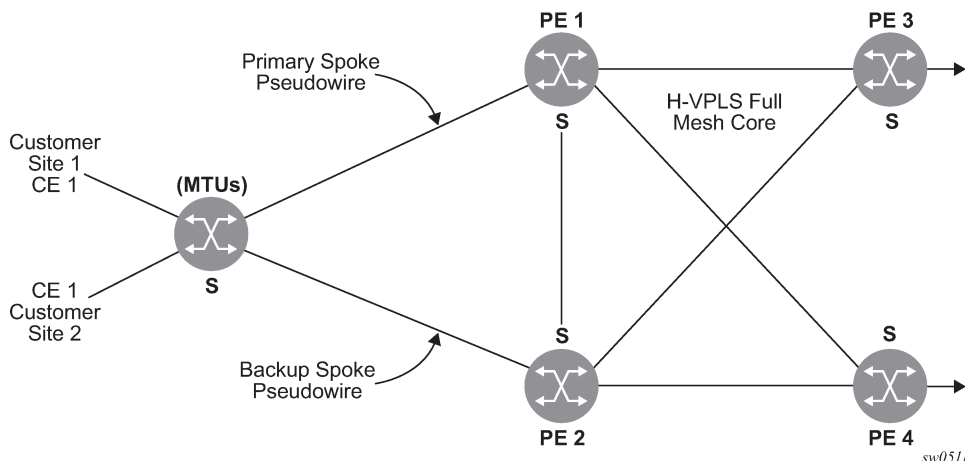
A second application of hierarchical VPLS is using MTUs that are MPLS-enabled which must have spoke-SDPs to the closest PE node. To protect against failure of the PE node, an MTU can be dual-homed.

The following are several mechanisms to resolve a loop in an access network where 7210s are used:

- STP-based access, with or without mVPLS
- Ethernet APS using G.8032

5.2.11.1 STP-based redundant access to VPLS using PWs on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 managed by the specific STP instance.

5.2.11.2 Redundant access to VPLS without STP and using MPLS uplinks on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

The Nokia implementation also alternative methods for providing a 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 L2-B). Upon detection of a link failure PE-C will send MAC-Address-Withdraw messages, which will indicate to all LDP peers that they should flush all MAC addresses learned from PE-C. This will lead that to a broadcasting of packets addressing affected hosts and 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. Per the standard definition, upon receipt of a MAC withdraw message, all MAC addresses, except the ones learned from the source PE, are flushed.

This section specifies that all MAC addresses learned from the source are flushed. This message has been implemented as an LDP address message with vendor-specific type, length, value (TLV), and is called the flush-mine message.

The advantage of this approach (as compared to RSTP based methods) is that only MAC-affected addresses are flushed and not the full forwarding database. While this method does not provide a mechanism to secure alternative loop-free topology, the convergence time is dependent on the speed of the specific CE device will open alternative link (L2-B switch) as well as on the speed PE routers will flush their FDB.

In addition, this mechanism is effective only if PE and CE are directly connected (no hub or bridge) as it reacts to physical failure of the link.

5.2.11.3 MAC flush message processing in VPLS services with MPLS uplinks on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 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-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. Note, that 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.

5.2.11.4 Dual homing to a VPLS service on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Figure 69: 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 L2-B). Upon detection of a link failure PE-C will send MAC-Address-Withdraw messages, which will indicate to all LDP peers that they should flush all MAC

addresses learned from PE-C. This will lead that to a broadcasting of packets addressing affected hosts and relearning process in case an alternative route exists.

Note that the message described here is different from the message described in draft-ietf-l2vpnpls-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 (L2-B switch in [Figure 69: Dual homed CE connection to VPLS](#)) 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.

5.2.12 VPLS service considerations

This section describes various 7210 SAS service features and any special capabilities or considerations as they relate to VPLS services.

5.2.12.1 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 on an access port:

- Ethernet null (Supported on all platforms)
- Ethernet Dot1q (Supported on all platforms)
- Ethernet QinQ (Supported on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

The following encapsulations are supported on an access-uplink port:

- Ethernet QinQ (Supported on all platforms)

5.2.12.2 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

Any VLAN tags are ignored and the packet goes to a default service for the SAP.

2. Dot1q encapsulation defined on ingress

Only first VLAN tag is considered.

3. QinQ encapsulation defined on ingress

Both VLAN tags are considered.

Note that the SAP can be defined with a wild-card for the inner label (for example, "100.*"). In this situation all packets with an outer label of 100 will be treated as belonging to the SAP. If, on the same physical link, there is also a SAP defined with a QinQ encapsulation of 100.1, then traffic with 100.1 will go to that SAP and all other traffic with 100 as the first label will go to the SAP with the *100.* definition.

In situations 2 and 3, traffic encapsulated with tags for which there is no definition are discarded.

5.3 BGP-AD VPLS

The following section provides information about BGP-AD VPLS.

5.3.1 BGP-AD and Target LDP (T-LDP) interaction on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C




Note:

This feature is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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 70: 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.

5.3.2 BGP-AD VPLS - SDP usage

Service Access Points (SAP) are linked to transport tunnels using Service Distribution Points (SDP). The service architecture of the 7210 platform allows services to be abstracted from the transport network.

MPLS transport tunnels are signaled using the Resource Reservation Protocol (RSVP-TE) or by the Label Distribution Protocol (LDP). The capability to automatically create an SDP only exists for LDP based transport tunnels. Using a manually provisioned SDP is available for both RSVP-TE and LDP transport tunnels. See the 7210 SAS-K 2F6C4T, K 3SFP+ 8C MPLS Guide for more information about MPLS, LDP, and RSVP.

5.3.3 BGP-AD VPLS - automatic creation of SDPs

When BGP AD is used for LDP VPLS and LDP is used as the transport tunnel there is no requirement to manually create an SDP. The LDP SDP can be automatically instantiated using the information advertised by BGP AD. This simplifies the configuration on the service node.

Enabling LDP on the IP interfaces connecting all nodes between the ingress and the egress, builds transport tunnels based on the best IGP path. LDP bindings are automatically built and stored in the hardware. These entries contain an MPLS label pointing to the best next hop along the best path toward the destination.

When two endpoints need to connect and no SDP exists, a new SDP will automatically be constructed. New services added between two endpoints that already have an automatically created SDP will be immediately used. No new SDP will be constructed.

The far-end information is gleaned from the BGP next hop information in the NLRI. When services are withdrawn with a BGP_Unreach_NLRI, the automatically established SDP will remain up as long as at least one service is connected between those endpoints. An automatically created SDP will be removed and the resources released when the only or last service is removed.

5.3.4 BGP-AD VPLS - manually provisioned SDP

The carrier is required to manually provision the SDP if they create transport tunnels using RSVP-TE. Operators have the option to choose a manually configured SDP, if they use LDP as the tunnel signaling protocol. The functionality is the same regardless of the signaling protocol.

Creating a BGP-AD enabled VPLS service on an ingress node with the manually provisioned SDP option causes the Tunnel Manager to search for an existing SDP that connects to the far-end PE. The far-end IP information is gleaned from the BGP next hop information in the NLRI. If a single SDP exists to that PE, it is used. If no SDP is established between the two endpoints, the service remains down until a manually configured SDP becomes active.

When multiple SDPs exist between two endpoints, the tunnel manager selects the appropriate SDP. The algorithm preferred SDPs with the best (lower) metric. Should there be multiple SDPs with equal metrics, the operational state of the SDPs with the best metric is considered. If the operational state is the same, the SDP with the higher sdp-id is used. If an SDP with a preferred metric is found with an operational state that is not active, the tunnel manager flags it as ineligible and restarts the algorithm.

5.3.5 BGP-AD VPLS - automatic instantiation of pseudowires (SDP bindings)

The choice of manual or auto provisioned SDPs has limited impact on the amount of required provisioning. Most of the savings are achieved through the automatic instantiation of the pseudowire infrastructure (SDP bindings). This is achieved for every auto-discovered VSIs through the use of the pseudowire template concept. Each VPLS service that uses BGP AD contains the "pw-template-binding" option defining specific Layer 2 VPN parameters. This command references a "pw-template" which defines the pseudowire parameters. The same "pwtemplate" may be referenced by multiple VPLS services. As a result, changes to these pseudowire templates have to be treated with great care as they may impact many customers at once.

The 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:

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.

5.3.6 BGP-AD VPLS - mixing statically configured and auto-discovered pseudowires in a VPLS service

The services implementation allows for manually provisioned and auto-discovered pseudowire (SDP bindings) to 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 may 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.

5.3.7 BGP-AD VPLS - resiliency schemes

The use of BGP-AD on the network side, or in the backbone, does not affect the different resiliency schemes 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

Routed VPLS (R-VPLS) allows a VPLS instance to be associated with an IP interface on the 7210 SAS-D, 7210 SAS-Dxp, or 7210 SAS-K 2F1C2T. IPv4 addressing and forwarding is supported on the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T platforms. IPv6 addressing and forwarding is supported on the 7210 SAS-Dxp platform.

R-VPLS allows a VPLS instance to be associated with an IP interface on the 7210 SAS-K 2F6C4T or 7210 SAS-K 3SFP+ 8C. Only IPv4 addressing and forwarding is supported on these platforms.

Within an R-VPLS service, traffic with a destination MAC matching that of the associated IP interface is routed based on the IP forwarding table; all other traffic is forwarded based on the VPLS forwarding table.

On the 7210 SAS-D and 7210 SAS-K 2F1C2T, 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 inband management of the 7210 SAS nodes in a ring deployment using a single IPv4 subnet, reducing the number of IP subnets needed.

On the 7210 SAS-Dxp, an R-VPLS service can be associated with an IPv4 or IPv6 interface and supports only static routing. It is primarily designed for use of in-band management of the node. It allows for inband management of the 7210 SAS nodes in a ring deployment using a single IPv4 or IPv6 subnet, reducing the number of IP subnets needed.

On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, R-VPLS service can be associated with an IPv4 interface and supports static routing and other routing protocols. It can be used to provide a service to the customer or for in-band management of the node.

5.4.1 IES IP interface binding

A standard IP interface within an existing IES 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) may have other IP interfaces bound to other VPLS service contexts.

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 is enabled. If a name is applied to the VPLS service while the flag is set, the system scans the existing IES or VPRN services for an IP interface bound to the specified service name. If an IP interface is found, the IP interface is attached to the VPLS service associated with the name. Only one interface can be bound to the specified name.

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

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

5.4.3 Service binding requirements

If the defined service name is created on the system, the system checks to ensure that the service type is VPLS. If the created service type is VPLS, the IP interface is eligible to enter the operationally upstate.



Note:

On the 7210 SAS-K 2F6C4T, 7210 SAS-K 2F1C2T, and 7210 SAS-K 3SFP+ 8C, use the R-VPLS tag in the creation of the VPLS service, otherwise the service cannot be bound to an IP interface. This is not required on the 7210 SAS-D and 7210 SAS-Dxp.

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 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 any time. Only one interface can be bound to a service. When an IP interface is bound to a service name and the IP interface is administratively up, the system scans for a VPLS service context using the name and takes the following actions:

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

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 (that is, SAP) is operational.

5.4.3.5 IP interface MTU and fragmentation on 7210 SAS-D and 7210 SAS-Dxp

In 7210 SAS-D Access-Uplink mode and 7210 SAS-Dxp Access-Uplink mode, VPLS service MTU is not supported. The user must ensure that the port MTU is configured appropriately so that the largest packet traversing through any of the SAPs (virtual ports) of the VPLS service can be forwarded out of any of the SAPs. VPLS services do not support fragmentation and can discard packets larger than the configured port MTU.

When an IP interface is associated with a VPLS service, the IP-MTU is based on either the administrative value configured for the IP interface or an operational value derived from port MTU of all the SAPs configured in the service. The port MTU excluding the Layer 2 Header and tags for all the ports which have SAPs configured in this VPLS service are considered and the minimum value among those are computed (which is called computed MTU). The operational value of the IP interface is set as follows:

- If the configured (administrative) value of IP MTU is greater than the computed MTU, then the operational IP MTU is set to the computed MTU.
- If the configured (administrative) value of IP MTU is lesser than or equal to the computed MTU, then operational IP MTU is set to the configured (administrative) value of IP MTU.

5.4.3.6 IP interface MTU and fragmentation on 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8CC

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

VPLS services do not support fragmentation and must discard any Layer 2 packet larger than the service MTU after the service delineation overhead is removed.

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

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

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

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.



Note:

Static ARP is only supported on the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

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 for the MAC (routed or bridged) is targeted to the specific SAP where the MAC has been learned. As with ARP entries, static MAC entries may be created in the VPLS FIB. Dynamically learned MAC addresses can age out or be flushed from the VPLS FIB while static MAC entries always remain associated with a specific virtual port. Dynamic MACs may also be relearned on another VPLS SAP than the current SAP in the FIB. In this case, the system automatically moves the MAC FIB entry to the new VPLS SAP.



Note:

- On the 7210 SAS-D and 7210 SAS-Dxp, 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. This restriction does not apply to the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.
- On the 7210 SAS-D and 7210 SAS-Dxp, if the VPLS FIB limit is reached and we are no longer able to learn new MAC address, ARP will also not be learnt. This restriction does not apply to the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

5.4.5 Specific ARP cache behavior

In typical routing behavior, the system uses the IP route table to select the egress interface, an ARP entry is used forward the packet to the appropriate Ethernet MAC. With routed VPLS, the egress IP interface may be represented by multiple egress (VPLS service SAPs).

The following table describes how the ARP cache and MAC FIB entry states interact.

Table 38: 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 typically happen on the 7210 SAS-D and 7210 SAS-Dxp, as when an L2 entry is removed from the FDB, the matching MAC address is also removed from the ARP cache. On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, 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 SAPs only supported on standard Ethernet ports

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

5.4.6.1 LAG port membership constraints

If a LAG has a non-supported port type as a member, a SAP for the routing-enabled VPLS service cannot be created on the LAG. When one or more routing enabled VPLS SAPs are associated with a LAG, a non-supported Ethernet port type cannot be added to the LAG membership.

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 (The flag also cannot be enabled while any of these features are applied to the VPLS service). The following restrictions apply to both network mode and access-uplink mode unless called out separately:

- On the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, spoke or mesh SDP bindings cannot be configured.
- On the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the VPLS service type cannot be M-VPLS.
- On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the VPLS Service type must be 'r-vpls' and any other VPLS service is not allowed.
- MVR from Routed VPLS and to another SAP is not supported.
- Default QinQ SAPs is not supported in R-VPLS service.
- The "allow-ip-int-binding" command cannot be used in a VPLS service acting as the G.8032 control instance.
- IPv4 filters (ingress and egress) can be used with the R-VPLS SAPs. Additionally, IP ingress override filters are supported which affects the behavior of the IP filters attached to the R-VPLS SAPs. Please see the following for more information about use of ingress override filters.
- MAC filters (ingress and egress) are not supported for use with R-VPLS SAPs.
- VPLS IP interface is not allowed in a R-VPLS service. The converse also holds.
- On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, during creation of the VPLS service the keyword 'rvpls' must be used. It allows the software know that this is a VPLS service to which an IP interface will be associated.
- On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the VPLS service can be configured either access SAP or Access-Uplink SAPs.
- On the 7210 SAS-D and 7210 SAS-Dxp, 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, except that default QinQ SAPs cannot be used in a R-VPLS service. The allowed SAP combinations are similar to that available in a plain VPLS service and is as specified in the preceding table in the services Chapter (with the exception noted before).
- On 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, VPLS service can use the following 'svc-sap-type' values: any. Only specific SAP combinations are allowed for a specific svc-sap-type, except that default QinQ SAPs cannot be used in a R-VPLS service. The allowed SAP combinations are similar to that available in a plain VPLS service and is as specified in the preceding table in the services Chapter (with the exception noted before).
- G.8032 or mVPLS/STP based protection mechanism can be used with R-VPLS service. A separate G.8032 control instance or a separate mVPLS/STP instance needs to be used and the R-VPLS SAPs needs to be associated with these control instances such that the R-VPLS SAP's forwarding state is driven by the control instance protocols.
- IP multicast is not supported in the R-VPLS service.
- On the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, DHCP snooping is not supported for the SAPs configured in the routed VPLS service. Instead, DHCP relay can be enabled on the IES service associated with the routed VPLS service.
- 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 on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

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 (IPv4) 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 a R-VPLS service:

Table 39: 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
Destination MAC = IES IP interface MAC and destination IP not on same subnet as IES IP interface and route to destination IP does not exist	No	No	No for 7210 SAS-D and 7210 SAS-Dxp devices Yes for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C devices
Destination MAC = IES IP interface MAC and destination IP not on same subnet as IES IP interface and route to destination IP exists	No	No	Yes
Destination MAC = IES IP interface MAC and IP TTL = 1	No	No	No for 7210 SAS-D and 7210 SAS-Dxp devices Yes for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C devices
Destination MAC = IES IP interface MAC and IPv4 packet with Options	No	No	No for 7210 SAS-D and 7210 SAS-Dxp devices

Type of traffic	SAP Ingress IPv4 filter	SAP Egress IPv4 filter	Ingress override IPv4 filter
			Yes for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C devices
Destination MAC = IES IP interface MAC and IPv4 Multicast packet	No	No	No for 7210 SAS-D and 7210 SAS-Dxp devices Yes for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C devices

Table 40: ACL lookup behavior without ingress override filter attached to an IES interface in a R-VPLS service

Type of traffic	SAP ingress IPv4 filter	SAP egress IPv4 filter
Destination MAC != IES IP interface MAC	Yes	Yes
Destination MAC = IES IP interface MAC and Destination IP on same subnet as IES IP interface	Yes	No
Destination Mac = IES IP interface mac and destination IP not on same subnet as IES IP interface and route to destination IP does not exist	No for 7210 SAS-D and 7210 SAS-Dxp devices Yes for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C devices	No
Destination Mac = IES IP interface MAC and destination IP not on same subnet as IES IP interface and route to destination IP exists	Yes	No
Destination MAC = IES IP interface MAC and IP TTL = 1	No for 7210 SAS-D and 7210 SAS-Dxp devices Yes for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C devices	No
Destination MAC = IES IP interface MAC and IPv4 packet with Options	No for 7210 SAS-D and 7210 SAS-Dxp devices Yes for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C devices	No

Type of traffic	SAP ingress IPv4 filter	SAP egress IPv4 filter
Destination MAC = IES IP interface MAC and IPv4 Multicast packet	No for 7210 SAS-D and 7210 SAS-Dxp devices Yes for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C devices	No

5.4.7.1 QoS support for VPLS SAPs and IP interface

- SAP ingress classification (IPv4 and MAC criteria) is supported for SAPs configured in the service. SAP ingress policies cannot be associated with IES IP interface.
- On the 7210 SAS-D and 7210 SAS-Dxp, egress port based queuing and shaping are available. It is shared among all the SAPs on the port.
- On the 7210 SAS-D and 7210 SAS-Dxp, port-based egress marking is supported for both routed packets and bridged packets. The existing access egress QoS policy can be used for Dot1p marking and DSCP marking.
- On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, per-SAP egress queuing, shaping and scheduling is available. Per SAP egress Dot1p marking is supported for both routed packet and bridged packets.
- On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, 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.

5.4.7.2 Routing related protocols on 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Routed VPLS is supported only in the base routing instance on the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T. It is supported in both the base routing instance and VPRN services on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C. The following table lists the support available for routing protocols on IP interfaces bound to a VPLS service for different platforms.

Table 41: Routing protocols on IPv4 interfaces bound to a VPLS service

Protocol	7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T	7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
Static routing	Supported	Supported
BGP	Not Supported	Supported
OSPF	Not Supported	Supported

Protocol	7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T	7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
IS-IS	Not Supported	Supported
BFD	Not Supported	Not Supported
VRRP	Not Supported	Supported
ARP and Proxy-Arp	ARP is supported	Both are supported
DHCP Relay ¹⁰	Supported	Supported

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 service.

5.4.8 R-VPLS support and caveats

Routed VPLS supported functionality and restrictions for both access-uplink and network mode is listed as follows. The following is applicable to both the modes, unless called out explicitly:

- On the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T, Routed VPLS can be bound to only IES IP interface. It cannot be bound to VPRN IP interface.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, routed VPLS can be bound to IES or VPRN IP interface.
- On the 7210 SAS-D, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, only IPv4 addressing and forwarding is supported with Routed VPLS services. IPv6 addressing and forwarding is not supported.
- The 7210 SAS-Dxp supports both IPv4 and IPv6 addressing and forwarding with R-VPLS services.
- On the 7210 SAS-D and 7210 SAS-Dxp, static ARP cannot be configured with an IES IP interface associated with an R-VPLS, though static MAC can be configured in an R-VPLS service.
- On the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T, only static routes are supported. No dynamic routing protocols are supported.
- On the 7210 SAS-D and 7210 SAS-Dxp, 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.
- On the 7210 SAS-D and 7210 SAS-Dxp, multiple SAPs configured on the same port cannot be part of the same R-VPLS Service. A single service can only be configured with a single SAP on a specific port. This restriction does not apply to the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C. On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, multiple SAPs configured on the same port can be part of the same service.

¹⁰ DHCP relay can be configured for the IES interface associated with the Routed VPLS service. DHCP snooping cannot be configured on the VPLS SAPs in the routed VPLS Service.

- The 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C support service MTU configuration for R-VPLS.
- On the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, service-based SHGs are not supported in an R-VPLS service.

5.5 Configuring a VPLS service with CLI

This section provides information to configure VPLS services using the command line interface.

5.5.1 Basic configuration

The following fields require specific input (there are no defaults) to configure a basic VPLS service:

- Customer ID (See [Configuring customer accounts](#))
- For a local service, configure two SAPs, specifying local access ports and encapsulation values.
- For a distributed service, configure a SAP and an SDP (only for 7210 SAS devices that support MPLS uplinks) for each far-end node.

The **svc-sap-type** option must be specified at service creation. The default option is **any**. The following is a sample configuration output of a local VPLS service with a null-star SAP type on ALA-1. This service comprises of a null access SAP and an Uplink LAG SAP.

```
*A:ALA-1>config>service# info
-----
...
    vpls 7 customer 7 create
        stp
            shutdown
        exit
        sap 1/1/21 create
        exit
        sap lag-1:700 create
        exit
        no shutdown
    exit
...
-----
*A:ALA-1>config>service#
```

The following is a sample configuration output of a distributed VPLS service between ALA-1, ALA-2, and ALA-3.

```
vpls 100 customer 1 svc-sap-type any create
    shutdown
    bgp
    exit
    stp
        shutdown
    exit
exit
epipe 101 customer 1 svc-sap-type any create
    shutdown
exit
-----
```

```
*A:ALA-1>config>service#
*A:ALA-2>config>service# info
-----
...
vpls 9000 customer 6 create
description "This is a distributed VPLS."
stp
shutdown
exit
sap 1/1/5:16 create
description "VPLS SAP"
exit
spoke-sdp 2:22 create
exit
no shutdown
exit
...
-----
*A:ALA-2>config>service#
```

5.5.2 Common configuration tasks

This section provides a brief overview of the tasks that must be performed to configure local VPLS services and provides the CLI commands.

For VPLS services:

1. Associate VPLS service with a customer ID
2. Define SAPs:
 - Select 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)
3. Modify STP default parameters (optional) (see [VPLS and Spanning Tree Protocol](#))
4. Enable service

5.5.3 Configuring VPLS components

5.5.3.1 Creating a VPLS service

Use the following syntax to create a VPLS service (for 7210 SAS-D).

```
config>service# vpls service-id [customer customer-id] [create] [vpn <vpn-id>] [m-vpls] [svc-
sap-type {null-star|dot1q-preserve|any }] [customer-vid <vlan-id>]
```

The following is a sample VPLS configuration (for 7210 SAS-D) 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 svc-sap-type any create
description "This is a Distributed VPLS with ANY SAP"
stp
shutdown
exit
no shutdown
exit
vpls 3000 customer 8 svc-sap-type dot1q-preserve customer-vid 300 create
description "This is a VPLS with QinQ Uplink SAP"
stp
shutdown
exit
no shutdown
exit
...
-----
*A:ALA-1>config>service>vpls#
```

Use the following syntax to create a VPLS service (for 7210 SAS-Dxp).

```
- config>service# vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls]
[customer-vid vlan-id][svc-sap-type {null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-
vpls]
```

The following is a sample VPLS configuration (for 7210 SAS-Dxp) 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 svc-sap-type any create
description "This is a Distributed VPLS with ANY SAP"
stp
shutdown
exit
no shutdown
exit
vpls 3000 customer 8 svc-sap-type dot1q-preserve customer-vid 300 create
description "This is a VPLS with QinQ Uplink SAP"
stp
shutdown
exit
no shutdown
exit
...
-----
*A:ALA-1>config>service>vpls#
```

Use the following syntax to create a VPLS service (for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C).

```
– config>service# vpls service-id [customer customer-id] [create] [vpn <vpn-id>] [m-vpls] [svc-  
sap-type {dot1q-range | any}] [customer-vid <vlan-id>]
```

The following is a sample VPLS configuration (for 7210 SAS-2F1C2T) output.

```
*A:sask_duti>config>service>vpls$ info
-----
      shutdown
      stp
          no shutdown
      exit
-----
*A:sask_duti>config>service>vpls$ info detail
-----
      shutdown
      no description
      service-mtu-check
      service-mtu 1514
      no def-mesh-vc-id
      no disable-learning
      no disable-aging
      no discard-unknown
      fdb-table-size 250
      fdb-table-high-wmark 95
      fdb-table-low-wmark 90
      no host-connectivity-verify
      no shcv-policy-ipv4
      no send-flush-on-failure
      local-age 300
      no mfib-table-size
      mfib-table-high-wmark 95
      mfib-table-low-wmark 90
      remote-age 900
      no propagate-mac-flush
      stp
          priority 32768
          hello-time 2
          forward-delay 15
          max-age 20
          hold-count 6
          mode rstp
          mst-max-hops 20
          no mst-name
          mst-revision 0
          no shutdown
      exit
      igmp-snooping
          shutdown
          query-interval 125
          robust-count 2
          report-src-ip 0.0.0.0
          no query-src-ip
          mvr
              shutdown
              no description
              no group-policy
          exit
      exit
      mac-move
```

```

        move-frequency 2
        retry-timeout 10
        shutdown
    exit
    static-mac
    exit
    no vsd-domain
-----
*A:sask_duti>config>service>vpls$

```

5.5.3.1.1 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

```

The following is sample 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#

```

5.5.3.1.2 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 parameters must be done in the constraints of the following two formula:

$$2 \times (\text{Bridge_Forward_Delay} - 1.0 \text{ seconds}) \geq \text{Bridge_Max_Age}$$

$$\text{Bridge_Max_Age} \geq 2 \times (\text{Bridge_Hello0_Time} + 1.0 \text{ seconds})$$

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.

5.5.3.1.2.1 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 device. When STP on the VPLS is administratively enabled, but the administrative state of a SAP is down, BPDUs received on such a SAP are discarded.

```
config>service>vpls service-id# stp
no shutdown
```

5.5.3.1.2.2 Mode

To be compatible with the different iterations of the IEEE 802.1D standard, the 7210 SAS devices support several variants of the Spanning Tree protocol:

- **rstp**
Rapid Spanning Tree Protocol (RSTP) compliant with IEEE 802.1D-2004 - default mode.
- **dot1w**
Compliant with IEEE 802.1w.
- **comp-dot1w**
Operation as in RSTP but backwards compatible with IEEE 802.1w. This mode was introduced for interoperability with some MTU types.
- **mstp**
Compliant with the Multiple Spanning Tree Protocol specified in IEEE 802.1Q REV/D5.0-09/2005. This mode of operation is only supported in an mVPLS.

See section [Spanning tree operating modes](#) for details on these modes.

```
config>service>vpls service-id# stp
mode {rstp | comp-dot1w | dot1w | mstp}
```

Default: rstp

5.5.3.1.2.3 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

5.5.3.1.2.4 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

5.5.3.1.2.5 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 (e.g. on shared links, see the following), 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

5.5.3.1.2.6 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

5.5.3.1.2.7 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

5.5.3.1.2.8 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 its own tree within the region, therefore making sure different VLANs follow different paths.

You can assign non overlapping VLAN ranges to each instance. VLANs not assigned to an instance are implicitly assumed to be in instance 0, which is also called the CIST. This CIST cannot be deleted or created.

The parameter that can be defined per instance are **mst-priority** and **vlan-range**:

- **mst-priority**

The bridge-priority for this specific mst-instance. It follows the same rules as bridge-priority. For the CIST, the bridge-priority is used.

- **vlan-range**

The VLANs are mapped to this specific mst-instance. If no VLAN-ranges are defined in any mst-instances, then all VLANs are mapped to the CIST.

5.5.3.1.2.9 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.

5.5.3.1.2.10 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.

5.5.3.1.2.11 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.

5.5.3.2 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.

5.5.3.2.1 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.

Example

```
*A:ALA-1>config>service# info
-----
vpls 1000 customer 1 create
  description "This is a Local VPLS with NULL SAP"
  stp
    shutdown
  exit
  sap 1/1/1 create
  exit
  sap 1/1/2 create
  exit
  sap 1/1/3:500.* create
  exit
  no shutdown
  exit
vpls 2000 customer 6 create
```

```
description "This is a Local VPLS with DOT1Q SAP"
stp
  shutdown
exit
sap 1/1/4:100 create
exit
sap 1/1/5:200 create
exit
sap 1/1/3:900.* create
exit
no shutdown
exit
vpls 3000 customer 8 create
  description "This is a Local VPLS"
  stp
    shutdown
  exit
  sap 1/1/4:300 create
  exit
  sap 1/1/5:300 create
  exit
  sap 1/1/3:1200 create
  exit
  no shutdown
exit
-----
*A:ALA-1>config>service#
```

5.5.3.2.2 Distributed VPLS service with SAPs (on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP + 8C)

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.

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 100 customer 1 svc-sap-type any create
  shutdown
  no description
  service-mtu-check
  service-mtu 1514
  no def-mesh-vc-id
  no disable-learning
  no disable-aging
  no discard-unknown
  fdb-table-size 250
  fdb-table-high-wmark 95
  fdb-table-low-wmark 90
  no host-connectivity-verify
  no shcv-policy-ipv4
  no send-flush-on-failure
  local-age 300
  no mfib-table-size
```

```
mfib-table-high-wmark 95
mfib-table-low-wmark 90
remote-age 900
no propagate-mac-flush
bgp
  no route-distinguisher
  no vsi-export
  no vsi-import
  no route-target
exit
-----
*A:ALA-3>config>service#
```

5.5.3.2.3 Configuring default QinQ SAPs to pass all traffic from access to access-uplink port without any tag modifications

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
      received-count
  exit
exit
exit
sap 1/1/6:*. * create
  statistics
    ingress
      received-count
  exit
exit
exit
exit
```

5.5.3.2.4 Configuring SAP-specific STP parameters

When a VPLS has STP enabled, each SAP within the VPLS has STP enabled by default. The operation of STP on each SAP is governed by the following.

5.5.3.2.4.1 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)

5.5.3.2.4.2 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 its own virtual port number that is unique to every other SAP defined on the VPLS. The virtual port number is assigned at the time that the SAP is added to the VPLS.

Since the order in which SAPs are added to the VPLS is not preserved between reboots of the system, the virtual port number may change between restarts of the STP instance. To achieve consistency after a reboot, the virtual port number can be specified explicitly.

```
config>service>vpls>sap# stp  
port-num number
```

Range: 1 — 2047

Default: automatically generated

Restore Default: no port-num

5.5.3.2.4.3 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 details on the virtual port number.

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

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

```
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

5.5.3.2.4.4 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. Since SAPs are controlled by complex queuing dynamics, the STP path cost is a static configuration on the 7210 SAS.

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

5.5.3.2.4.5 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**

5.5.3.2.4.6 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

5.5.3.2.4.7 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

5.5.3.2.4.8 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.

5.5.3.2.5 STP SAP operational states

The operational state of STP within a SAP controls how BPDUs are transmitted and handled when received.

5.5.3.2.5.1 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.

5.5.3.2.5.2 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.

5.5.3.2.5.3 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.

5.5.3.2.5.4 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.

5.5.3.2.5.5 SAP BDU encapsulation state

STP is associated with a VPLS service like PVST is associated per VLAN. The main difference resides in the Ethernet and LLC framing and a type-length-value (TLV) field trailing the BDU.

IEEE 802.1d (referred as dot1d) and Cisco per VLAN Spanning Tree (PVST) BDU encapsulations are supported on a per SAP basis. The STP is associated with a VPLS service like PVST is per VLAN. The difference between the two encapsulations is in the Ethernet and LLC framing and a type-length-value

(TLV) field trailing the BPDU. The encapsulation format cannot be configured by the user, the system automatically determines the encapsulation format based on the BPDUs received on the port.

[Table 42: SAP BPDU encapsulation states](#) shows differences between Dot1d and PVST Ethernet BPDU encapsulations based on the interface encap-type field.

Table 42: SAP BPDU encapsulation states

Field	dot1d encap-type null	dot1d encap-type dot1q	PVST encap-type null	PVST encap-type dot1q
Destination MAC	01:80:c2:00:00:00	01:80:c2:00:00:00	N/A	01:00:0c:cc:cc:cd
Source MAC	Sending Port MAC	Sending Port MAC	N/A	Sending Port MAC
EtherType	N/A	0x81 00	N/A	0x81 00
Dot1p and CFI	N/A	0xe	N/A	0xe
Dot1q	N/A	VPLS SAP ID	N/A	VPLS SAP encap value
Length	LLC Length	LLC Length	N/A	LLC Length
LLC DSAP SSAP	0x4242	0x4242	N/A	0xaaaa (SNAP)
LLC CNTL	0x03	0x03	N/A	0x03
SNAP OUI	N/A	N/A	N/A	00 00 0c (Cisco OUI)
SNAP PID	N/A	N/A	N/A	01 0b
CONFIG	Standard 802.1d	Standard 802.1d	N/A	Standard 802.1d
TLV: Type & Len	N/A	N/A	N/A	58 00 00 00 02
TLV: VLAN	N/A	N/A	N/A	VPLS SAP encap value
Padding	As Required	As Required	N/A	As Required

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

- **Dot1d**

This state specifies that the switch is currently sending IEEE 802.1d standard BPDUs. The BPDUs are tagged or non-tagged based on the encapsulation type of the egress interface and the encapsulation value defined in the SAP. A SAP defined on an interface with encapsulation type Dot1q continues in the dot1d BPDU encapsulation state until a PVST encapsulated BPDU is received. When a PVST-encapsulated BPDU is received, the SAP converts to the PVST encapsulation state. Each received BPDU must be correctly IEEE 802.1q tagged, if the interface encapsulation type is defined as dot1q. PVST BPDUs are silently discarded, if received, when the SAP is on an interface defined with encapsulation type null.

- **PVST**

This state specifies that the switch is currently sending proprietary encapsulated BPDUs. PVST BPDUs are only supported on Ethernet interfaces with the encapsulation type set to dot1q. The SAP continues in the PVST BPDU encapsulation state until a dot1d encapsulated BPDU is received. When a dot1d-encapsulated BPDU is received, the SAP reverts to the dot1d encapsulation state. Each received BPDU must be 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 BPDU encapsulation state for SAPs defined on Ethernet interface with encapsulation type set to null.

5.5.3.2.6 Configuring VPLS SAPs with per-service split horizon



Note:

Split Horizon group per service is only supported on the 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

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#
```

5.5.3.2.6.1 Configuring SDP bindings



Note:

SDPs are only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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 data-link specific information of the VC.

5.5.3.2.6.2 Configuring VPLS spoke-SDPs with split horizon



Note:

Split Horizon group with spoke-SDPs is only supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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#
```

5.5.4 Configuring VPLS redundancy

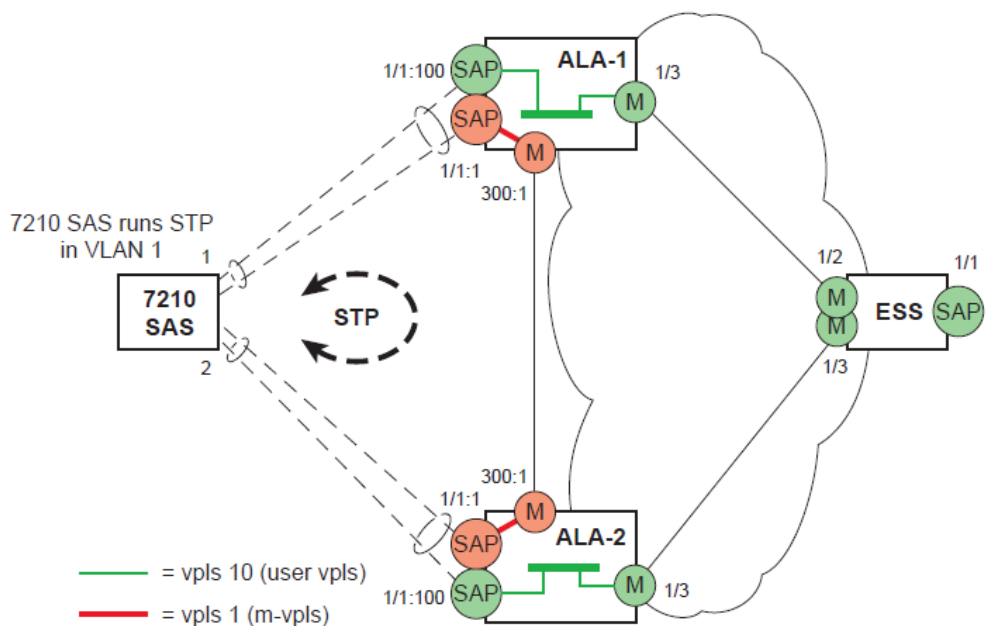
This section describes the service management tasks.

5.5.4.1 Creating a management VPLS for SAP protection

This section provides a brief overview of the tasks that must be performed to configure a management VPLS for SAP protection and provides the CLI commands, see [Figure 71: Example configuration for protected VPLS SAP](#). The following tasks should be performed on both nodes providing the protected VPLS service:

1. Create an access uplink SAPs to the peer node.
2. Create a management VPLS.
3. Define a SAP in the m-vpls on the port 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.
4. Optionally modify STP parameters for load balancing.
5. Create access uplink SAPs in the m-vpls using the access uplink SAPs defined in Step 1.
6. Enable the management VPLS service and verify that it is operationally up.
7. Create a list of VLANs on the port managed by this management VPLS.
8. Create one or more user VPLS services with SAPs on VLANs in the range defined by Step 6.

Figure 71: Example configuration for protected VPLS SAP



OSSG0477210SASE

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 is a sample VPLS configuration output.

```
*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#
```

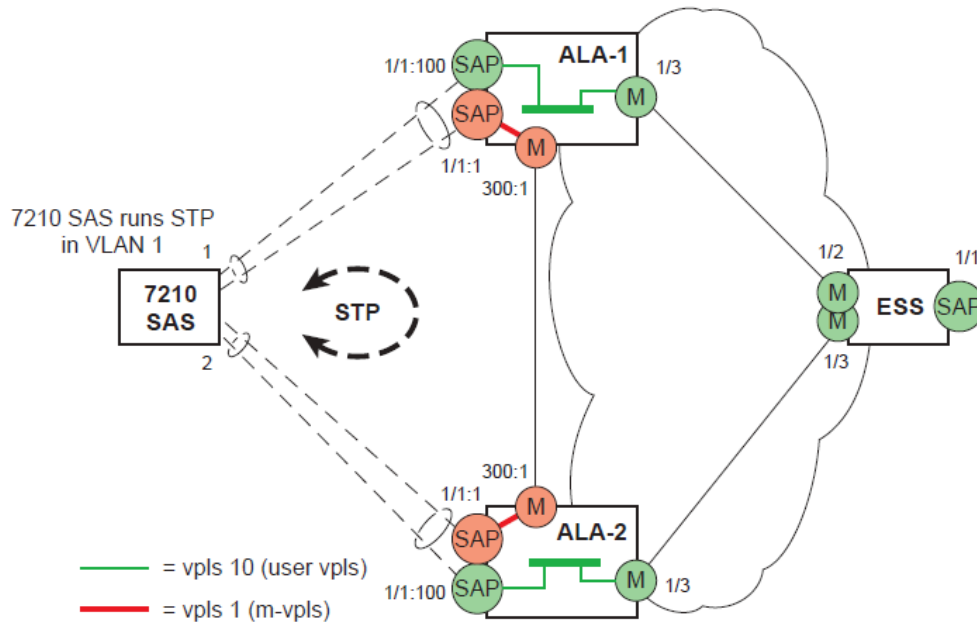
5.5.4.2 Configuring load balancing with management VPLS

With the concept of management VPLS, it is possible to load balance the user VPLS services across the two protecting nodes. This is done by creating two management VPLS instances, where both instances have different active QinQ SAPs (by changing the STP path-cost). When different user VPLS services are associated with either the two management VPLS services, the traffic will be split across the two QinQ SAPs. Load balancing can be achieved in SAP protection scenarios.

Example

The following figure is an example configuration for load balancing across management VPLS.

Figure 72: Example configuration for load balancing across with management VPLS



Example

Use the following commands to configure load balancing with management VPLS



Note: the STP path costs in each peer node should be reversed.

```
(for 7210 SAS-D)config>service# vpls service-id [customer customer-id] [create][m-vpls]
[svc-sap-type {null-star | any | dot1q-preserve}] [customer-vid vlan-id]
description description-string
sap sap-id create
managed-vlan-list
    range vlan-range
stp
no shutdown
```

The following 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
-----
```

5.5.4.3 Configuring a BGP Auto-Discovery on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 are required to start using BGP AD. At the end of this section, it will be apparent that a very simple

configuration will automatically generate the required values used by BGP and LDP. In most cases, deployments will provide full mesh connectivity between all nodes across a VPLS instance. However, capabilities are available to influence the topology and build hierarchies or hub and spoke models.

BGP AD automatically creates SDP-bindings using a template to configure SDP-binding configuration parameters. The **I2-auto-bind** command initiates 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 "I2-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 I2-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 I2-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

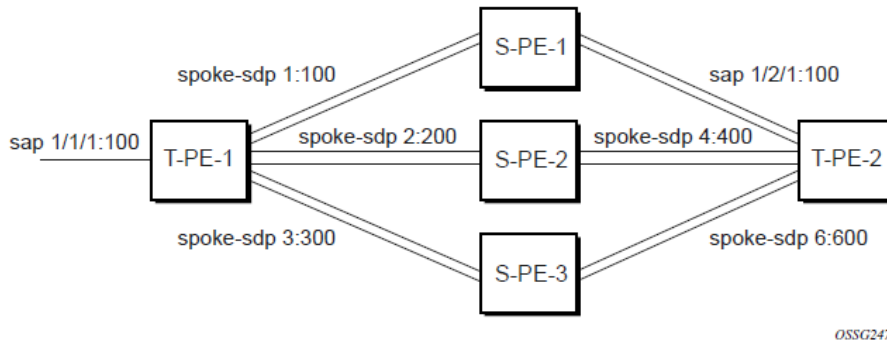
The keyword **allow-service-impact** enables service impacting changes. If this keyword is not configured, an error message is generated if the parameter changes are service impacting.

5.5.4.4 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 73: 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.

The following CLI example shows the BGP AD CLI command tree.

```

- config>service>vpls>bgp-ad
  - [no] pw-template-bind
  - [no] route-target
  - [no] shutdown
  - vpls-id
  - [no] vsi-export
  - vsi-id
  - [no] vsi-import
    
```

A helpful command displays the service information, the BGP parameters and the SDP bindings in use. When the discovery process is completed successfully each endpoint will have an entry for the service.

```

*A:H-SASK# show service l2-route-table

=====
Services: L2 Route Information - Summary
=====
Svc Id      L2-Routes (RD-Prefix)      Next Hop      Origin
           Sdp Bind Id                PW Temp Id
-----
No. of L2 Route Entries: 0
=====

=====
Services: L2 Multi-Homing Route Information - Summary
=====
Svc Id      L2-Routes (RD-Prefix)      Next Hop      SiteId      State      DF
-----
No. of L2 Multi-Homing Route Entries: 0
=====
    
```

```

=====
Services: L2 Bgp-Vpls Route Information - Summary
=====
Svc Id      L2-Routes (RD)      Next Hop      Ve-Id
           Sdp Bind Id         PW Temp Id
-----
No. of L2 Bgp-Vpls Route Entries: 0
=====
    
```

```

=====
Services: L2 Bgp-Vpws Route Information - Summary
=====
Svc Id      L2-Routes (RD)      Next Hop      Ve-Id
           Sdp Bind Id         PW Temp Id
-----
    
```

When only one of the endpoints has an entry for the service in the l2-routing-table, it is most likely a problem with the RT values used for import and export. This would most likely happen when different import and export RT values are configured using a router policy or the route-target command.

Service specific commands continue to be available to display service specific information, including status.

```

*A:H-SASK# show service sdp-using

=====
SDP Using
=====
SvcId      SdpId              Type   Far End      Opr   I.Label  E.Label
           :GlobalId
-----
1          1:1                Spok   4.4.4.4      Up    130769  130872
2          2:2                Spok   4.4.4.4      Up    130770  130873
3          3:3                Spok   4.4.4.4      Up    130771  130874
4          4:4                Spok   4.4.4.4      Up    130772  130875
5          5:5                Spok   4.4.4.4      Up    130773  130876
6          6:6                Spok   4.4.4.4      Up    130774  130877
7          7:7                Spok   4.4.4.4      Up    130775  130878
8          8:8                Spok   4.4.4.4      Up    130776  130879
9          9:9                Spok   4.4.4.4      Up    130777  130880
10         10:10              Spok   4.4.4.4      Up    130778  130881
11         11:11              Spok   4.4.4.4      Up    130779  130882
12         12:12              Spok   4.4.4.4      Up    130780  130883
    
```

BGP AD advertises the VPLS-ID in the extended community attribute, VSI-ID in the NLRI and the local PE ID in the BGP next hop. At the receiving PE, the VPLS-ID is compared against locally provisioned information to determine whether the two PEs share a common VPLS. If it is found that they do, the BGP information is used in the signaling phase.

5.5.5 BGP- AD- LDP signaling

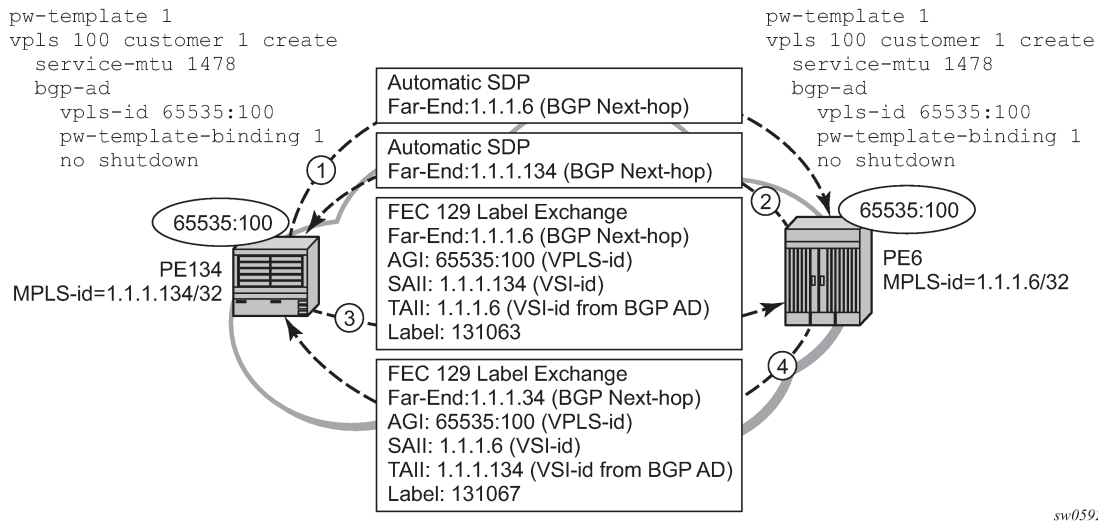
T-LDP is triggered when the VPN endpoints have been discovered using BGP. The T-LDP session between the PEs is established when one does not exist. The far-end IP address required for the T-LDP identification is gleaned from the BGP AD next hop information. The pw-template and pw-template-binding

configuration statements are used to establish the automatic SDP or to map to the appropriate SDP. The FEC129 content is built using the following values:

- AGI from the locally configured VPLS-ID.
- The SAll from the locally configured VSI-ID.
- The TAll from the VSI-ID contained in the last 4 bytes of the received BGP NLRI.

The following figure shows the different detailed phases of the LDP signaling path, post BGP AD completion. It also indicates how some fields can be auto generated when they are not specified in the configuration.

Figure 74: BGP AD triggering LDP functions



The first command shows the LDP peering relationships that have been established (as shown in the following figure). The type of adjacency is displayed in the "Adj Type" column. In this case the type is "Both" meaning link and targeted sessions have been successfully established.

Figure 75: Show router LDP session output

```
PERs6# show router ldp session
LDP Sessions
```

Peer LDP Id	Adj Type	State	Msg Sent	Msg Recv	Up Time
1.1.1.134:0	Both	Established	21482	21482	0d 15:38:44

No. of Sessions: 1

The second command shows the specific LDP service label information broken up per FEC element type, 128 or 129, basis (as shown in the following figure). The information for FEC element 129 includes the AGI, SAll and the TAll.

Figure 76: Show router LDP bindings FEC-type services

```
PERs6# show router ldp bindings fec-type services
LDP LSR ID: 1.1.1.6
```

```
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
        S - Status Signaled Up, D - Status Signaled Down
        E - Epipe Service, V - VPLS Service, M - Mirror Service
        A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
        P - Ipipe Service, C - Cpipe Service
        TLV - (Type, Length: Value)
```

```
LDP Service FEC 128 Bindings
```

Type	VCId	SvcId	SDPId	Peer	IngLbl	EgrLbl	LMTU	RMTU
No Matching Entries Found								

```
LDP Service FEC 129 Bindings
```

AGI			SAII		TAII			
Type	SvcId	SDPId	Peer	IngLbl	EgrLbl	LMTU	RMTU	
65535:100			1.1.1.6		1.1.1.134			
V-Eth	100	17406	1.1.1.134	131063U	131067S	1464	1464	
No. of FEC 129s: 1								

0989

5.5.6 BGP AD- pseudowire template

The pseudowire template is defined under the top-level service command (**config>service> pw-template**) and specifies whether to use an automatically generated SDP or manually configured SDP. It also provides the set of parameters required for establishing the pseudowire (SDP binding) as follows.

```
PERs6>config>service# pw-template 1 create
-[no] pw-template <policy-id> [use-provisioned-sdp | prefer-provisioned-sdp]
<policy-id> : [1..2147483647]
<use-provisioned-s*> : keyword
<prefer-provisioned*> : keyword

[no] accounting-pol* - Configure accounting-policy to be used
[no] auto-learn-mac* - Enable/Disable automatic update of MAC protect list
[no] block-on-peer-* - Enable/Disable block traffic on peer fault
[no] collect-stats - Enable/disable statistics collection
[no] control word - Enable/Disable the use of Control Word
[no] disable-aging - Enable/disable aging of MAC addresses
[no] disable-learn-* - Enable/disable learning of new MAC addresses
[no] discard-unknow* - Enable/disable discarding of frames with unknown source
MAC address
    egress + Spoke SDP binding egress configuration
[no] force-qinq-vc-* - Forces qinq-vc-type forwarding in the data-path
[no] force-vlan-vc-* - Forces vlan-vc-type forwarding in the data-path
[no] hash-label - Enable/disable use of hash-label
    igmp-snooping + Configure IGMP snooping parameters
    in gr ess + Spoke SDP binding ingress configuration
[no] l2pt-terminati* - Configure L2PT termination on this spoke SDP
[no] limit-mac-move - Configure mac move
[no] mac-pinning - Enable/disable MAC address pinning on this spoke SDP
[no] max-nbr-mac-ad* - Configure the maximum number of MAC entries in the FDB
from this SDP
[no] restrict-prote* - Enable/disable protected src MAC restriction
[no] sdp-exclude - Configure excluded SDP group
```

```
[no] sdp-include      - Configure included SDP group
[no] split-horizon-* + Configure a split horizon group
    stp              + Configure STP parameters
    vc-type          - Configure VC type
[no] vlan-vc-tag     - Configure VLAN VC tag
```

A **pw-template-binding** command configured within the VPLS service under the **bgp-ad** subcommand is a pointer to the pw-template that should be used. If a VPLS service does not specify an import-rt list, then that binding applies to all route targets accepted by that VPLS. The **pw-template-bind** command can select a different template on a per import-rt basis. It is also possible to specify specific pw-templates for some route targets with a VPLS service and use the single **pw-template-binding** command to address all unspecified but accepted imported targets.

The following figure shows the PW-template-binding CLI syntax.

Figure 77: PW-template-binding CLI syntax

```
PERs6>config>service>vpls>bgp-ad# pw-template-binding
- pw-template-binding <policy-id> [split-horizon-group <group-name>] [import-rt
  {ext-community, ... (upto 5 max)}]
- no pw-template-binding <policy-id>

<policy-id>          : [1..2147483647]
<group-name>        : [32 chars max]
<ext-community>     : target:{<ip-addr:comm-val>|<as-number:ext-comm-val>}
                    ip-addr      - a.b.c.d
                    comm-val     - [0..65535]
                    as-number    - [1..65535]
                    ext-comm-val - [0..4294967295]
```

0990

It is important understand the significance of the split-horizon-group used by the pw-template. Traditionally, when a VPLS instance was manually created using mesh-sdp bindings, these were automatically placed in a common split-horizon-group to prevent forwarding between the pseudowire in the VPLS instances. This prevents loops that would have otherwise occurred in the Layer 2 service. When automatically discovering VPLS service using BGP AD the service provider has the option of associating the auto-discovered pseudowire with a split-horizon group to control the forwarding between pseudowires.

5.5.7 Configuring a VPLS management interface

Use the following syntax to create a VPLS management interface.

```
config>service>vpls# interface ip-int-name
  address ip-address[/mask] [netmask]
  arp-timeout seconds
  description description-string
  mac ieee-address
  no shutdown
  static-arp ip-address ieee-address
```

The following is a sample configuration output.

```
A:ALA-49>config>service>vpls>interface# info detail
```

```
-----  
no description  
mac 14:31:ff:00:00:00  
address 10.231.10.10/24  
no arp-timeout  
no shutdown  
-----  
A:ALA-49>config>service>vpls>interface#
```

5.5.7.1 Creating a management VPLS for SAP protection using MPLS uplinks

This section provides a brief overview of the tasks that must be performed to configure a management VPLS for SAP protection and provides the CLI commands. The following tasks should be performed on both nodes providing the protected VPLS service.

Before configuring a management VPLS, first read [Configuring VPLS redundancy](#) for an introduction to the concept of management VPLS and SAP redundancy:

1. Create an SDP to the peer node.
2. Create a management VPLS.
3. Define a SAP in the m-vpls on the port 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.
4. Optionally modify STP parameters for load balancing (see [Configuring Load Balancing with Management VPLS](#) on page 414).
5. Create an SDP in the m-vpls using the SDP defined in Step 1. Ensure that this SDP runs over a protected LSP.
6. Enable the management VPLS service and verify that it is operationally up.
7. Create a list of VLANs on the port managed by this management VPLS.
8. Create one or more user VPLS services with SAPs on VLANs in the range defined by Step 6.

```
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 is a sample VPLS configuration output.

```
*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#
```

5.5.7.2 Creating a management VPLS for spoke-SDP protection, using MPLS uplinks (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

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. The following tasks should be performed on all four nodes providing the protected VPLS service.

Before configuring a management VPLS, please first read [Configuring VPLS redundancy](#) [Configuring VPLS redundancy](#) for an introduction to the concept of management VPLS and spoke-SDP redundancy:

1. Create an SDP to the local peer node (node ALA-A2 in the following sample).
2. Create an SDP to the remote peer node (node ALA-B1 in the following sample).
3. Create a management VPLS.
4. Create a spoke-SDP in the m-vpls using the SDP defined in Step 1. Ensure that this mesh-spoke-SDP runs over a protected LSP (see following note).
5. Enable the management VPLS service and verify that it is operationally up.
6. Create a spoke-SDP in the m-vpls using the SDP defined in Step 2. Optionally, modify STP parameters for load balancing.
7. Create one or more user VPLS services with spoke-SDPs on the tunnel SDP defined by Step 2.

As long as the user spoke-SDPs created in step 7 are in this same tunnel SDP with the management spoke-SDP created in step 6, the management VPLS will protect them.

Use the following 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
```

The following is a sample VPLS configuration output.

```
*A:ALA-A1>config>service# info
-----
...
sdp 1 mpls create
shutdown
no description
signaling tldp
no ldp
no sr-isis
no sr-ospf
no bgp-tunnel
```

```

no path-mtu
no adv-mtu-override
keep-alive
shutdown
hello-time 10
hold-down-time 10
max-drop-count 3
timeout 5
no message-length
exit

...
-----
*A:ALA-A1>config>service#
    
```

5.5.7.3 Configuring load balancing with management VPLS using MPLS uplinks (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

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.

Example

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
    
```

Example

The following outputs are sample configurations for load balancing across two protected VPLS spoke-SDPs.

The following is a sample configuration output on ALA-A (7210 SAS).

```

# MVPLS 100 configs
*A:ALA-A# configure service vpls 100
*A:ALA-A>config>service>vpls# info
-----
description "Default tls description for service id 100"
stp
no shutdown
exit
sap lag-3:100 create
description "Default sap description for service id 100"
managed-vlan-list
    
```

```
range 101-110
exit
exit
spoke-sdp 1201:100 create
stp
path-cost 100
exit
exit
spoke-sdp 1401:100 create
exit
no shutdown
-----
*A:ALA-A>config>service>vpls#

# UVPLS 101 configs
*A:ALA-A>config>service# vpls 101
*A:ALA-A>config>service>vpls# info
-----
description "Default tls description for service id 101"
sap lag-3:101 create
description "Default sap description for service id 101"
exit
spoke-sdp 1201:101 create
exit
spoke-sdp 1401:101 create
exit
no shutdown
-----
*A:ALA-A>config>service>vpls#

# MVPLS 200 configs
*A:ALA-A# configure service vpls 200
*A:ALA-A>config>service>vpls# info
-----
description "Default tls description for service id 200"
stp
no shutdown
exit
sap lag-3:200 create
description "Default sap description for service id 200"
managed-vlan-list
range 201-210
exit
exit
spoke-sdp 1202:200 create
exit
spoke-sdp 1402:200 create
stp
path-cost 100
exit
exit
no shutdown
-----
*A:ALA-A>config>service>vpls#

# UVPLS 201 configs
*A:ALA-A>config>service# vpls 201
*A:ALA-A>config>service>vpls# info
-----
description "Default tls description for service id 201"
sap lag-3:201 create
description "Default sap description for service id 201"
exit
spoke-sdp 1202:201 create
```

```
exit
spoke-sdp 1402:201 create
exit
no shutdown
-----
*A:ALA-A>config>service>vpls# exit all
```

The 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#
```

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#
```

5.5.7.4 Configuring selective MAC Flush, when using MPLS uplinks (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Example

Use the following syntax to enable selective MAC Flush in a VPLS.

```
config>service# vpls service-id
send-flush-on-failure
```

Example

Use the following syntax to disable selective MAC Flush in a VPLS.


```
config>service# vpls service-id
no send-flush-on-failure
```

5.5.7.5 Configuring AS pseudowire in VPLS when using MPLS uplinks (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Example

In the following figure, pseudowire is configured on MTU.

Figure 78: Sample AS pseudowire



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#
```

5.6 Service management tasks

This section describes the service management tasks.

5.6.1 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#
```

5.6.2 Modifying management VPLS parameters

To modify the range of VLANs on an access port managed by an existing management VPLS, first the new range should be entered and afterwards the old range removed. If the old range is removed before a new range is defined, all customer VPLS services in the old range will become unprotected and may be disabled.

```
config>service# vpls service-id  
  sap sap-id  
  managed-vlan-list  
  [no] range vlan-range
```

5.6.3 Deleting a management VPLS

As with normal VPLS service, a management VPLS cannot be deleted until SAPs are unbound (deleted), interfaces are shutdown, and the service is shutdown on the service level.

Use the following syntax to delete a management VPLS service.

```
config>service  
  [no] vpls service-id  
  shutdown  
  [no] sap sap-id  
  shutdown
```

5.6.4 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 desired, first unmanage the user VPLS service by removing them from the managed-vlan-list.

```
config>service  
  vpls service-id  
  shutdown
```

Example: Disabling a management VPLS

```
config>service# vpls 1  
  config>service>vpls# shutdown  
  config>service>vpls# exit
```


5.6.5 Deleting a VPLS service

A VPLS service cannot be deleted until SAPs are unbound (deleted), interfaces are shutdown, and the service is shutdown on the service level.

Use the following syntax to delete a VPLS service.

```
config>service
  [no] vpls service-id
  shutdown
  sap sap-id
  no sap sap-id
  shutdown
```

5.6.6 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: Disabling a VPLS service

```
config>service# vpls 1
config>service>vpls# shutdown
config>service>vpls# exit
```

5.6.7 Re-enabling a VPLS service

Use the following syntax to re-enable a VPLS service that was shutdown.

```
config>service> vpls service-id [no] shutdown
```

Example

```
config>service# vpls 1
config>service>vpls# no shutdown
config>service>vpls# exit
```

5.7 VPLS services command reference

5.7.1 Command hierarchies

- [VPLS service configuration commands](#)
- [Global commands for 7210 SAS-D and 7210 SAS-K 2F1C2T](#)
- [Global commands for 7210 SAS-Dxp](#)
- [Global commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)

- VPLS service xSTP commands
- VPLS service SAP xSTP commands
- VPLS service SAP DHCP snooping commands
- SAP commands
- VPLS SAP QoS and filter commands for 7210 SAS-D and 7210 SAS-Dxp
- VPLS SAP QoS and filter commands for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C
- VPLS service SAP IGMP snooping and MVR commands
- VPLS SAP meter override commands
- VPLS SAP statistics commands for 7210 SAS-D
- VPLS mesh SDP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
- VPLS spoke-SDP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
- Routed VPLS commands for 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C
- Show commands
- Clear commands
- Debug commands

5.7.1.1 VPLS service configuration commands

5.7.1.1.1 Global commands for 7210 SAS-D and 7210 SAS-K 2F1C2T

```
config
- service
- vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-star |
dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id] (7210 SAS-D)
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [customer-
vid vlan-id] [svc-sap-type {null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls]
(7210 SAS-Dxp)
- vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {any | dot1q-
range}] [r-vpls] (7210 SAS-K 2F1C2T)
- no vpls service-id
- description description-string
- no description
- [no] disable-aging
- [no] disable-learning
- [no] discard-unknown
- [no] fdb-table-high-wmark high-water-mark
- [no] fdb-table-low-wmark low-water-mark
- fdb-table-size table-size
- no fdb-table-size [table-size]
- local-age aging-timer
- no local-age
- [no] mac-move
- move-frequency frequency
- no move-frequency
- retry-timeout timeout
- no retry-timeout
```

```
- [no] shutdown
- remote-age aging-timer
- no remote-age
- service-mtu octets
- no service-mtu
```

5.7.1.1.2 Global commands for 7210 SAS-Dxp

```
config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [customer-
vid vlan-id] [svc-sap-type {null-star | dot1q-preserve | any}] [b-vpls | i-vpls | r-vpls]
- no vpls service-id
- allow-ip-int-bind
- no allow-ip-int-bind
- description description-string
- no description
- [no] disable-aging
- [no] disable-learning
- [no] discard-unknown
- fdb-table-high-wmark high-water-mark
- no fdb-table-high-wmark [high-water-mark]
- fdb-table-low-wmark low-water-mark
- no fdb-table-low-wmark [low-water-mark]
- fdb-table-size table-size
- no fdb-table-size [table-size]
- igmp-snooping
- local-age aging-timer
- no local-age
- [no] mac-move
- move-frequency frequency
- no move-frequency
- retry-timeout timeout
- no retry-timeout
- [no] shutdown
- 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
- remote-age aging-timer
- no remote-age
- sap sap-id [create][capture-sap] [eth-ring ring-index] [g8032-shg-enable]
- no sap sap-id
- service-name service-name
- no service-name
- [no] shutdown
- stp
```

5.7.1.1.3 Global commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
config
- service
- vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {any | dot1q-
range}] [r-vpls]
- no vpls service-id
- bgp
```

```

- pw-template-binding policy-id [split-horizon-group group-name] [import-rt
{ext-community... (up to 5 max)}}
- no pw-template-binding policy-id
community}}}
- 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
- [no] fdb-table-high-wmark high-water-mark
- [no] fdb-table-low-wmark low-water-mark
- fdb-table-size table-size
- no fdb-table-size [table-size]
- local-age aging-timer
- no local-age
- [no] mac-move
  - move-frequency frequency
  - no move-frequency
  - retry-timeout timeout
  - no retry-timeout
  - [no] shutdown
- [no] propagate-mac-flush
- remote-age aging-timer
- no remote-age
- [no] send-flush-on-failure
- service-mtu octets
- no service-mtu
- no service-name
- [no] shutdown
- split-horizon-group group-name [create]
- description description-string
- no description
    
```

5.7.1.2 VPLS service xSTP commands

```

config
- service
    
```

```

- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q | dot1q-preserve}] [customer-vid vlan-id]
- vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-star |
dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
- vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {any | dot1q-
range}] [r-vpls]
- 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}
    - 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

```

5.7.1.3 VPLS service SAP xSTP commands

```

config
- service
  - vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q | dot1q-preserve}] [customer-vid vlan-id]
  - vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-star |
dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
  - vpls service-id [customer customer-id] [create] [svc-sap-type {any | dot1q-range}]
[r-vpls]
  - no vpls service-id
    - sap sap-id [split-horizon-group group-name] [eth-ring ring-index]
    - no sap sap-id
      - stp
        - [no] auto-edge
        - [no] edge-port
        - link-type {pt-pt | shared}
        - no link-type [pt-pt | shared]
        - mst-instance mst-inst-number
          - mst-path-cost inst-path-cost
          - no mst-path-cost
          - mst-port-priority stp-priority
          - no mst-port-priority
        - path-cost sap-path-cost
        - no path-cost
        - [no] port-num virtual-port-number
        - priority stp-priority
        - no priority
        - no root-guard

```

- **root-guard**
- [no] **shutdown**
- **managed-vlan-list**
- [no] **default-sap**
- [no] **range** *vlan-range*

5.7.1.4 VPLS service SAP DHCP snooping commands

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q | dot1q-preserve}] [customer-vid vlan-id]
- vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-star |
dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
- vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {any | dot1q-
range}] [r-vpls]
- no vpls service-id
- sap sap-id [split-horizon-group group-name] [eth-ring ring-index]
- no sap sap-id
- dhcp
- description description-string
- no description
- [no] option
- action [dhcp-action]
- no action
- [no] circuit-id [ascii-tuple | vlan-ascii-tuple]
- [no] remote-id [mac | string string]
- [no] vendor-specific-option
- [no] client-mac-address
- [no] sap-id
- [no] service-id
- string text
- no string
- [no] system-id
- [no] shutdown
- [no] snoop
    
```

5.7.1.5 SAP commands

```

config
- service
- vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [svc-sap-type
{null-star | dot1q | dot1q-preserve}] [customer-vid vlan-id]
- vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {null-star |
dot1q-preserve | dot1q-range | any}] [customer-vid vlan-id]
- vpls service-id [customer customer-id] [create] [m-vpls] [svc-sap-type {any | dot1q-
range}] [r-vpls]
- no vpls service-id
- sap sap-id [split-horizon-group group-name] [eth-ring ring-index]
- no sap sap-id
- accounting-policy acct-policy-id
- no accounting-policy
- bpdu-translation {auto | pvst | stp}
- no bpdu-translation
- [no] collect-stats
- description description-string
- no description
- [no] disable-aging
    
```

```

- [no] disable-learning
- [no] discard-unknown-source
- eth-cfm
  - mep mep-id domain md-index association ma-index [direction {up | down}]
  - no mep mep-id domain md-index association ma-index
    - [no] ais-enable
      - client-meg-level [level [level...]]
      - no client-meg-level
      - [no] description
      - interval {1| 60}
      - no interval
      - priority priority-value
      - no priority
      - no send-ais-on-port-down
      - send-ais-on-port-down
    - [no] ccm-enable
    - ccm-ltm-priority priority
    - no ccm-ltm-priority
    - description description-string
    - no description
    - [no] eth-test-enable
      - bit-error-threshold bit-errors
      - test-pattern {all-zeros | all-ones} [crc-enable]
      - no test-pattern
    - mac-address {allDef | macRemErrXcon | remErrXcon | errXcon | xcon |
noXcon}
      - mac-address mac-address
      - no mac-address
      - one-way-delay-threshold seconds
      - [no] shutdown
- l2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp]
- no l2pt-termination
- limit-mac-move [blockable | non-blockable]
- no limit-mac-move
- [no] mac-pinning
- mac-swap-enable
- no mac-swap-enable
- max-nbr-mac-addr table-size
- no max-nbr-mac-addr
- tod-suite tod-suite-name
- no tod-suite
    
```

5.7.1.6 VPLS SAP QoS and filter commands for 7210 SAS-D and 7210 SAS-Dxp

```

config
- service
  - vpls
  - no vpls
    - sap sap-id [create] [eth-ring ring-index]
    - no sap sap-id
      - egress
        - aggregate-meter-rate rate-in-kbps [burst burst-in-kbits] [enable-stats]
        - 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
```

5.7.1.7 VPLS SAP QoS and filter commands for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C

```
config
- service
- vpls
- no vpls
- sap sap-id [create] [split-horizon-group group-name]
- no sap sap-id
- egress
- agg-shaper-rate agg-rate
- no agg-shaper-rate
- dot1p-inner dot1p-inner
- no dot1p-inner
- no dot1p-outer
- dot1p-outer dot1p-outer
- filter [ip ip-filter-id]
- filter [ipv6 ipv6 -filter-id]
- filter [mac mac-filter-id]
- no filter [ip ip-filter-id] [ipv6 ipv6 -filter-id] [mac mac-filter-id]
- qos policy-id
- no qos
- ingress
- agg-shaper-rate agg-rate
- no agg-shaper-rate
- aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
- no aggregate-meter-rate
- filter [ip ip-filter-id]
- filter [ ipv6 ipv6-filter-id]
- filter [mac mac-filter-id]
- no filter [ip ip-filter-id] [ ipv6 ipv6-filter-id] [mac mac-filter-id]
- qos policy-id
- no qos
```

5.7.1.8 VPLS service SAP IGMP snooping and MVR commands

```
config
- service
- [no] vpls
- igmp-snooping
- mvr
- description description-string
- no description
- group-policy policy-name
- no group-policy
- no shutdown
- shutdown
- query-interval seconds
- no query-interval
- no query-src-ip
- query-src-ip ip-address
```



```
- no report-src-ip
- report-src-ip ip-address
- robust-count count
- no robust-count
- no shutdown
- shutdown
- sap sap-id [split-horizon-group group-name] [eth-ring ring-index]
- no sap sap-id
  - igmp-snooping
    - disable-router-alert-check
    - no disable-router-alert-check
    - [no] fast-leave
    - import policy-name
    - no import
    - last-member-query-interval interval
    - no last-member-query-interval
    - max-num-groups max-num-groups
    - no max-num-groups
    - max-num-sources max-num-sources
    - no max-num-sources
    - [no] mrouter-port
    - mvr
      - from-vpls service-id
      - no from-vpls
      - to-sap sap-id
      - no to-sap
    - query-interval seconds
    - no query-interval
    - query-response-interval seconds
    - no query-response-interval
    - robust-count count
    - no robust-count
    - [no] send-queries
    - static
      - [no] group group-address
      - [no] source (ip-address applicable only in access-uplink mode)
      - [no] starg
    - version version
    - no version
  - mfib-table-high-wmark high-water-mark
  - no mfib-table-high-wmark
  - mfib-table-low-wmark low-water-mark
  - no mfib-table-low-wmark
  - mfib-table-size table-size
  - no mfib-table-size
```

5.7.1.9 VPLS SAP meter override commands

```
config
- service
  - vpls
  - no vpls
    - sap sap-id [split-horizon-group group-name] [eth-ring ring-index ]
    - no sap sap-id
      - ingress
        - meter-override
          - meter meter-id [create]
          - no meter meter-id
            - adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
            - cbs size [kbits | bytes | kbytes]
            - no cbs
```

```
- mbs size [kbits | bytes | kbytes]
- no mbs
- mode
- no mode
- rate cir cir-rate [pir pir-rate]
```

5.7.1.10 VPLS SAP statistics commands for 7210 SAS-D

```
config
- service
  - vpls
  - no vpls
    - sap sap-id [split-horizon-group group-name] [eth-ring ring-index] [create]
    - no sap sap-id
      - statistics
        - egress
          - [no] forwarded-count
        - ingress
          - counter-mode {in-out-profile-count | forward-drop-count}
          - [no] received-count
```

5.7.1.11 VPLS mesh SDP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
config
- service
  - [no] vpls
    - mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
    - no mesh-sdp sdp-id[:vc-id]
      - accounting-policy acct-policy-id
      - no accounting-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
          - mac-address {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 [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] starg
  - version version
  - no version
- ingress
  - vc-label egress-vc-label
- [no] mac-pinning
- [no] static-mac ieee-address
- [no] shutdown
- suppress-vc-vlan
- no suppress-vc-vlan
- vlan-vc-tag 0..4094
- no vlan-vc-tag [0..4094]
- dhcp
  - description description-string
  - no description
  - no snoop
  - snoop
```

5.7.1.12 VPLS spoke-SDP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```
config
- service
  - vpls
  - no vpls
    - spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] [split-horizon-
group group-name] endpoint endpoint-name
    - no spoke-sdp sdp-id[:vc-id]
      - accounting-policy acct-policy-id
      - no accounting-policy
      - adv-service-mtu octets
      - no adv-service-mtu
      - [no] block-on-mesh-failure
      - bpdu-translation {auto | pvst | stp}
      - no bpdu-translation
      - [no] collect-stats
      - [no] control-word
      - description description-string
      - no description
```

```

- [no] disable-aging
- [no] disable-learning
- [no] discard-unknown-source
- dhcp
  - description description-string
  - no description
  - no snoop
  - snoop
- eth-cfm
  - mep mep-id domain md-index association ma-index [direction {up}{down}]
  - no mep mep-id domain md-index association ma-index
    - [no] ais-enable
    - client-meg-level [[level [level...]]
    - no client-meg-level
    - interval {1 | 60}
    - no interval
    - priority priority-value
    - no priority
  - [no] ccm-enable
  - ccm-ltm-priority priority
  - no ccm-ltm-priority
  - [no] description description string
  - no eth-test-enable
    - bit-error-threshold bit-errors
    - test-pattern {all-zeros | all-ones} [crc-enable]
  - no test-pattern mac-address {allDef | macRemErrXcon | remErrXcon | err
Xcon | 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
- 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] starg
  - version version
  - no version
- [no] ignore-standby-signaling
- ingress
  - vc-label egress-vc-label
  - no vc-label [egress-vc-label]

```

```

- [no] l2pt-termination [cdp] [dtp] [pagp] [stp] [udld] [vtp]
- limit-mac-move [blockable | non-blockable]
- no limit-mac-move
- [no] mac-pinning
- max-nbr-mac-addr table-size
- no max-nbr-mac-addr
- precedence precedence-value | primary
- no precedence
- [no] shutdown
- [no] static-mac ieee-address
- statistics
  - ingress
  - [no] drop-count-extra-vlan-tag-pkts
- stp
  - [no] auto-edge
  - [no] edge-port
  - link-type {pt-pt | shared}
  - no link-type [pt-pt | shared]
  - path-cost sap-path-cost
  - no path-cost
  - [no] port-num virtual-port-number
  - priority stp-priority
  - no priority
  - no root-guard
  - root-guard
  - [no] shutdown
- suppress-vc-vlan
- no suppress-vc-vlan
- vlan-vc-tag vlan-id
- no vlan-vc-tag [vlan-id]
    
```

5.7.1.13 Routed VPLS commands for 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

```

config
- service
  - vpls service-id [customer customer-id] [vpn vpn-id] [r-vpls] [create]
  - [no] allow-ip-int-bind
    
```

5.7.1.14 Show commands

```

show
- service
  - fdb-info
  - fdb-mac ieee-address [expiry]
  - id service-id
    - all
    - arp [ip-address] | [mac ieee-address] | [sap sap-id] | [interface ip-int-name] |
[summary]
  - base
  - bgp-ad
  - dhcp
    - statistics [sap sap-id] [interface interface-name]
    - summary [interface interface-name | saps]
  - fdb [sap sap-id] | [mac ieee-address | endpoint endpoint | detail] [expiry]
  - igmp-snooping
    - all
    
```

```
- base
- mroute [detail]
- mvr
- port-db sap sap-id [detail]
- port-db sap sap-id group grp-address
- proxy-db [detail]
- proxy-db [group grp-ip-address]
- querier
- static [sap sap-id]
- statistics[sap sap-id | sdp sdp-id:vc-id]
- l2pt disabled
- l2pt [detail]
- mac-move
- mrib [brief]
- mrib [group grp-address | mstp-configuration]
- sap [sap-id [detail | stp]]
- stp [detail]
- sap-using [sap sap-id]
- sap-using [ingress | egress] filter filter-id
- sap-using [ingress] qos-policy qos-policy-id
- service-using [vpls]
```

5.7.1.15 Clear commands

```
clear
- service
  - id service-id
    - fdb {all | mac ieee-address | sap sap-id |}
    - igmp-snooping
      - port-db sap sap-id [group grp-address]
      - querier
    - statistics [all | sap sap-id | sdp sdp-id:vc-id]
    - stp
      - detected-protocols [all | sap sap-id]
  - statistics
    - id service-id
      - counters
      - stp
    - sap sap-id {all | counters | stp}
```

5.7.1.16 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
```

5.7.2 Command descriptions

5.7.2.1 VPLS service configuration commands



Note:

References in this section to spoke-SDPs, mesh SDPs, or VPRN only apply to the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

5.7.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

config>service>vpls

config>service>vpls>igmp-snooping

config>service>vpls>sap

config>service>vpls>sap>stp

config>service>vpls>mesh-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>stp

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.

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.

description

Syntax

description *description-string*

no description

Context

```
config>service>vpls
config>service>vpls>split-horizon-group
config>service>vpls>igmp-snooping>mvr
config>service>vpls>sap
config>service>vpls>mesh-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
config>service>pw-template>split-horizon-group (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
```

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 **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

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.

5.7.2.1.2 DHCP commands

```
dhcp
```

Syntax

```
dhcp
```

Context

```
config>service>vpls>sap
config>service>vpls>mesh-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure DHCP parameters.

option

Syntax

[no] option

Context

config>service>vpls>sap>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables DHCP Option 82 (Relay Agent Information Option) parameter processing and enters the context for configuring Option 82 suboptions.

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

Default

no option

action

Syntax

action [*dhcp-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

Description

This command configures the Relay Agent Information Option (Option 82) processing.

By default, existing information is kept intact.

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

Parameters

dhcp-action

Specifies the DHCP option action.

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 more information. In the downstream direction, the Option 82 field is not stripped and is sent on toward the client.

The behavior is slightly different in 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 done only when the incoming message already has 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>vpls>sap>dhcp>option

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** suboption of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by "|".

To send a tuple in the circuit ID, the **action replace** command must be configured in the same context.

If disabled, the **circuit-id** suboption of the DHCP packet is left empty.

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

Default

circuit-id

Parameters

ascii-tuple

Specifies that the ASCII-encoded concatenated tuple will be used, which consists of the *access-node-identifier*, *service-id*, and *interface-name*, separated by "|".

ifindex

Specifies that the interface index will be used. (The If Index of a router interface can be displayed using the **show router interface detail** command.) This parameter is not supported on 7210 SAS-Dxp.

sap-id

Specifies that the SAP identifier will be used. This parameter is not supported on 7210 SAS-Dxp.

vlan-ascii-tuple

Specifies that the format will include VLAN ID and dot1p bits in addition to what is already included in *ascii-tuple*. The format is supported on dot1q-encapsulated ports only. Therefore, when the option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

remote-id

Syntax

```
[no] remote-id [mac | string string]
```

Context

```
config>service>vpls>sap>dhcp>option
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies that information that goes into the remote ID suboption in the DHCP relay packet. If disabled, the **remote-id** suboption of the DHCP packet is left empty. The **no** form of this command reverts to the default value.

Default

```
no remote-id
```

Parameters

mac

Keyword to specify that the MAC address of the remote end is encoded in the suboption.

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

Description

This command configures the vendor-specific suboption of the DHCP relay packet.

client-mac-address

Syntax

[no] client-mac-address

Context

config>service>vpls>sap>dhcp>option>vendor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the sending of the MAC address in the vendor-specific suboption of the DHCP relay packet.

The **no** form of this command disables the sending of the MAC address in the vendor-specific suboption of the DHCP relay packet.

sap-id

Syntax

[no] sap-id

Context

config>service>vpls>sap>dhcp>option>vendor

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the sending of 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

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the sending of 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

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Specifies a string that 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

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies whether the system ID is encoded in the vendor-specific suboption of Option 82.

snoop

Syntax

[no] snoop

Context

config>service>vpls>sap>dhcp

config>service>vpls>mesh-sdp>dhcp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>spoke-sdp>dhcp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables DHCP snooping for DHCP messages on the SAP. Enabling DHCP snooping on VPLS interfaces (SAPs) is required for DHCP messages where Option 82 information is to be inserted.

This includes interfaces in the path to receive messages from either DHCP servers or from subscribers.

The **no** form of this command disables DHCP snooping on the specified VPLS SAP.

Default

no snoop

5.7.2.1.3 VPLS service commands

vpls

Syntax

vpls *service-id* [**customer** *customer-id*] [**create**] [**m-vpls**] [**svc-sap-type** {**null-star** | **dot1q-preserve** | **dot1q-range** | **any**}] [**customer-vid** *vlan-id*]

vpls *service-id* [**customer** *customer-id*] [**create**] [**m-vpls**] [**customer-vid** *vlan-id*] [**svc-sap-type** {**null-star** | **dot1q-preserve**}]

vpls *service-id* [**customer** *customer-id*] [**create**] [**m-vpls**] [**svc-sap-type** {**any** | **dot1q-range**}] [**r-vpls**]

vpls *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**m-vpls**] [**customer-vid** *vlan-id*] [**svc-sap-type** {**null-star** | **dot1q-preserve** | **any**}] [**b-vpls** | **i-vpls** | **r-vpls**]

no vpls *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 LAN Services (VPLS) instance. The **vpls** command is used to create or maintain a VPLS service. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites together acting like a zero-hop, Layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context. When a service is created, the **customer** keyword and *customer-id* must be specified and must 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 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 VPLS service may be created for a single customer ID.

By default, no VPLS instances exist until they are explicitly created.

- 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T support local service.
- 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C support both local and distributed service.

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs defined within the service ID have been shutdown and deleted, and the service has been shut down.

Parameters

service-id

Specifies the unique service 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.

create

This keyword is mandatory while creating a VPLS service. Keyword used to create the service instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

customer-vid *vlan-id*

Specifies the dot1q VLAN ID for creating the local dot1q SAP for **svc-sap-type dot1q-preserve**.

Values 1 to 4094

svc-sap-type

Keyword to specify the type of access SAPs and access-uplink SAPs allowed in the service.

Values **dot1q-range** — Keyword to specify that the access SAP in the service can use VLAN ranges as the SAP tags. The VLAN ranges are configured using the CLI **configure connection-profile** command. On ingress of the access dot1q SAP using VLAN ranges, the received tag on the SAP is preserved. A VPLS service with **svc-sap-type** set to **dot1q-range** can be used for Epipe emulation with G.8032 for protection. For more information about the capabilities and restrictions, see [Epipe emulation using dot1q VLAN range SAP in VPLS with G.8032](#). This parameter is supported on the 7210 SAS-D, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

dot1q-preserve — Specifies that the allowed SAP in the service are dot1q. The dot1q ID is not stripped after packets match the SAP. This parameter is supported only on 7210 SAS-D and 7210 SAS-Dxp.

null-star — Specifies that the allowed SAP in the service can be null SAPs, dot1q default, Q.* SAP, 0.* SAP or default QinQ SAP. This parameter is supported only on 7210 SAS-D and 7210 SAS-Dxp.

any — Specifies that the service can be configured with null, dot1q and QinQ SAPs. This parameter is supported only on 210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

Default any

r-vpls

Allows this VPLS instance to be associated with an IP interface to provide R-VPLS functionality. This parameter is supported only on the 7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C.

bgp

Syntax

bgp

Context

config>service>vpls

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure the BGP related parameters to BGP AD.

block-on-mesh-failure

Syntax

[no] block-on-mesh-failure

Context

config>service>vpls>spoke-sdp

config>service>vpls>endpoint

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables blocking after all configured SDPs or endpoints are in operationally down state.

Blocking brings the entity to an operationally down state. This event is signaled to corresponding T-LDP peer by withdrawing the 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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the translation of bridge protocol data units (BPDUs) to a specific format, meaning that all BPDUs transmitted on a specific SAP or spoke-SDP will have a specified format.

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

Default

no bpdu-translation

Parameters

auto

Keyword to specify that the appropriate format will be detected automatically, based on type of BPDUs received on such port.

pvst

Keyword to specify the BPDU format as PVST. The correct VLAN tag is included in the payload (depending on encapsulation value of outgoing SAP).

stp

Keyword to specify the BPDU format as STP.

l2pt-termination

Syntax

l2pt-termination [**cdp**] [**dtp**] [**pagp**] [**stp**] [**udld**] [**vtp**]

no l2pt-termination

Context

config>service>vpls>sap

config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command enables Layer 2 Protocol Tunneling (L2PT) termination on a specific SAP. L2PT termination is supported for STP, CDP, DTP, PAGP, UDLD, and VTP PDUs.

This command can be enabled only if STP is disabled in the context of the specific VPLS service.

Default

no l2pt-termination

Parameters

cdp

Keyword that specifies the Cisco discovery protocol.

dtp

Keyword that specifies the dynamic trunking protocol.

pagp

Keyword that specifies the port aggregation protocol.

stp

Keyword that specifies all spanning tree protocols: stp, rstp, mstp, pvst (default).

udld

Keyword that specifies unidirectional link detection.

vtp

Keyword that specifies the VLAN trunking protocol.

disable-aging

Syntax

[no] disable-aging

Context

config>service>vpls

config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command disables MAC address aging across a VPLS service or on a VPLS service SAP.

Similar to a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In a VPLS service instance, the local age timer is applicable to both the local learned and remote learned MAC entries in the VPLS forwarding database (FDB). The `disable-aging` command at the service level turns off aging for local and remote learned MAC addresses.

When **no disable-aging** is specified for a VPLS, it is possible to disable aging for specific SAPs by entering the **disable-aging** command at the appropriate level.

When the **disable-aging** command is entered at the VPLS level, the **disable-aging** state of individual SAPs 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>vpls>sap
```

```
config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command disables learning of new MAC addresses in the VPLS FDB for the service instance.

When **disable-learning** is enabled, new source MAC addresses are not entered in the VPLS service forwarding database.

When **disable-learning** is disabled, new source MAC addresses are learned and entered into the VPLS forwarding database.

This command 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

discard-unknown

Syntax

[no] discard-unknown

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables, at the VPLS level, packets with unknown destination MAC address to be dropped instead of flooded (even when configured FIB size limits for VPLS or SAP are not yet reached). By default, packets with unknown destination MAC addresses are flooded.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a service endpoint.

Parameters

endpoint-name

Specifies an endpoint name, up to 32 characters.

create

Keyword that is mandatory while creating a service endpoint.

description

Syntax

description *description-string*

no description

Context

config>service>vpls>endpoint

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

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

This 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

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, etc.), the entire string must be enclosed within double quotes.

ignore-standby-signaling

Syntax

[no] ignore-standby-signaling

Context

config>service>vpls>endpoint

config>service>vpls>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

When this command is enabled, the node ignores the standby-bit received from TLDP peers for the specific spoke-SDP and performs internal tasks without taking it into account.

This command is available at the endpoint level and the spoke-SDP level. If the spoke-SDP is part of the explicit endpoint, it is not possible to change this setting at the spoke-SDP level. The existing spoke-SDP

will become part of the explicit endpoint only if the setting is not conflicting. The newly created spoke-SDP, which is part of the specific explicit endpoint, inherits this setting from the endpoint configuration.

Default

disabled

split-horizon-group

Syntax

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

Context

config>service>vpls

config>service>pw-template (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command creates a 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.

An 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 instance. The same group name can be reused in different VPLS instances.



Note:

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.

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

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

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

Description

This command specifies the maximum number of MAC entries in the FDB for the VPLS instance on this node.

This command 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 the table size to the default value.

Default

250

Parameters

table-size

Specifies the maximum number of MAC entries in the FDB.

Values	1 to 15359 (7210 SAS-Dxp)
	1 to 30719 (7210 SAS-D, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C)

bgp-ad

Syntax

[no] **bgp-ad**

Context

config>service>vpls

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures BGP auto-discovery.

vpls-id

Syntax

vpls-id *vpls-id*

Context

```
config>service>vpls>bgp-ad
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the VPLS ID component that will be signaled in one of the extended community attributes (ext-comm). Values and format (6 bytes, other 2 bytes of type-subtype will be automatically generated).

Parameters

vpls-id

Specifies a globally unique VPLS ID for BGP auto-discovery in this VPLS service.

Values vpls-id : <ip-addr:comm-val>|<as-number:ext-comm-val>
ip-addr a.b.c.d
comm-val 0 to 65535

vsi-export

Syntax

```
vsi-export policy-name [policy-name...(up to 5 max)]  
no vsi-export
```

Context

```
config>service>vpls>bgp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the name of the VSI export policies to be used for BGP auto-discovery, when this command 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
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the name of the VSI import policies used for BGP auto-discovery, when this command is configured in the VPLS service. If multiple policy names are configured, the policies are evaluated in the order they are specified. The first policy that matches is applied. The policy name list is handled by the SNMP agent as a single entity.

route-target

Syntax

```
route-target {ext-community}[[export ext-community][import ext-community]]
```

```
no route-target
```

Context

```
config>service>vpls>bgp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the route target (RT) component that is signaled in the related MPBGP attribute to be used for BGP auto-discovery, when this command is configured in the VPLS service.

If this command is not used, the RT is built automatically using the VPLS ID. The ext-comm can have the same two formats as the VPLS ID, a two-octet AS-specific extended community, IPv4 specific extended community.

The following rules apply.

- If the BGP AD VPLS ID is configured and no RT is configured under BGP node, RT = VPLS-ID.
- If the BGP AD VPLS ID is not configured, an RT value must be configured under the BGP node. (This is the case when only BGP VPLS is configured.)
- If the BGP AD VPLS ID is configured and an RT value is also configured under BGP node, the configured RT value prevails.

Parameters

export ext-community

Specifies the communities allowed to be sent to remote PE neighbors.

import ext-community

Specifies the communities allowed to be accepted from remote PE neighbors.

pw-template-binding

Syntax

pw-template-binding *policy-id* [**split-horizon-group** *group-name*] [**import-rt** {*ext-community*,...(up to 5 max)}]

no pw-template-bind *policy-id*

Context

config>service>vpls>bgp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command binds the advertisements received with the 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.

This command 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 a group name that overrides the split horizon group template settings.

import-rt ext-community

Specifies communities allowed to be accepted from remote PE neighbors. An extended BGP community is expressed 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

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

Description

This command configures the aging time for locally learned MAC addresses in the FDB for the VPLS instance.

In a VPLS service, MAC addresses are associated with a SAP. MACs associated with a SAP are classified as local MACs, and MACs associated with are remote MACs QinQ / access uplink SAPs.

Similar to 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 the local aging timer to the default value.

Default

local age 300

Parameters

aging-timer

Specifies the aging time for local MACs, 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

Description

Commands in this context configure 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** command 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 using the **config>service>vpls>sap** context. This means that when the relearn rate has exceeded the limit, another (blockable) SAP will be disabled instead.

The **mac-move** command enables the feature at the service level for SAPs, because only those objects can be blocked by this feature.

The operation of this feature is the same on the SAP; for example, if a MAC address moves from SAP to SAP, one will be blocked to prevent thrashing.

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 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

Description

This command configures 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)

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

Description

This command configures the time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled.

It is recommended that the **retry-timeout** value be larger than 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 value of zero indicates that the SAP will not be automatically reenabled after being disabled. If, after the SAP is reenabled 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

Description

This command specifies the multicast FIB (MFIB) high watermark. When the percentage filling level of the MFIB exceeds the configured value, a trap is generated and a log entry is added.

Parameters

high-water-mark

Specifies the MFIB 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

Description

This command specifies the MFIB low watermark. When the percentage filling level of the MFIB drops below the configured value, the corresponding trap is cleared and a log entry is added.

Parameters

low-water-mark

Specifies the MFIB 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

Description

This command specifies the maximum number of (s,g) entries in the MFIB for this VPLS instance.

The *table-size* parameter specifies the maximum number of multicast database entries for both learned and static multicast addresses for the VPLS instance. When a *table-size* limit is set on the MFIB of a service that is lower than the current number of dynamic entries present in the MFIB, the number of entries remains above the limit.

The **no** form of this command removes the configured maximum MFIB table size.

Parameters

table-size

Specifies the maximum number of (s,g) entries allowed in the MFIB.

Values	1 to 999 (7210 SAS-Dxp)
	1 to 1024 (7210 SAS-D)
	1 to 1000 (7210 SAS-K 2F1C2T)
	1 to 1015 (7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

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

Description

This command configures the aging time for remotely learned MAC addresses in the FDB for the VPLS instance. In a VPLS service, MAC addresses are associated with a SAP. MACs associated with a SAP are classified as local MACs.

Similar to 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 value larger than the **local-age** timer.

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

Default

remote age 900

Parameters

seconds

Specifies the aging time for remote MACs, in seconds.

Values 60 to 86400

send-flush-on-failure

Syntax

[no] send-flush-on-failure

Context

config>service>vpls

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables sending out "flush-all-from-ME" messages to all LDP peers included in the affected VPLS, in the event of physical port failures or "oper-down" events of individual SAPs.

This feature provides an LDP-based mechanism for recovering a physical link failure in a dual-homed connection to a VPLS service. This method provides an alternative to 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 the **send-flush-on-failure** command is enabled, "flush-all-from-ME" messages are 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

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

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** command 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.

When a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, all associated SAP and SDP binding operational states are automatically reevaluated.

The **no** form of this command reverts the default **service-mtu** for the indicated service type to the default value.



Note:

To disable the service MTU check, execute the **no service-mtu-check** command. Disabling service MTU check allows the packets to pass to the egress if the packet length is less than or equal to the MTU configured on the port.

Default

VPLS: 1514

The following table displays MTU values for specific VC types.

Table 43: MTU values for VC types

VC-type	Example service MTU	Advertised MTU
Ethernet	1514	1500

VC-type	Example service MTU	Advertised MTU
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

root-guard

Syntax

[no] root-guard

Context

config>service>vpls>sap>stp

config>service>vpls>spoke-sdp>stp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

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 a collection of policies (ACLs, QoS), including time-ranges, that define the full or partial behavior of a SAP. The suite can be applied to more than one SAP.

vsi-id

Syntax

vsi-id

Context

config>service>vpls>bgp-ad

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the low-order 4 bytes used to compose the 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

Description

This command configures an optional service name, up to 64 characters, that adds a name identifier to a specific service to then use that service name in configuration references and in show commands throughout the system. This helps service providers and administrators identify and manage services.

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).

allow-ip-int-bind

Syntax

[no] **allow-ip-int-bind**

Context

config>service>vpls

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command sets a flag on the VPLS service that enables the ability to attach an 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.

On the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T, the VPLS can be bound to only an IES IP interface. On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the VPLS can be bound to an IES or VPRN IP interface. Only IPv4 addressing and forwarding is supported in the R-VPLS context.

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 5.0 the following VPLS features must be disabled or not configured to set the **allow-ip-int-bind** flag:

- SAP ingress QoS policies applied to the VPLS SAPs cannot have MAC match criteria defined
- VPLS service type cannot be M-VPLS
- MVR from routed VPLS and to another SAP is not supported

When the VPLS **allow-ip-int-bind** flag is set on a VPLS service, the preceding features cannot be enabled on the VPLS service.

When 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** or **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 attached, the **no allow-ip-int-bind** command fails. When the **allow-ip-int-bind** flag is reset on the VPLS service, the configuration restrictions associated with setting the flag are removed.

5.7.2.1.4 VPLS STP commands

stp

Syntax

stp

Context

config>service>vpls

config>service>vpls>sap

config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure the Spanning Tree Protocol (STP) parameters. The Nokia implementation of the 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 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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>pw-template (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures automatic detection of the edge port characteristics of the SAP.

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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the SAP 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 parameter is set to edge-port. This value changes if the following is true.

- A BPDU is received on that port. This means that there is another bridge connected to this port. The edge-port becomes disabled.
- If auto-edge is configured and no BPDU is received within a certain period of time, RSTP concludes that it is on an edge and enables the edge-port.

The **no** form of this command reverts the edge port setting to the default value.

Default

no edge-port

forward-delay

Syntax

forward-delay *seconds*

no forward-delay

Context

config>service>vpls>stp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

RSTP, as defined in the IEEE 802.1D-2004 standards, usually transitions 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** command 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

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the STP hello time for the VPLS STP instance.

This command defines the default timer value that controls the sending interval between BPDU configuration messages by this bridge, on ports where this bridge assumes the designated role.

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 seconds

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

Platforms

Supported on all 7210 SAS platforms as described in this document

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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum number of bridges for STP behind this SAP. If there is only a single bridge, transitioning to forwarding state will be based on handshaking (fast transitions).

If more than two bridges are connected via shared media, their SAP should all be configured as shared, and timer-based transitions are used.

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

Default

pt-pt

Parameters

pt-pt

Specifies that a maximum of one bridge can exist behind this SAP or spoke-SDP.

shared

Specifies that two or more bridges can exist behind this SAP or spoke-SDP.

mst-instance

Syntax

```
mst-instance mst-inst-number
```

Context

```
config>service>vpls>sap>stp
```

```
config>service>vpls>stp
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies an MSTI for which to configure related parameters at the SAP level. This context can be open only for existing MSTIs 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 config>service>vpls>stp>mst-instance

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures 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

config>service>vpls>stp>mst-instance

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command 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

Description

This command configures 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. Every 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 seconds

Parameters

seconds

Specifies the max info age for the STP instance in seconds. Allowed values are integers.

Values 6 to 40

mode

Syntax

mode {**rstp** | **comp-dot1w** | **dot1w** | **mstp**}

no mode

Context

config>service>vpls>stp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the version of STP the bridge is currently running. See [Spanning tree operating modes](#) for information about these modes.

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

Default

rstp

Parameters

rstp

Corresponds to the Rapid STP 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 STP fully conformant to IEEE 802.1w.

mstp

Sets MSTP as the STP mode of operation. Corresponds to the Multiple STP specified in 802.1Q REV/D5.0-09/2005.

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

Description

This command specifies an MSTI for which to configure 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 Multiple Spanning Tree 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

Description

This command specifies the bridge priority for the specific MSTI for this service. The *bridge-priority* value reflects likelihood that the switch will be chosen as the regional root switch (65535 represents the least likely). It is used as the highest 4 bits of the bridge ID included in the MSTP BPDUs 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 created by the [vlan-range](#) command do not have explicit definition of bridge-priority and will inherit the default value.

Parameters

bridge-priority

Specifies the priority of this specific MSTI 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

Description

This command specifies a range of VLANs associated with a certain MST instance. This range applies to all SAPs of the mVPLS.

Every VLAN range that is not assigned within any of the instances created by the [mst-instance](#) command 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 shut down.

The **no** form of this command removes the **vlan-range** from the specific MST instance.

Parameters

vlan-range

The first VLAN range specifies the left-bound (minimum) value of a range of VLANs associated with the mVPLS SAP. This value must be smaller than or equal to the second VLAN range value. The second VLAN range specifies the right-bound (maximum) value of a range of VLANs 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

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

Description

This command defines an MST region name. Two bridges are considered a part of the same MST region as soon as their configuration of the MST region name, the MST-revision and VLAN-to-instance assignment are 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

Description

This command defines the MST configuration revision number. Two bridges are considered 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 MST configuration revision 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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the Spanning Tree Protocol (STP) path cost for the SAP.

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

STP suggests that the path cost is defined as a function of the link bandwidth. Because SAPs are controlled by complex queuing dynamics, in the 7210 SAS the STP path cost is a static configuration.

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

Parameters

path-cost

Specifies the path cost for the SAP.

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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the virtual port number that uniquely identifies a SAP within configuration BPDUs. The internal representation of a SAP is unique to a system and has a reference space much bigger than the 12 bits definable in a configuration BPDU. STP takes the internal representation value of a SAP and identifies it with its own virtual port number that is unique to every other SAP defined on the 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 *stp-priority*

no priority

Context

```
config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
config>service>vpls>sap>stp
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the STP priority for the SAP.

STP priority is a configurable parameter associated with a SAP. When configuration BPDUs are received, the priority is used in some circumstances as a tie breaking mechanism to determine whether the 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.

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 example, 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 the STP priority to the default value.

Default

128

Parameters

stp-priority

Specifies the STP priority value for the SAP. Allowed values are integer in the range of 0 to 255, 0 being the highest priority. The actual value used for STP priority (and stored in the configuration) is the result of masking out the lower 4 bits, therefore the actual value range is 0 to 240 in increments of 16.

Default 128

5.7.2.1.5 VPLS SAP commands

```
sap
```

Syntax

```
sap sap-id [eth-ring ring-index]
```

```
no sap sap-id
```

Context

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters that identify the service access point on the interface and within the 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 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.

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.

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. For IES, the IP interface must be shut down before the SAP on that interface may be removed.

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). The 7210 SAS does not support explicit null encapsulation for a 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 or disabled in the **environment>create** context.

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.

discard-unknown-source

Syntax

[no] discard-unknown-source

Context

config>service>vpls>sap

config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

When this command is enabled, packets received on a SAP or with an unknown source MAC address will be dropped only if the maximum number of MAC addresses for that SAP (see [max-nbr-mac-addr](#)) has been reached. If the **max-nbr-mac-addr** command has not been set for the SAP, enabling the **discard-unknown-source** command has no effect.

When disabled, the packets are forwarded based on the destination MAC addresses.

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

Default

no discard-unknown-source

5.7.2.1.6 ETH-CFM service commands

eth-cfm

Syntax

eth-cfm

Context

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context configure ETH-CFM parameters.

mep

Syntax

mep *mep-id* **domain** *md-index* **association** *ma-index* [**direction** {**up** | **down**}]

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

Context

config>service>vpls>sap>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document

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**}

Specifies the direction in which the MEP faces on the bridge port. Direction is not supported when a MEP is created directly under the **vpls>eth-cfm** context (vMEP).

Values **down** — Keyword that sends ETH-CFM messages away from the MAC relay entity.

up — Keyword that sends ETH-CFM messages toward the MAC relay entity.

ais-enable

Syntax

[no] ais-enable

Context

config>service>vpls>mesh-sdp>eth-cfm>mep (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>epipe>spoke-sdp>eth-cfm>mep (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>sap>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document

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>vpls>mesh-sdp>eth-cfm>mep>ais-enable (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>sap>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document

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

level

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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>sap>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the transmission interval of AIS messages.

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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>epipe>spoke-sdp>eth-cfm>mep>ais-enable (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>sap>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

Values 0 to 7

Default 1

ccm-enable

Syntax

[no] **ccm-enable**

Context

config>service>vpls>mep

config>service>vpls>sap>eth-cfm>mep

config>service>vpls>mesh-sdp>mep (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>epipe>spoke-sdp>eth-cfm>mep (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP + 8C) config>service>epipe>spoke-sdp>eth-cfm>mep (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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

By default, the highest priority on the bridge port is configured.

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

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>sap>eth-cfm>mep

config>service>vpls>spoke-sdp>eth-cfm>mep (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

For ETH-test to work, operators need to configure ETH-test parameters on both sender and receiver nodes. The ETH-test can then be performed 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 performed for both the provisioning and test to ensure the MEP is an Y.1731 MEP (MEP provisioned with domain format none, association format icc-based). If not, the operation fails. An error message in the CLI and SNMP indicates the problem.

test-pattern

Syntax

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

no test-pattern

Context

config>service>vpls>sap>eth-cfm>mep>eth-test-enable

config>service>vpls>mesh-sdp>eth-cfm>mep>eth-test-enable

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Keyword that specifies to use all zeros in the test pattern.

all-ones

Keyword that specifies to use all ones in the test pattern.

crc-enable

Keyword that generates a CRC checksum.

Default all-zeros

bit-error-threshold

Syntax

bit-error-threshold *bit-errors*

no bit-error-threshold

Context

config>service>vpls>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the threshold value of bit errors.

Parameters

bit-errors

Specifies the threshold value in bits.

Values 0 to 11840

Default 1

fault-propagation-enable

Syntax

fault-propagation-enable {*use-if-tlv* | *suspend-ccm*}

no fault-propagation-enable

Context

config>service>epipe>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the fault propagation for the MEP.

Parameters

use-if-tlv

Keyword that specifies to use the interface TLV.

suspend-ccm

Keyword that specifies to suspend the continuity check messages.

mac-address

Syntax

mac-address *mac-address*

no mac-address

Context

config>service>epipe>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

low-priority-defect

Syntax

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

Context

config>service>epipe>spoke-sdp>eth-cfm>mep (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>eth-cfm>mep (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>epipe>sap>eth-cfm>mep

Platforms

Supported on all 7210 SAS platforms as described in this document

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, DefErrorCCM, and DefXconCCM
	remErrXcon	Only DefRemoteCCM, DefErrorCCM, and DefXconCCM
	errXcon	Only DefErrorCCM and DefXconCCM

xcon	Only DefXconCCM; or
noXcon	No defects DefXcon or lower are to be reported

mac-address

Syntax

mac-address *mac-address*

no mac-address

Context

config>service>vpls>sap>eth-cfm>mep

config>service>vpls>mesh-sdp>eth-cfm>mep (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>epipe>spoke-sdp>eth-cfm>mep (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

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>sap>eth-cfm>mep

config>service>epipe>spoke-sdp>eth-cfm>mep (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables or 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*] [**primary-vlan-enable** *vlan-id*]

mip default-mac [**primary-vlan-enable** *vlan-id*]

no mip [**primary-vlan-enable** *vlan-id*]

Context

config>service>vpls>sap>eth-cfm

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command allows Maintenance Intermediate Points (MIPs) to be created if mhf-creation for the MA is configured using the default option.

For more information about the ETH-CFM matrix support for different platforms, see the 7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C OAM and Diagnostics Guide.

The **no** form of this command deletes the MIP.

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

Keyword that is useful if the user needs to change the MAC back to the default MAC without having to delete the MIP and reconfiguring.

Default no mip

primary-vlan-enable *vlan-id*

Specifies 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 supported only under Ethernet SAPs.

This parameter is supported only with UP MEP on 7210 SAS-D with a dot1q range SAP.

Values 0 to 4094

5.7.2.1.7 VPLS commands

limit-mac-move

Syntax

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

no limit-mac-move

Context

config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures whether the MAC-move agent, when enabled using the **config service vpls mac-move** or **config service epipe mac-move** commands will limit the MAC relearn (move) rate on this SAP.

Default

blockable

Parameters

blockable

Keyword to specify the agent will monitor the MAC relearn rate on the SAP, and it will block it when the relearn rate is exceeded.

non-blockable

Keyword to specify this SAP will not be blocked, and another blockable SAP will be blocked instead.

mac-pinning

Syntax

[no] mac-pinning

Context

config>service>vpls>sap

config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>endpoint

config>service>pw-template (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command disables relearning of MAC addresses on other SAPs within the VPLS. The MAC address remains attached to a specific SAP for the 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 SAP with **mac-pinning** enabled remains in the FIB on this SAP forever. Every event that would otherwise result in relearning will be logged (MAC address, original SAP, new SAP).

max-nbr-mac-addr

Syntax

max-nbr-mac-addr *table-size*

no max-nbr-mac-addr

Context

config>service>vpls>sap

config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>endpoint

config>service>pw-template (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum number of FDB entries for both learned and static MAC addresses for this SAP.

When the configured limit has been reached, and the [discard-unknown-source](#) command has been enabled for this SAP or spoke-SDP, packets with unknown source MAC addresses are discarded.

The **no** form of this command restores the global MAC learning limitations for the SAP.

Default

no max-nbr-mac-addr

Parameters

table-size

Specifies the maximum number of learned and static entries allowed in the FDB of this service.

Values 1 to 30719 (7210 SAS-D, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C)
1 to 15359 (7210 SAS-Dxp)

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

Description

This command configures the time to wait before reverting to the primary spoke-SDP.

The **no** form of this command removes the wait time.

Parameters

revert-time

Specifies the wait time in seconds.

Values 0 to 600

infinite

Keyword to make the endpoint non-revertive.

statistics

Syntax

statistics

Context

config>service>vpls>sap

Platforms

7210 SAS-D

Description

Commands in this context configure the counters associated with SAP ingress and egress.

egress

Syntax

egress

Context

config>service>vpls>sap>statistics

Platforms

7210 SAS-D

Description

Commands in this context configure the egress SAP statistics counter and set the mode of the counter. This counter counts the number of packets forwarded through the SAP.

ingress

Syntax

ingress

Context

config>service>vpls>sap>statistics

Platforms

7210 SAS-D

Description

Commands in this context configure the ingress SAP statistics counter.

For access-uplink SAPs, the ingress counters are not enabled by default. For access SAPs, if the ingress counter is enabled by default, it can be disabled.

The types of ingress SAP counters are the following:

- a counter that counts the total packets or octets received on the SAP
- a counter associated with meters defined in the QoS policy of the SAP. This counter counts the in-profile and out-of-profile packets or octets received on the SAP.

forwarded-count

Syntax

[no] forwarded-count

Context

config>service>vpls>sap>statistics>egress

Platforms

7210 SAS-D

Description

This command associates a counter with the SAP. The counter counts the number of packets forwarded through the SAP.

A limited number of such counters are available for use with access SAPs and access-uplink SAPs.

Use this command before enabling applicable accounting record collection on the SAP to associate a counter with the SAP.

The **no** form of this command disables the packet count.

counter-mode

Syntax

counter-mode {in-out-profile-count | forward-drop-count}

Context

config>service>vpls>sap>statistics>ingress

Platforms

7210 SAS-D

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 will be written into the current accounting file.

Run the following sequence of commands on the specified SAP to ensure the correct statistics are collected when the counter-mode is changed.

1. Run the **config service vpls sap no collect-stats** command to disable writing of accounting records for the SAP.
2. Change the **counter-mode** to the desired option by running the **config service vpls sap counter-mode {in-out-profile-count | forward-drop-count}** command.
3. Run the **config service vpls 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

forward-drop-count

Keyword to specify 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 are not individually available when operating in this mode.

in-out-profile-count

Keyword to specify 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.

received-count

Syntax

[no] received-count

Context

```
config>service>vpls>sap>statistics>ingress
```

Platforms

7210 SAS-D

Description

This command associates a counter with the SAP. It counts the number of packets and octets received on the SAP (ingress).

A limited number of such counters are available for use with access-uplink SAPs.

Use this command before enabling applicable accounting record collection on the SAP.

The **no** form of this command disables counter.

static-mac

Syntax

```
[no] static-mac ieee-mac-address [create]
```

Context

```
config>service>vpls>sap
```

```
config>service>vpls>mesh-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
```

```
config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command creates a local static MAC entry in the VPLS forwarding database (FDB) associated with the SAP.

In a VPLS service, MAC addresses are associated with a SAP or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Local static MAC entries create a permanent MAC address to SAP association in the forwarding database for the VPLS instance so that MAC address will not be learned on the edge device.

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 for each 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

Keyword that is mandatory 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

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 will be affected when the SAP changes state.

This command is only valid when the VPLS in which it is entered was created as a management VPLS.

default-sap

Syntax

[no] default-sap

Context

config>service>vpls>sap>managed-vlan-list

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds a default SAP to the managed VLAN list.

The **no** form of this command removes the default SAP from 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

Description

This command configures a range of VLANs on an access port managed by an existing management VPLS.

This command is valid only when the VPLS in which it is entered was created as a management VPLS, and when the SAP in which it was entered was created on an Ethernet port with encapsulation type of dot1q or QinQ, or on a Sonet/SDH port with encapsulation type of bcp-dot1q.

To modify the range of VLANs, first the new range should be entered and afterwards the old range removed. See [Modifying VPLS service parameters](#) for more information.

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

5.7.2.1.8 VPLS filter and QoS policy commands

egress

Syntax

egress

Context

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

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

Platforms

Supported on all 7210 SAS platforms as described in this document.

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 that has 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.

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

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. The filter policy must already exist within the created MAC filters.

Values 1 to 65535

qos

Syntax

qos *policy-id*

no qos

Context

config>service>vpls>sap>ingress

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command associates a QoS policy with an ingress SAP or IP interface.

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 ingress policies. The **qos** command only allows ingress policies to be associated on SAP ingress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specific type returns an error.

By default, if no specific QoS policy is associated with the SAP for ingress, 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 policy ID to associate with SAP on ingress. The policy ID must already exist.

Values 1 to 65535

agg-shaper-rate

Syntax

agg-shaper-rate *agg-rate*

no agg-shaper-rate

Context

config>service>vpls>sap>ingress

Platforms

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

Description

This command specifies the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues. Users can specify the CIR rate and the PIR rate. Users must not oversubscribe the total bandwidth available for use by ingress queues.

The **no** form of this command disables the SAP aggregate rate shaper; that is, the SAP can use up the maximum bandwidth available.

Default

no agg-shaper-rate

Parameters

agg-rate

Specifies the rate in kilobits per second.

Values 50 to 3000000 | max (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
64 to 20000000 | max (7210 SAS-K 3SFP+ 8C)

Default max

agg-shaper-rate

Syntax

agg-shaper-rate *agg-rate*

no agg-shaper-rate

Context

config>service>vpls>sap>egress

Description

7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, 7210 SAS-K 3SFP+ 8C

This command configures the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues. Users can specify the CIR rate and the PIR rate. Users must not oversubscribe the total bandwidth available for use by ingress queues.

The **no** form of this command disables the use of SAP aggregate rate shaper; that is, the SAP can use up the maximum bandwidth available.

Default

no agg-shaper-rate

Parameters

agg-rate

Specifies the rate in kilobits per second.

Values 50 to 1000000 | max (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
64 to 10000000 | max (7210 SAS-K 3SFP+ 8C)

Default max

aggregate-meter-rate

Syntax

aggregate-meter-rate *rate-in-kbps* [**burst** *burst-in-kbits*]

no aggregate-meter-rate

Context

```
config>service>epipe>sap>ingress
config>service>vpls>sap>ingress
```

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures the SAP ingress aggregate policer. The rate of the SAP ingress aggregate policer must be specified. Users can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the ingress traffic on different FCs and determines the final disposition of the packet. The packet is either forwarded to an identified profile or dropped.



Note:

The sum of CIR of the individual FCs configured under the SAP cannot exceed the PIR rate configured for the SAP. Although the 7210 SAS software does not block this configuration, it is not recommended.

The following table lists the final disposition of the packet based on the operating rate of the per-FC policer and the per-SAP aggregate policer.

Table 44: 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 ¹¹	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

¹¹ This configuration is not recommended for use. For more information, see the [aggregate-meter-rate](#) description.

Per FC meter operating rate	Per FC assigned color	SAP aggregate meter operating rate	SAP aggregate meter color	Final packet color
				Dropped

When the SAP aggregate policer is configured, the per-FC policer can be configured only 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

Parameters

rate-in-kbps

Specifies the rate in kilobits per second.

Values 0 to 20000000, max

Default max

burst-in-kilobits

Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 to 2146959

Default 512

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

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures the access SAP egress aggregate policer. The rate (PIR) of the SAP egress aggregate policer must be specified. Users can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the traffic sent out of the SAP and determines the final disposition of the packet, which is either forwarded or dropped.

Users can optionally associate a set of two counters to count total forwarded packets and octets, and total dropped packets and octets. When this counter is enabled, the amount of resources required increases by twice the amount of resources taken up when the counter is not used. If the **enable-stats** keyword is specified during the creation of the meter, the counter is allocated by software (if available). To free up the counter and relinquish its use, the user can use the **no aggregate-meter-rate** command and then recreate the meter using the **aggregate-meter rate** command.

If egress frame-based accounting is used, the SAP egress aggregate meter rate accounts for the Ethernet frame overhead. The system accounts for 12 bytes of IFG and 8 bytes of start delimiter. Frame-based accounting does not affect the count of octets maintained by the counter (if in use).



Note:

- Before enabling this command for a SAP, resources must be allocated to this feature from the egress-internal-tcam resource pool using the **config system resource-profile egress-internal-tcam egress-sap-aggregate-meter** command. See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide* for information.
- The egress aggregate meter is not FC-aware. The forward and drop decisions are made based on the order in which 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 4000000, max (7210 SAS-D)
1 to 20000000, max (7210 SAS-Dxp)

Default max

burst-in-kbits

Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 to 16384, default (7210 SAS-D)
4 to 2146959, default (7210 SAS-Dxp)

Default 512

enable-stats

Keyword to specify if the counter to count forwarded and dropped packets and octets is allocated. If this keyword is used while configuring the meter, the counter is allocated.

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy

Context

config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document

Description



Note: This command is not applicable for access uplink SAPs.

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

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>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>sap

Platforms

Supported on all 7210 SAS platforms as described in this document

Description



Note:

This command is not applicable for access uplink SAPs.

This command enables accounting and statistical data collection for either the SAP, network port, or IP interface. By default, when applying accounting policies the data is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default

no collect-stats

5.7.2.1.9 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command binds a VPLS service to an existing Service Distribution Point (SDP). Mesh SDPs bound to a service are logically treated like a single bridge "port" for flooded traffic where flooded traffic received on any mesh SDP on the service is replicated to other "ports" (spoke-SDPs and SAPs) and not transmitted on any mesh SDPs.

This command creates a binding between a service and an SDP. The SDP has an operational state that determines the operational state of the SDP within the service; for example, if the SDP is administratively or operationally down, the SDP for the service 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 removed, no packets are forwarded to the far-end router.

Special Cases

VPLS

Several SDPs can be bound to a VPLS. Each SDP must be destined for a different router. If two *sdp-id* bindings terminate on the same router, an error occurs and the second SDP is 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

Keyword that overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15-bit quantity containing a value that represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the bindings VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

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

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

ether

Keyword that 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. (hex 5)

vlan

Keyword that 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] endpoint no-endpoint
```

```
no spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create] endpoint no-endpoint
```

Context

```
config>service>vpls
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command binds a service to an existing SDP. A spoke-SDP is treated like the equivalent of a traditional bridge "port" where flooded traffic received on the spoke-SDP is replicated on all other "ports" (other spoke and mesh SDPs and SAPs) and not transmitted on the port on which it was received.



Note:

On the 7210 SAS-K 2F6C4T, the **spoke-sdp** command is supported on a routed VPLS when it is bound to a VPRN service. All configuration using the **config>service>spoke-sdp** context is supported on the routed VPLS.

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 exist in the **config>service>sdp** context before it can be associated with a VPLS 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 service is created.

SDPs must be explicitly associated and bound to a service to allow far-end 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 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 unique *vc-ids*. An error message is generated if two SDP bindings with identical *vc-ids* terminate on the same router. Split horizon groups can only be created in the scope of a VPLS service.

Parameters

sdp-id

Specifies the SDP identifier.

Values 1 to 17407

vc-id

Specifies the virtual circuit identifier.

Values 1 to 4294967295

create

Keyword that is mandatory for creating a spoke-SDP.

endpoint *endpoint-name*

Specifies the service endpoint to which this SDP bind is attached. The service ID of the SDP binding must match the service ID of the service endpoint.

no endpoint

Removes the association of a spoke-SDP with an explicit endpoint name.

ether

Keyword to specify 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)

split-horizon-group *group-name*

Specifies the name of the split horizon group to which the SDP belongs.

vc-type

Specifies 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 disabled, a change of the binding 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.

Values ether, vlan

vlan

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

The VLAN VC-type requires at least one dot1q tag within each encapsulated Ethernet packet transmitted to the far end.

adv-service-mtu

Syntax

adv-service-mtu *octets*

no adv-service-mtu

Context

config>service>vpls>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the MTU value signaled in a targeted LDP for the spoke-SDP and the value used to validate the MTU value signaled by the far-end PE. If configured, this value is used instead of the service MTU. However, the configuration does not affect the locally enforced value, which is still based on the service MTU. This command for the MTU cannot be configured on a spoke-SDP that is bound to an SDP with the **config>service>sdp>adv-mtu-override** command.

When unconfigured, an adjusted service MTU is used. See the **config>service>vpls>service-mtu** command for more information.

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

Default

no adv-service-mtu

Parameters

octets

Specifies the MTU size, in octets, expressed as a decimal integer.

Values 0 to 9212

control-word

Syntax

[no] control word

Context

config>service>vpls>mesh-sdp

config>service>vpls>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 the 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

egress

Syntax

egress

Context

config>service>vpls>mesh-sdp

config>service>vpls>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the egress SDP context.

hash-label

Syntax

hash-label [signal-capability]

no hash-label

Context

config>service>vpls>mesh-sdp

config>service>vpls>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the hash label on VLL or VPLS services that are bound to RSVP SDP, 3107 BGP SDP, segment routing, or LDP SDP, using the auto-bind mode with the **ldp**, **rsvp-te**, or **mpls** options. When this feature 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. The ingress datapath appends the hash label at the bottom of the stack (BoS) and sets the S-bit to one (1).



Note:

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

Packets generated in the CPM that are forwarded with a label 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 0.

Signaling of the hash label capability is enabled by adding the **signal-capability** option under the VLL spoke-SDP, VPLS spoke-SDP or mesh SDP interface, or PW template instance. In this case, the decision of the local PE to insert the hash label on the user and control plane packets is determined by the outcome of the signaling process and can override the local PE configuration. The following process flow applies when the **hash-label** and **signal-capability** options 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 the specific spoke-SDP or mesh SDP.
- If a 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 the user and control plane packets that it forwards on the spoke-SDP or mesh SDP. The local PE also drops user and control plane packets received from a remote PE if they include a hash label. The dropped packets may be caused by:
 - a remote 7210 SAS PE that does not support the **hash-label** command
 - a remote 7210 SAS PE that has the **hash-label** command enabled but does not support the **signal-capability** option
 - a remote 7210 SAS PE that supports the **hash-label** command and the **signal-capability** option, but the user did not enable them due to a misconfiguration
- If the remote PE sends a flow label sub-TLV in the PW ID FEC element with T=TRUE and R=TRUE, the local PE enables the hash label capability. Consequently, the local PE node inserts a hash label in the user and control plane packets that it forwards on the spoke-SDP or mesh SDP. The local PE node also accepts user and control plane packets from the remote PE with a hash label. The local PE node drops user and control plane packets from the remote PE without a hash label.

If the **hash-label** command is enabled on the local PE with the **signal-capability** option configured and on the remote PE without the **signal-capability** option configured on the spoke-SDP or mesh-SDP, the hash label is included in the pseudowire packets received by the local PE. These packets must be dropped. To resolve this situation, you must disable the **signal-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 is not enabled on the local configuration of the spoke-SDP or mesh-SDP at the remote PE, the hash label is not included in the pseudowire received by the local PE.

If the **signal-capability** option is enabled or disabled in the CLI, 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 mLDP LSP in both the base router instance and in the multicast VPN (mVPN) instance.
- In 7750 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 to 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. 7210 SAS devices do not set the MSB in the hash label value for service traffic. Therefore, the user must ensure that both ends are correctly configured to either process hash labels or disable them. The MSB bit is set for MPLS/OAM traffic on 7210 SAS devices.
- The **cpe-ping**, **mac-ping**, and **svc-ping** commands are not supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C when the **hash-label** command is enabled.

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

Default

no hash-label

Parameters

signal-capability

Keyword that enables the signaling and negotiation of the use of the hash label between the local and remote PE nodes.

ingress

Syntax

ingress

Context

config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the ingress SDP context.

vc-label

Syntax

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

Context

config>service>vpls>mesh-sdp>egress

config>service>vpls>spoke-sdp>egress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the egress VC label.

Parameters

vc-label

Specifies a 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the ingress VC label.

Parameters

vc-label

Specifies a VC ingress value that indicates a specific connection.

Values 2048 to 18431

suppress-vc-vlan

Syntax

suppress-vc-vlan
no suppress-vc-vlan

Context

config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command suppresses the processing of the **vlan-vc-tag** value in the dataplane. When this command is enabled, the dataplane does not add a VLAN tag with a value of 0 or the configured **vlan-vc-tag** value when forwarding the packet out of the PW. When this command is enabled, the dataplane does not remove the outermost VLAN tag in the encapsulated packet received on the PW.

The **no** form of this command disables the suppression.

Default

no suppress-vc-vlan

vlan-vc-tag

Syntax

vlan-vc-tag *vlan-id*
no vlan-vc-tag [*vlan-id*]

Context

config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

vlan-id

Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

Values 0 to 4094

5.7.2.1.10 IGMP snooping commands

disable-router-alert-check

Syntax

[no] **disable-router-alert-check**

Context

config>service>vpls>sap>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the Internet Group Management Protocol (IGMP) router alert check option.

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

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

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.

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

Description

This command enables fast leave. When the IGMP fast leave processing is enabled, the 7210 SAS immediately removes a SAP from the multicast group when it detects an IGMP "leave" on that SAP.

Fast-leave processing allows the switch to remove a SAP that sends a "leave" from the forwarding table without first sending out group-specific queries to the SAP, and 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.

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

Description

This command configures the VPLS from which multicast traffic is copied upon receipt of an IGMP join request. IGMP snooping must be enabled on the MVR VPLS.

Default

no from-vpls

Parameters

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

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds a static multicast group as a (*, g). When a static IGMP group is added, multicast data for that (*,g) or (s,g) is forwarded to the specific SAP without receiving any membership report from a host.

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

Description

This command identifies filter policy of multicast groups to be applied to this VPLS entity. The sources of the multicast traffic must be a member of the VPLS. By default, no group policy is identified.

The **no** form of this command removes the policy association from the VPLS configuration.

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.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command forces VC VLAN-type forwarding in the datapath for spoke-SDPs or mesh SDPs that 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

igmp-snooping

Syntax

igmp-snooping

Context

config>service>vpls

config>service>vpls>sap

config>service>vpls>spoke-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>pw-template (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command enables the IGMP snooping context.

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

Description

This command configures the maximum number of multicast sources allowed per group.

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

Parameters

max-num-sources

Specifies the maximum number of multicast sources allowed per group.

Values 1 to 2047 (7210 SAS-D)
1 to 900 (7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T,
and 7210 SAS-K 3SFP+ 8C)

import

Syntax

import *policy-name*

no import

Context

config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>pw-template>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>pw-template>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

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 *count*

no max-num-groups

Context

config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>pw-template>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the maximum number of multicast groups that can be joined on this SAP or SDP. If the node receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

Default

no max-num-groups

Parameters

count

Specifies the maximum number of groups that can be joined on this SAP.

Values	1 to 1024 (7210 SAS-D) 1 to 900 (7210 SAS-Dxp, 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C)
---------------	--

mrouter-port

Syntax

[no] mrouter-port

Context

config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures whether a multicast router is attached behind this SAP.

Configuring a SAP or SDP as an mrouter-port has a double effect. Firstly, all multicast traffic received on another SAP or SDP is copied to this SAP or SDP. Secondly, 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. While the other multicast router (non-querier) stops sending IGMP queries, it should still receive reports to keep its multicast trees up to date. To support this, the **mrouter-port** command should be enabled on all SAPs or SDPs connecting to a multicast router.

The IGMP version to be used for the reports (v1 or v2) can only be determined after an initial query has been received. Until that time, no reports are sent on the SAP, even if **mrouter-port** is enabled.

If the **send-queries** command is enabled on this SAP or SDP, the **mrouter-port** command cannot be configured.

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

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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>pw-template>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the IGMP query interval. If the **send-queries** command is enabled, this command 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

query-interval 125

Parameters

seconds

Specifies the time interval, in seconds, that the router transmits general host-query messages.

Values

1 to 65535

config service vpls igmp-snooping
config service vpls spoke-sdp igmp-snooping
config service vpls mesh-sdp igmp-snooping

2 to 1024

config service vpls sap igmp-snooping
config service pw-template igmp-snooping

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

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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>pw-template>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures the IGMP query response interval. If the **send-queries** command is enabled, this command specifies the maximum response time advertised in IGMP queries.

The **query-response-interval** must be smaller than the **query-interval**.

If **send-queries** is not enabled on this SAP or SDP, the **query-response-interval** value is ignored.

Default

query-response-interval 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

report-src-ip

Syntax

report-src-ip *ip-address*

no report-src-ip

Context

```
config>service>vpls>igmp-snooping
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures 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

```
report-src-ip 0.0.0.0
```

Parameters

ip-address

Specifies the source IP source address in transmitted IGMP reports.

robust-count

Syntax

```
robust-count robust-count
```

```
no robust-count
```

Context

```
config>service>vpls>igmp-snooping
```

```
config>service>vpls>sap>igmp-snooping
```

```
config>service>vpls>spoke-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
```

```
config>service>vpls>mesh-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
```

```
config>service>pw-template>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
```

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

If the **send-queries** command is enabled, this command allows tuning for the expected packet loss on a SAP or SDP. The *robust-count* value 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 command is ignored.

Default

robust-count 2

Parameters

robust-count

Specifies the robust count for the SAP.

Values **config service vpls sap igmp-snooping:** 2 to 7
 config service vpls igmp-snooping: 1 to 255

precedence

Syntax

precedence *precedence-value* | **primary**

no precedence

Context

config>service>vpls>spoke-sdp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the spoke-SDP precedence.

Default

precedence 4

Parameters

precedence-value

Specifies the spoke-SDP precedence.

Values 0 to 4

primary

Keyword to specify the precedence is primary.

propagate-mac-flush

Syntax

[no] propagate-mac-flush

Context

config>service>vpls

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures 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 (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>pw-template>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures whether to send IGMP general query messages on the SAP or SDP.

When the **send-queries** command is configured, all query reports generated locally are of the type belonging to the configured version. If a report of a version higher than the configured version is received, the report is dropped and a new counter to track the wrong version is incremented. If **send-queries** is not configured, the version command has no effect, and the version used is the version of the querier.

Default

no send-queries

group

Syntax

[no] group *grp-address*

Context

config>service>vpls>sap>igmp-snooping>static

config>service>vpls>spoke-sdp>snooping>static (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>snooping>static (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds a static multicast group either as a (*, g) or as one or more (s,g) records. When a static IGMP group is added, multicast data for that (*,g) or (s,g) is forwarded to the specific SAP or SDP without receiving a membership report from a host.

Parameters

grp-address

Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

SOURCE

Syntax

[no] source *ip-address*

Context

config>service>vpls>sap>igmp-snooping>static>group

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command adds a static (s,g) entry, to allow multicast traffic for a multicast group from a specified source. For a multicast group, more than one source address can be specified. Static (s,g) entries cannot be added, if a starg is previously created.

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

Parameters

ip-address

Specifies the IPv4 unicast address.

starg

Syntax

[no] starg

Context

config>service>vpls>sap>igmp-snooping>static>group

config>service>vpls>spoke-sdp>igmp-snooping>static>group (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>igmp-snooping>static>group (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

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

static

Syntax

static

Context

config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context 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 are forwarded even if no join message was registered for the specific group.

version

Syntax

version *version*

no version

Context

config>service>vpls>sap>igmp-snooping

config>service>vpls>mesh-sdp>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>vpls>mesh-sdp>snooping>static (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

config>service>pw-template>igmp-snooping (only supported on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command specifies the version of IGMP that is running on this SAP. This object can be used to configure a router capable of running either value. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN.

When the **send-queries** command is configured, all query reports generated locally are of the type belonging to the configured version. If a report of a version higher than the configured version is received, the report is dropped and a new counter to track the wrong version is incremented. If **send-queries** is not configured, the version command has no effect, and the version used is the version of the querier.

Parameters

version

Specifies the IGMP version.

Values 1, 2, or 3

to-sap

Syntax

to-sap *sap-id*

no to-sap

Context

config>service>vpls>sap>igmp-snooping>mvr

Platforms

Supported on all 7210 SAS platforms as described in this document

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 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.

5.7.2.2 VPLS show commands

fdb-info

Syntax

fdb-info

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Displays global FDB usage information.

Output

The following output is an example of FDB usage information, and [Table 45: Output fields: FDB-info](#) describes the output fields.

Sample output

```
A:7210-SASE# show service fdb-info
=====
Forwarding Database(FDB) Information
=====
Service Id : 1 Mac Move : Disabled
Primary Factor : 3 Secondary Factor : 2
Mac Move Rate : 2 Mac Move Timeout : 10
Mac Move Retries : 3
Table Size : 250 Total Count : 1
Learned Count : 0 Static Count : 0
Host-learned Count: 1
Remote Age : 900 Local Age : 300
High Watermark : 95% Low Watermark : 90%
Mac Learning : Enabled Discard Unknown : Disabled
Mac Aging : Enabled Relearn Only : False
Mac Subnet Len : 48
-----
Total Service FDBs : 1
Total FDB Configured Size : 250
Total FDB Entries In Use : 1
PBB MAC Address Indices In Use : 0
-----
A:7210-SASE#
```

Table 45: Output fields: FDB-info

Label	Description
Service ID	The value that identifies a service.
Mac Move	Indicates the administrative state of the MAC movement feature associated with the service.
Mac Move Rate	The maximum rate at which MACs can be relearned in this TLS service, before the SAP where the moving MAC was last seen is automatically disabled to protect the system against undetected loops or duplicate MACs. 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	Indicates the time in seconds to wait before a SAP that has been disabled after exceeding the maximum relearn rate is reenabled. A value of zero indicates that the SAP will not be automatically re-enabled after being disabled. If after the SAP is reenabled it is disabled again, the effective retry timeout is doubled to avoid thrashing.

Label	Description
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 use of the FDB table of this service at which a "table full" alarm is raised by the agent.
Low WaterMark	The use of the FDB table of this service at which a "table full" alarm is cleared by the agent.
Mac Learning	Specifies whether the MAC learning process is enabled in this service.
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded in this service.
MAC Aging	Specifies whether the MAC aging process is enabled in this service.
MAC Pinning	Specifies whether MAC pinning is enabled in this service.
Relearn Only	When enabled, indicates that either the FDB table of this service is full or that the maximum system-wide number of 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

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 MAC address FDB entry information, and [Table 46: 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 46: 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.

Label	Description
Type/Age	Static — FDB entries created by management. Learned — Dynamic entries created by the learning process. OAM — Entries created by the OAM process. H — Host, the entry added by the system for a static configured subscriber host. D or DHCP — DHCP-installed MAC. Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease. P — Indicates the MAC is protected by the MAC protection feature.

sap-using

Syntax

sap-using [ingress | egress] filter *filter-id*

sap-using [sap *sap-id*]

Context

show>service

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Keyword that specifies matching an ingress policy.

egress

Keyword that 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

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 SAP information, and [Table 47: 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  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 47: 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.

service-using

Syntax

service-using [epipe] [ies] [vpls] [mirror] [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.

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 48: 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:SetupCLI# show service service-using
=====
Services
=====
ServiceId  Type      Adm   Opr      CustomerId  Last Mgmt Change
-----
23         mVPLS    Up    Down     2            09/25/2007 21:45:58
100        Epipe    Up    Down     2            09/25/2007 21:45:58
101        Epipe    Up    Down     2            09/25/2007 21:45:58
102        Epipe    Up    Down     2            09/25/2007 21:45:58
105        Epipe    Up    Down     2            09/25/2007 21:45:58
110        Epipe    Up    Down     1            09/25/2007 21:45:58
990        IES      Up    Down     1            09/25/2007 21:45:58
1000       Mirror   Up    Down     1            09/25/2007 21:45:59
1001       Epipe    Up    Down     1            09/25/2007 21:45:58
1002       Epipe    Up    Down     1            09/25/2007 21:45:58
1003       Epipe    Up    Down     1            09/25/2007 21:45:58
1004       Epipe    Up    Down     1            09/25/2007 21:45:58
2000       Mirror   Up    Down     1            09/25/2007 21:45:59
...
-----
Matching Services : 27
-----
*A:SetupCLI#

*A:SetupCLI# show service service-using
=====
Services
=====
ServiceId  Type      Adm   Opr      CustomerId  Last Mgmt Change
-----
23         mVPLS    Up    Down     2            09/25/2007 21:45:58
100        Epipe    Up    Down     2            09/25/2007 21:45:58
    
```

```

101      Epipe    Up    Down    2          09/25/2007 21:45:58
102      Epipe    Up    Down    2          09/25/2007 21:45:58
105      Epipe    Up    Down    2          09/25/2007 21:45:58
110      Epipe    Up    Down    1          09/25/2007 21:45:58
990      IES      Up    Down    1          09/25/2007 21:45:58
1000     Mirror    Up    Down    1          09/25/2007 21:45:59
1001     Epipe    Up    Down    1          09/25/2007 21:45:58
1002     Epipe    Up    Down    1          09/25/2007 21:45:58
1003     Epipe    Up    Down    1          09/25/2007 21:45:58
...
-----
Matching Services : 27
-----
*A:SetupCLI#
    
```

Table 48: Output fields: service-using

Label	Description
Service Id	The service identifier.
Type	Specifies the service type configured for the service ID.
Adm	The administrative state of the service.
Opr	The operating state of the service.
CustomerID	The ID of the customer who owns this service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this service.

id

Syntax

id *service-id*

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.

Values *service-id*: 1 to 214748364
svc-name: A string up to 64 characters.

- all**
Displays detailed information about the service.
- arp**
Displays ARP entries for the service.
- base**
Displays basic service information.
- fdb**
Displays FDB entries.
- igmp-snooping**
Displays IGMP snooping information.
- interface**
Displays service interfaces.
- mstp-configuration**
Displays MSTP information.
- sap**
Displays SAPs associated to 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

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 49: Output fields: service ID all](#) describes the output fields.

Sample output

```
*A:Dut-G# show service id 1 all
=====
Service Detailed Information
=====
Service Id       : 1                Vpn Id           : 0
Service Type    : VPLS
Name            : (Not Specified)
Description     : Default tls description for service id 1
Customer Id     : 1                Creation Origin  : manual
Last Status Change: 03/06/2020 06:12:06
Last Mgmt Change : 03/06/2020 06:12:02
Admin State    : Up                Oper State       : Up
Admin State    :                   Def. Mesh VC Id : 1
SAP Count      : 2
SAP Type       : Any                Customer vlan    : n/a
VSD Domain     : <none>
SPI load-balance : Disabled
TEID load-balance : Disabled

-----
BGP Information
-----

Split Horizon Group specifics
-----

ETH-CFM service specifics
-----
Tunnel Faults   : ignore                V-Mep Extensions : Enabled
-----

Service Access Points
-----

SAP 1/1/7:1
-----
Service Id      : 1
SAP             : 1/1/7:1                Encap           : q-tag
Description     : SAP 1/1/7:1 on Ethernet Ring 1
Admin State    : Up                Oper State       : Up
Flags          : None
VLAN Authentication: Disabled          VLAN Auth Status : N/A
Last Status Change : 03/06/2020 06:12:13
Last Mgmt Change : 03/06/2020 06:12:06
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
OAM MAC Addr    : 0                DHCP MAC Addr   : 0
Host MAC Addr   : 0                Intf MAC Addr   : 0
SPB MAC Addr    : 0                Cond MAC Addr   : 0
Admin MTU       : 9212              Oper MTU        : 9212
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          Limit Unused BW : Disabled
Mac Aging       : Enabled          Discard Unkwn Srce: Disabled
BPDU Translation : Disabled        Mac Pinning     : Disabled
L2PT Termination : 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)
Cflowd           : Disabled
-----
Eth-Ring Information
-----
Ethernet Ring    : 1                      Path          : b
Control Flag     : Yes
-----
Stp Service Access Point specifics
-----
Stp Admin State  : Down                   Stp Oper State  : Down
Core Connectivity : Down
Port Role        : N/A                   Port State      : Forwarding
Port Number      : N/A                   Port Priority    : 128
Port Path Cost   : 10                    Auto Edge       : Enabled
Admin Edge       : Disabled              Oper Edge       : N/A
Link Type        : Pt-pt                 BPDU Encap      : Dot1d
Root Guard       : Disabled              Active Protocol  : N/A
Last BPDU from   : N/A
CIST Desig Bridge : N/A                   Designated Port : N/A
Forward transitions: 0                   Bad BPDUs rcvd  : 0
Cfg BPDUs rcvd   : 0                   Cfg BPDUs tx    : 0
TCN BPDUs rcvd   : 0                   TCN BPDUs tx    : 0
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                   Egress qos-policy : 1
-----
Rate              : n/a                   Burst              : n/a
-----
Egress Aggregate Meter
-----
Rate              : n/a                   Burst              : n/a
-----
DHCP
-----
Description       : (Not Specified)
Admin State       : Down                   Lease Populate     : 0
DHCP Snooping     : Down                   Action             : Keep
Proxy Admin State : Down
Proxy Lease Time  : N/A
Emul. Server Addr : Not Configured
-----
Subscriber Management
-----
Admin State       : Down                   MAC DA Hashing    : False
Def Sub-Id        : None
Def Sub-Profile   : None
Def SLA-Profile   : None
Def Inter-Dest-Id : None
Sub-Ident-Policy  : None
Subscriber Limit   : 1
```

```

Single-Sub-Parameters
Prof Traffic Only : False
Non-Sub-Traffic   : N/A
Static host management
MAC learn options : N/A
-----
Ingress QoS Classifier Usage
-----
Classifiers Allocated: 2           Meters Allocated   : 1
Classifiers Used      : 1           Meters Used        : 1
-----
Sap Statistics
-----
Ingress Stats:      Packets      Octets
                   319           20416
Egress Stats:       311           19904
Ingress Drop Stats: 0             0
-----
Sap per Meter stats (in/out counter mode)
-----
Ingress Meter 1
For. InProf         : 64           4096
For. OutProf        : 255          16320
-----
SAP 1/1/9:1
-----
Service Id          : 1
SAP                 : 1/1/9:1      Encap                : q-tag
Description         : SAP 1/1/9:1 on Ethernet Ring 1
Admin State         : Up           Oper State            : Up
Flags               : None
VLAN Authentication: Disabled     VLAN Auth Status     : N/A
Last Status Change : 03/06/2020 06:12:08
Last Mgmt Change   : 03/06/2020 06:12:06
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
OAM MAC Addr       : 0           DHCP MAC Addr        : 0
Host MAC Addr      : 0           Intf MAC Addr        : 0
SPB MAC Addr       : 0           Cond MAC Addr        : 0
Admin MTU          : 9212        Oper MTU              : 9212
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
Limit Unused BW   : Disabled
Mac Learning      : Enabled      Discard Unkwn Srce   : Disabled
Mac Aging         : Enabled      Mac Pinning           : Disabled
BPDU Translation : Disabled
L2PT Termination : Disabled
Acct. Pol         : None         Collect Stats         : Disabled
Anti Spoofing     : None         Dynamic Hosts         : Enabled
Oper Group        : (none)       Monitor Oper Grp     : (none)
Host Lockout Plcy: n/a
Lag Link Map Prof : (none)
Cflowd            : Disabled
-----
Eth-Ring Information
-----
Ethernet Ring      : 1           Path                  : a
Control Flag       : Yes
    
```

```
-----  
Stp Service Access Point specifics  
-----  
Stp Admin State      : Down                Stp Oper State      : Down  
Core Connectivity    : Down  
Port Role           : N/A                Port State         : Forwarding  
Port Number         : N/A                Port Priority      : 128  
Port Path Cost      : 10                 Auto Edge         : Enabled  
Admin Edge          : Disabled           Oper Edge         : N/A  
Link Type           : Pt-pt              BPDU Encap       : Dot1d  
Root Guard          : Disabled           Active Protocol   : N/A  
Last BPDU from      : N/A  
CIST Desig Bridge   : N/A                Designated Port   : N/A  
Forward transitions : 0                 Bad BPDUs rcvd   : 0  
Cfg BPDUs rcvd     : 0                 Cfg BPDUs tx     : 0  
TCN BPDUs rcvd     : 0                 TCN BPDUs tx     : 0  
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                 Egress qos-policy : 1  
-----  
Rate                : n/a                 Burst              : n/a  
-----  
Egress Aggregate Meter  
-----  
Rate                : n/a                 Burst              : n/a  
-----  
DHCP  
-----  
Description         : (Not Specified)  
Admin State         : Down                Lease Populate     : 0  
DHCP Snooping       : Down                Action             : Keep  
Proxy Admin State   : Down  
Proxy Lease Time    : N/A  
Emul. Server Addr   : Not Configured  
-----  
Subscriber Management  
-----  
Admin State         : Down                MAC DA Hashing    : False  
Def Sub-Id          : None  
Def Sub-Profile     : None  
Def SLA-Profile     : None  
Def Inter-Dest-Id   : None  
Sub-Ident-Policy    : None  
Subscriber Limit    : 1  
Single-Sub-Parameters  
Prof Traffic Only   : False  
Non-Sub-Traffic    : N/A  
Static host management  
MAC learn options   : N/A  
-----  
Ingress QoS Classifier Usage  
-----
```

```

Classifiers Allocated: 2           Meters Allocated   : 1
Classifiers Used      : 1           Meters Used        : 1
-----
Sap Statistics
-----
Ingress Stats:      Packets      Octets
                   319          20416
Egress Stats:       311          19904
Ingress Drop Stats: 0            0
-----
Sap per Meter stats (in/out counter mode)
-----
Ingress Meter 1
For. InProf        : 64          4096
For. OutProf       : 255         16320
-----
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.ab:8f:80:f8:c0:01 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
Host-learned Count: 0            Intf MAC Count    : 0
Spb Count         : 0            Cond MAC 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          Oper MRtr Send Max   Max MVR      Num
Id           Stat Port Qrys Grps  Srcs From-VPLS Grps
-----
sap:1/1/7:1  Up   No   No   None  None Local  0
sap:1/1/9:1  Up   No   No   None  None Local  0
=====
-----
WLAN Gateway specifics
    
```

```

-----
Admin State           : disabled
Description           : (Not Specified)
SAP-template          : (Not Specified)
Last management change : (Not Specified)
Not supported on 7210
=====
VPLS VXLAN, Ingress VXLAN Network Id: 0
=====
Egress VTEP, VNI
=====
VTEP Address          Egress VNI      Num. MACs    Mcast    Oper State  L2 PBR
-----
No Matching Entries
=====
Service Endpoints
-----
No Endpoints found.
=====
VPLS Sites
=====
Site                  Site-Id  Dest          Mesh-SDP  Admin   Oper   Fwdr
-----
No Matching Entries
=====
    
```

Table 49: Output fields: service ID all

Label	Description
Service Id	The service identifier.
VPN Id	The number that identifies the VPN.
Service Type	Specifies the type of service.
SDP Id	The SDP identifier.
Description	Generic information about the service.
Customer Id	The customer identifier.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
SAP Count	The number of SAPs specified for this service.
SDP Bind Count	The number of SDPs bound to this service.
Split Horizon Group	Name of the split-horizon group for this service.
Description	Description of the split-horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split-horizon group.

Label	Description
SDP Id	The SDP identifier.
Type	Indicates whether this service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP: MPLS.
Admin State	The administrative state of this SDP or SAP.
Oper State	The operational state of this SDP or SAP.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP.
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP.
Ingress Filter	The ID of the ingress filter policy.
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field.
Number of SDPs	The total number SDPs applied to this service ID.
Service Id	The service identifier.

Label	Description
Port Id	The ID of the access port where this SAP is defined.
Description	Generic information about the SAP.
Encap Value	The value of the label used to identify this SAP on the access port.
Admin State	The administrative state of the SAP.
Oper State	The operating state of the SAP.
Last Changed	The date and time of the last change.
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The SAP ingress QoS policy ID.
Egress qos-policy	The SAP egress QoS policy ID.
Ingress Filter-Id	The SAP ingress filter policy ID.
Egress Filter-Id	The SAP egress filter policy ID.
Multi Svc Site	Indicates the multi-service site that the SAP is a member.
Ingress sched-policy	Indicates the ingress QoS scheduler for the SAP.
Egress sched-policy	Indicates the egress QoS scheduler for the SAP.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
Dropped	The number of packets or octets dropped.
Offered Hi Priority	The number of high priority packets, as determined by the SAP ingress QoS policy.
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
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.

Label	Description
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Dropped In Profile	The number of in-profile packets or octets discarded.
Dropped Out Profile	The number of out-of-profile packets or octets discarded.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Ingress Queue 1	The index of the ingress QoS queue of this SAP.
High priority offered	The packets or octets count of the high priority traffic for the SAP.
High priority dropped	The number of high priority traffic packets or octets dropped.
Low priority offered	The packets or octets count of the low priority traffic.
Low priority dropped	The number of low priority traffic packets or octets dropped.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
Out profile forwarded	The number of out-of-profile octets (rate above CIR) forwarded.
Egress Queue 1	The index of the egress QoS queue of the SAP.
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded.
In profile dropped	The number of in-profile packets or octets dropped for the SAP.
Out profile forwarded	The number of out-of-profile packets or octets (rate above CIR) forwarded.
Out profile dropped	The number of out-of-profile packets or octets discarded.
State	Specifies whether DHCP relay is enabled on this SAP.
Info Option	Specifies whether Option 82 processing is enabled on this SAP.
Action	Specifies the Option 82 processing on this SAP or interface: keep, replace, or drop.
Circuit ID	Specifies whether the If Index is inserted in circuit ID suboption of Option 82.
Remote ID	Specifies whether the far-end MAC address is inserted in Remote ID suboption of Option 82.

Label	Description
Managed by Service	Specifies the service-id of the management VPLS managing this SAP.
Managed by MSTI	Specifies the MST instance inside the management VPLS managing this SAP.
Last BPDU from	The bridge ID of the sender of the last BPDU received on this SAP.
Managed by SAP	Specifies the SAP ID inside the management VPLS managing this SAP.
Prune state	Specifies the STP state inherited from the management VPLS.
Managed by Service	Specifies the service ID of the management VPLS managing this SAP.
Last BPDU from	The bridge ID of the sender of the last BPDU received on this SAP.
Managed by Spoke	Specifies the SAP ID inside the management VPLS managing this spoke-SDP.
Prune state	Specifies the STP state inherited from the management VPLS.

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

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

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

Keyword to specify 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

[Table 50: Output fields: ARP](#) describes show service ID ARP output fields.

Table 50: 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

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, and SAPs.

Output

The following output is an example of basic service ID information, and [Table 51: Output fields: base](#) describes the output fields.

Sample output

```
A:ALU-7210>config>service>vpls# show service id 700 base
=====
Service Basic Information
=====
Service Id       : 700                Vpn Id           : 0
Service Type     : VPLS
Customer Id      : 1
Last Status Change: 09/19/2005 16:25:28
Last Mgmt Change : 09/21/2005 14:07:07
Admin State      : Down              Oper State       : Down
SAP Count        : 1
Uplink Type:     : L2
SAP Type:        : Null-star         Customer vlan:   : n/a
-----
Service Access & Destination Points
-----
Identifier                Type      AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/2                  null      1514    1514    Up   Down
=====
A:ALU-7210>config>service>vpls#
*A:ALU_SIM2>config>service>vpls# show service id 600 base
=====
Service Basic Information
=====
Service Id       : 600                Vpn Id           : 0
Service Type     : uVPLS
Customer Id      : 1
Last Status Change: 11/17/2008 00:12:16
Last Mgmt Change : 11/14/2008 03:36:21
Admin State      : Up                Oper State       : Down
SAP Count        : 3
Uplink Type:     : L2
SAP Type:        : Dot1q            Customer vlan:   : n/a
-----
Service Access & Destination Points
-----
Identifier                Type      AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/1:950             q-tag     1518    1518    Down Down
sap:1/1/2:1450            q-tag     1518    1518    Down Down
sap:1/1/3:3000.*         qinq      1522    1522    Down Down
=====
*A:ALU_SIM2>config>service>vpls#
*A:7210SAS-D>show>service# id 15 base
```

```

=====
Service Basic Information
=====
Service Id       : 15                Vpn Id          : 0
Service Type    : Epipe
Description     : (Not Specified)
Customer Id     : 40
Last Status Change: 01/01/1970 00:00:14
Last Mgmt Change  : 01/11/1970 23:34:24
Admin State     : Down              Oper State      : Down
SAP Count       : 1
Uplink Type:    : L2
SAP Type:       : Any              Customer vlan:  : n/a

-----
Service Access & Destination Points
-----
Identifier                Type      AdmMTU  OprMTU  Adm  Opr
-----
sap:1/1/1                null      1514    1514    Up   Down
=====
*A:SAS-D>show>service#
*A:7210SAS# show service id 10 base

=====
Service Basic Information
=====
Service Id       : 10                Vpn Id          : 0
Service Type    : VPLS
Name            : (Not Specified)
Description     : (Not Specified)
Customer Id     : 1
Last Status Change: 02/06/2106 06:28:12
Last Mgmt Change  : 01/10/1970 01:55:31
Admin State     : Down              Oper State      : Down
MTU             : Not Applicable    Def. Mesh VC Id : 10
SAP Count       : 0
Uplink Type:    : L2
SAP Type:       : Dot1q Range      Customer vlan:  : n/a

-----
Service Access & Destination Points
-----
Identifier                Type      AdmMTU  OprMTU  Adm  Opr
-----
No Matching Entries
=====
*A:7210SAS# show service id 10 base
    
```

Table 51: Output fields: base

Label	Description
Service Id	The service identifier.
Service Type	Displays the type of service.
Description	Generic information about the service.
Customer Id	The customer identifier.

Label	Description
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Adm	The administrative state of the service.
Oper	The operational state of the service.
Mtu	The largest frame size (in octets) that the port can handle.
Adm	The largest frame size (in octets) that the SAP can handle.
SAP Count	The number of SAPs defined on the service.
SAP Type	The type of SAPs allowed in the service. It also describes the applied processing by the node to the packets received on these SAPs.
Identifier	Specifies the service access (SAP).
OprMTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this port, without requiring the packet to be fragmented.
Opr	The operating state of the SAP.

bgp-ad

Syntax

bgp-ad

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command shows BGP AD related information for the service.

Output

The following output is an example of BGP AD information, and [Table 52: Output fields: BGP-AD](#) describes the fields.

Sample output

```
*A:dut-a>show>service# id 5 bgp-ad
```

BGP Auto-discovery Information

```
-----  
Admin State      : Up  
Vpls Id         : 100:5  
Prefix          : 1.1.1.1  
-----
```

*A:dut-a>show>service#

Table 52: Output fields: BGP-AD

Label	Description
Admin state	Displays the administrative state.
VPLS ID	Displays the VPLS ID.
Prefix	Displays the prefix.

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

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 detailed information.

expiry

Displays time until the MAC is aged out.

Output

The following output is an example of FDB information, and [Table 53: Output fields: FDB](#) describes the output fields.

Sample output

```
A:ALU-7210# show service id 1 fdb
=====
Forwarding Database, Service 1
=====
Service Id : 1 Mac Move : Disabled
Primary Factor : 3 Secondary Factor : 2
Mac Move Rate : 2 Mac Move Timeout : 10
Mac Move Retries : 3
Table Size : 250 Total Count : 1
Learned Count : 0 Static Count : 0
Host-learned Count: 1
Remote Age : 900 Local Age : 300
High Watermark : 95% Low Watermark : 90%
Mac Learning : Enabled Discard Unknown : Disabled
Mac Aging : Enabled Relearn Only : False
Mac Subnet Len : 48
=====
A:ALU-7210#
```

Table 53: Output fields: FDB

Label	Description
ServID	Displays the service ID.
MAC	Displays the associated MAC address.
Mac Move	Displays the administrative state of the MAC movement feature associated with this service.
Primary Factor	Displays a factor for the primary ports defining how many MAC-relearn periods should be used to measure the MAC-relearn rate.
Secondary Factor	Displays a factor for the secondary ports defining how many MAC-relearn periods should be used to measure the MAC-relearn rate.
Mac Move Rate	Displays the maximum rate at which MACs can be relearned in this service, before the SAP where the moving MAC was last seen is automatically disabled to protect the system against undetected loops or duplicate 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 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 will not be automatically re-enabled after being disabled. If after the

Label	Description
	SAP is re-enabled it is disabled again, the effective retry timeout is doubled to avoid thrashing.
Mac Move Retries	Displays the number of times retries are performed for re-enabling the SAP.
Table Size	Specifies the maximum number of learned and static entries allowed in the FDB of this service.
Total Count	Displays the total number of learned entries in the FDB of this service.
Learned Count	Displays the current number of learned entries in the FDB of this service.
Static Count	Displays the current number of static entries in the FDB of this service.
OAM-learned Count	Displays the current number of OAM entries in the FDB of this service.
DHCP-learned Count	Displays the current number of DHCP-learned entries in the FDB of this service.
Host-learned Count	Displays the current number of host-learned entries in the FDB of this service.
Remote Age	Displays the number of seconds used to age out FDB entries learned on an SAP. These entries correspond to MAC addresses learned on remote SAPs.
Local Age	Displays the number of seconds used to age out FDB entries learned on local SAPs.
High Watermark	Displays the utilization of the FDB table of this service at which a table full alarm will be raised by the agent.
Low Watermark	Displays the utilization of the FDB table of this service at which a table full alarm will be cleared by the agent.
Mac Learning	Specifies whether the MAC learning process is enabled
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded.
Mac Aging	Indicates whether the MAC aging process is enabled.
Relearn Only	When enabled, displays that either the FDB table of this service is full, or that the maximum system-wide number of MAs supported by the agent has been reached, and therefore MAC learning is temporary disabled, and only MAC relearns can take place.

Label	Description
Mac Subnet Len	Displays the number of bits to be considered when performing MAC-learning or MAC-switching.
Source-Identifier	The location where the MAC is defined.
Type/Age	Type — Specifies the number of seconds used to age out TLS FDB entries learned on local SAPs.
	Age — Specifies the number of seconds used to age out TLS FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs.
	L — Learned - Dynamic entries created by the learning process.
	OAM — Entries created by the OAM process.
	H — Host, the entry added by the system for a static configured subscriber host.
	D or DHCP — DHCP-installed MAC. Learned addresses can be temporarily frozen by the DHCP snooping application for the duration of a DHCP lease.
	P — Indicates the MAC is protected” by the MAC protection feature.
	Static — Statically configured.
Last Change	Indicates the time of the most recent state changes.

I2pt

Syntax

I2pt disabled

I2pt [detail]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays Layer 2 Protocol Tunnel (L2-PT) route information associated with this service.

Parameters

disabled

Displays only entries with termination disabled. This helps identify configuration errors.

detail

Displays detailed information.

Output

The following output is an example of L2PT information, and [Table 54: 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 54: Output fields: L2PT

Label	Description
Service id	Displays the 24 bit (0 to 16777215) service instance identifier for the service.
L2pt-term enabled	Indicates if L2-PT-termination or BPDU-translation is in use in this service by at least one SAP or spoke-SDP binding. If in use, at least one of L2PT-termination or BPDU-translation is enabled. When enabled it is not possible to enable STP on this service.
L2pt-term disabled	Indicates that L2-PT-termination is disabled.

Label	Description
Bpdu-trans auto	Specifies the number of L2-PT PDUs 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.
Sapld	The ID of the access point where this SAP is defined.
L2pt-termination	Indicates whether L2pt termination is enabled or disabled.
Admin Bpdu-translation	Specifies whether BPDU translation is administratively enabled or disabled.
Oper Bpdu-translation	Specifies whether BPDU translation is operationally enabled or disabled.
SAP Id	Specifies the SAP ID.

mac-move

Syntax

mac-move

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays MAC move related information about the service.

mrollers

Syntax

mrollers [detail]

Context

show>service>id>mld-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays all multicast routers.

```
sap
```

Syntax

```
sap sap-id [sap-id [detail | stp]]
```

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

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/mdal/port[.channel]* form. See [Common CLI command descriptions](#) for command syntax.

detail

Displays detailed information for the SAP.

stp

Displays the STP information of the SAP.

Output

The following output is an example of SAP information, and [Table 55: Output fields: SAP](#) describes the output fields.

Sample output

```
*A:ALU_SIM2>config>service# show service id 1 sap 1/1/3:500.* detail
=====
Service Access Points(SAP)
=====
Service Id      : 1
SAP             : 1/1/3:500.*          Encap           : qinq
```

```

QinQ Dot1p      : Default
Dot1Q Ethertype : 0x8100
QinQ Ethertype  : 0x8100

Admin State     : Up
Flags           : PortOperDown
Oper State      : Down
Last Status Change : 11/17/2008 00:26:56
Last Mgmt Change  : 11/14/2008 02:45:15
Loopback Mode   : Internal No-svc-port used : 1/1/13
Loopback Src Addr : 00:00:00:22:22:22
Loopback Dst Addr : 00:00:00:11:11:11
Max Nbr of MAC Addr: No Limit
Learned MAC Addr : 0
Ingress qos-policy : 1
Shared Q pncy    : n/a
Ingr IP Fltr-Id : n/a
Ingr Mac Fltr-Id : n/a
tod-suite       : TodResourceUnavail
Egr Agg Rate Limit : max
Mac Learning    : Enabled
Mac Aging       : Enabled
L2PT Termination : Disabled
Vlan-translation : None
Total MAC Addr  : 0
Static MAC Addr : 0
Egress qos-policy : 1
Multipoint shared : Disabled
Egr IP Fltr-Id  : n/a
Egr Mac Fltr-Id : n/a
qinq-pbit-marking : both
Discard Unkwn Srce: Disabled
Mac Pinning     : Disabled
BPDU Translation : Disabled

Acct. Pol       : None
Collect Stats   : Disabled

Anti Spoofing   : None
Nbr Static Hosts : 0
-----
Stp Service Access Point specifics
-----
Mac Move        : Blockable
Stp Admin State : Up
Core Connectivity : Down
Port Role       : Disabled
Port Number     : 2049
Port Path Cost  : 10
Admin Edge      : Disabled
Link Type       : Pt-pt
Root Guard      : Disabled
Last BPDU from  : N/A
CIST Desig Bridge : This Bridge
Stp Oper State  : Up
Port State      : Discarding
Port Priority    : 128
Auto Edge       : Enabled
Oper Edge       : False
BPDU Encap      : Dot1d
Active Protocol : Rstp
Designated Port : 34817

Forward transitions: 1
Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd    : 0
Cfg BPDUs tx      : 0
TCN BPDUs rcvd    : 0
TCN BPDUs tx      : 0
RST BPDUs rcvd    : 0
RST BPDUs tx      : 124267
MST BPDUs rcvd    : 0
MST BPDUs tx      : 0
-----
Sap Statistics
-----
Ingress Packets rcvd: 0
Packets
-----
Sap per Meter stats
-----
Ingress Meter 1 (Unicast)
For. InProf      : 0
For. OutProf     : 0
Octets
-----
Ingress Meter 11 (Multipoint)
For. InProf      : 0
For. OutProf     : 0
Octets
=====
*A:ALU_SIM2>config>service#
    
```

```

*A:PE-A# show service id 10 sap 2/2/5:10 mrp
=====
Service Access Points(SAP)
=====
Service Id      : 10
SAP            : 2/2/5:10          Encap          : q-tag
Description    : Default sap description for service id 10
Admin State    : Up                Oper State     : Up
Flags          : None
Multi Svc Site: None
Last Status Change : 01/16/2008 09:37:57
Last Mgmt Change  : 01/16/2008 09:37:41
-----
SAP MRP Information
-----
Join Time      : 0.2 secs          Leave Time     : 1.0 secs
Leave All Time : 10.0 secs         Periodic Time  : 1.0 secs
Periodic Enabled : false
Rx Pdus       : 11                Tx Pdus       : 12
Dropped Pdus  : 0                 Tx Pdus       : 12
Rx New Event  : 0                 Rx Join-In Event : 150
Rx In Event   : 10                Rx Join Empty Evt : 10
Rx Empty Event : 10                Rx Leave Event  : 0
Tx New Event  : 0                 Tx Join-In Event : 140
Tx In Event   : 0                 Tx Join Empty Evt : 20
Tx Empty Event : 10                Tx Leave Event  : 0
-----
SAP MMRP Information
-----
MAC Address      Registered      Declared
-----
01:1e:83:00:00:65 Yes                Yes
01:1e:83:00:00:66 Yes                Yes
01:1e:83:00:00:67 Yes                Yes
01:1e:83:00:00:68 Yes                Yes
01:1e:83:00:00:69 Yes                Yes
01:1e:83:00:00:6a Yes                Yes
01:1e:83:00:00:6b Yes                Yes
01:1e:83:00:00:6c Yes                Yes
01:1e:83:00:00:6d Yes                Yes
01:1e:83:00:00:6e Yes                Yes
-----
Number of MACs=10 Registered=10 Declared=10
-----
*A:PE-A#
    
```

Table 55: Output fields: 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.

Label	Description
Oper State	The operational state of the SAP.
Flags	Specifies the conditions that affect the operating status of this SAP. Display output includes: ServiceAdminDown, SapAdminDown, InterfaceAdminDown, PortOperDown, L2OperDown, SapIngressQoSMismatch, SapEgressQoSMismatch, RelearnLimitExceeded, ParentIfAdminDown, TodResourceUnavail, TodMssResourceUnavail, SapParamMismatch, SapIngressNamedPoolMismatch, SapEgressNamedPoolMismatch, NoSapEpipeRingNode.
Last Status Change	Specifies the time of the most recent operating status change to this SAP.
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SAP.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether collect stats is enabled.
SAP per Meter stats	
Ingress Meter	Specifies the meter ID.
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded.
For. OutProf	The number of out-of-profile packets and octets. (rate above CIR and below PIR) forwarded by the ingress meter.
Ingress TD Profile	The profile ID applied to the ingress SAP.
Egress TD Profile	The profile ID applied to the egress SAP.
Alarm Cell Handling	The indication that OAM cells are being processed.
AAL-5 Encap	The AAL-5 encapsulation type.
Loopback Mode	Displays the Ethernet port loopback mode.
Loopback Src Addr	Displays the configured loopback source address.
Loopback Dst Addr	Displays the configured loopback destination address.
No-svc-port used	Displays the port ID of the port on which no service is configured.

Label	Description
	This port is used for the port loop back with MAC swap functionality.

stp

Syntax

stp [detail]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays information for the STP instance for the service.

Parameters

detail

Displays detailed information.

Output

The following output is an example of STP information, and [Table 56: Output fields: STP](#) describes the output fields.

Sample output

```
*A:ALU_SIM2>config>service>vpls>sap# show service id 1 stp
=====
Stp info, Service 1
=====
Bridge Id       : 00:00.00:45:67:32:10:ab  Top. Change Count : 2
Root Bridge    : This Bridge              Stp Oper State    : Up
Primary Bridge : N/A                      Topology Change   : Inactive
Mode           : Rstp                      Last Top. Change  : 1d 18:34:36
Vcp Active Prot. : N/A
Root Port      : N/A                      External RPC      : 0
=====
Stp port info
=====
Sap Id          Oper-   Port-   Port-   Port-   Oper-   Link-   Active
                State  Role   State  Num    Edge   Type   Prot.
-----
1/1/21:1       Up     Designated Forward  2048   True   Pt-pt  Rstp
lag-1:1.1      Up     Designated Forward  2049   False  Pt-pt  Rstp
lag-2:1.1      Up     Designated Forward  2050   False  Pt-pt  Rstp
=====
*A:ALU_SIM2>config>service>vpls>sap#
```



```
*A:ALU_SIM2>config>service>vpls>sap# show service id 1 stp detail
=====
Spanning Tree Information
=====
-----
VPLS Spanning Tree Information
-----
VPLS oper state      : Up
Stp Admin State     : Up
Mode                : Rstp
Stp Oper State      : Up
Vcp Active Prot.    : N/A

Bridge Id           : 00:00:00:45:67:32:10:ab
Bridge Priority      : 0
Topology Change     : Inactive
Last Top. Change    : 1d 18:35:18
Top. Change Count   : 2
Bridge Instance Id  : 0
Tx Hold Count       : 6
Bridge Hello Time   : 2
Bridge Max Age      : 20
Bridge Fwd Delay    : 15

Root Bridge         : This Bridge
Primary Bridge      : N/A

Root Path Cost      : 0
Rcvd Hello Time    : 2
Root Priority        : 0
Root Forward Delay  : 15
Root Max Age        : 20
Root Port           : N/A
-----
Spanning Tree Sap Specifics
-----
SAP Identifier      : 1/1/21:1
Port Role           : Designated
Port Number         : 2048
Port Path Cost      : 10
Admin Edge          : Disabled
Link Type           : Pt-pt
Root Guard          : Disabled
Last BPDUs from    : N/A
CIST Desig Bridge   : This Bridge
Forward transitions : 2
Cfg BPDUs rcvd     : 0
TCN BPDUs rcvd     : 0
RST BPDUs rcvd     : 0
MST BPDUs rcvd     : 0
Stp Admin State     : Up
Port State          : Forwarding
Port Priority        : 128
Auto Edge           : Enabled
Oper Edge           : True
BPDU Encap          : Dot1d
Active Protocol     : Rstp
Designated Port     : 34816
Bad BPDUs rcvd     : 0
Cfg BPDUs tx        : 0
TCN BPDUs tx        : 0
RST BPDUs tx        : 79126
MST BPDUs tx        : 0

SAP Identifier      : lag-1:1.1
Port Role           : Designated
Port Number         : 2049
Port Path Cost      : 10
Admin Edge          : Disabled
Link Type           : Pt-pt
Root Guard          : Disabled
Last BPDUs from    : 10:00:00:f3:f0:98:97:00
CIST Desig Bridge   : This Bridge
Forward transitions : 1
Cfg BPDUs rcvd     : 0
TCN BPDUs rcvd     : 0
RST BPDUs rcvd     : 5
MST BPDUs rcvd     : 0
Stp Admin State     : Up
Port State          : Forwarding
Port Priority        : 128
Auto Edge           : Enabled
Oper Edge           : False
BPDU Encap          : Dot1d
Active Protocol     : Rstp
Designated Port     : 34817
Bad BPDUs rcvd     : 0
Cfg BPDUs tx        : 0
TCN BPDUs tx        : 0
RST BPDUs tx        : 79128
MST BPDUs tx        : 0

SAP Identifier      : lag-2:1.1
Port Role           : Designated
Port Number         : 2050
Port Path Cost      : 10
Admin Edge          : Disabled
Link Type           : Pt-pt
Root Guard          : Disabled
Stp Admin State     : Up
Port State          : Forwarding
Port Priority        : 128
Auto Edge           : Enabled
Oper Edge           : False
BPDU Encap          : Dot1d
Active Protocol     : Rstp
```

```

Last BPDU from      : 20:00.00:e0:b1:99:cb:2a
CIST Desig Bridge  : This Bridge           Designated Port   : 34818
Forward transitions: 1                     Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                     Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                     TCN BPDUs tx      : 0
RST BPDUs rcvd     : 6                     RST BPDUs tx      : 78760
MST BPDUs rcvd     : 0                     MST BPDUs tx      : 0
=====
*A:ALU_SIM2>config>service>vpls>sap#
    
```

Table 56: Output fields: STP

Label	Description
Bridge-id	Specifies the MAC address used to identify this bridge in the network.
Bridge fwd delay	Specifies how fast a bridge changes its state when moving toward the forwarding state.
Bridge Hello time	Specifies the amount of time between the transmission of Configuration BPDUs.
Bridge max age	Specifies the maximum age of STP 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 STP 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 STP instance associated with this service.
Top. change count	Specifies the total number of topology changes detected by the STP 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 STP instance associated with this service. This value is used as the Root Identifier parameter in all configuration BPDUs originated by this node.
Root path cost	Specifies the cost of the path to the root bridge as seen from this bridge.
Root forward delay	Specifies how fast the root changes its state when moving toward the forwarding state.
Rcvd hello time	Specifies the amount of time between the transmission of configuration BPDUs.

Label	Description
Root max age	Specifies the maximum age of STP 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.
BPDU encap	Specifies the type of encapsulation used on BPDUs sent out and received on this SAP.
Port Number	Specifies the value of the port number field which is contained in the least significant 12 bits of the 16-bit port ID associated with this SAP.
Port Priority	Specifies the value of the port priority field which is contained in the most significant 4 bits of the 16-bit port ID associated with this SAP.
Port Path Cost	Specifies the contribution of this port to the path cost of paths toward the spanning tree root which include this port.
Designated Port	Specifies the port identifier of the port on the designated bridge for this port's segment.
Designated Bridge	Specifies the bridge identifier of the bridge which this port considers to be the designated bridge for this port's segment.

mstp-configuration

Syntax

mstp-configuration

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the MSTP specific configuration data. This command is valid only on a management VPLS.

Output

The following output is an example of MSTP configuration information, and [Table 57: Output Fields: MSTP configuration](#) describes output fields.

Sample output

```
A:SASD1>show>service>id# stp mst-instance 1

=====
MSTP specific info for service 1 MSTI 1
=====
Regional Root      : This Bridge          Root Port      : N/A
Internal RPC       : 0                    Remaining Hopcount: 20

=====
MSTP port info for MSTI 1
=====
Sap Id             Oper-   Port-   Port-   Port-   Same
                   State  Role   State  Num    Region
-----
1/1/2:0.*          Up     Designated Forward  2048   True
1/1/5:0.*          Up     Alternate Discard  2049   False
1/1/6:0.*          Up     Master   Forward  2050   False
=====
A:SASD1>show>service>id#
```

Table 57: Output Fields: MSTP configuration

Label	Description
Region Name	Displays the MSTP region name.
Region Revision	Displays the MSTP region revision.
MST Max Hops	Displays the MSTP maximum hops specified.
Instance	Displays the MSTP instance number.
Priority	Displays the MSTP priority.
Vlans mapped	Displays the VLAN range of the MSTP instance.

dhcp

Syntax

dhcp

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Commands in this context display DHCP information for the specified service.

statistics

Syntax

statistics [**sap** *sap-id*]

statistics [**sdp** *sdp-id:vc-id*]

statistics [**interface** *interface-name*]

Context

show>service>id>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

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 58: 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 58: 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 and 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 of 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.

summary

Syntax

summary [**interface** *interface-name*]

Context

show>service>id>dhcp

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Displays DHCP configuration summary information.

Parameters

interface *interface-name*

Displays information for the specified IP interface.

Output

The following output is an example of DHCP summary information, and [Table 59: 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 59: Output fields: DHCP summary

Label	Description
Interface Name	Name of the router interface.
Arp Populate	Specifies whether ARP populate is enabled. 7210 SAS does not support ARP populate.
Used/Provided	7210 SAS does not maintain lease state.

Label	Description
Info Option	Indicates whether Option 82 processing is enabled on the interface.
Admin State	Indicates the administrative state.

5.7.2.3 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

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

Description

This command displays more information for all aspects of IGMP snooping on the VPLS service.

Output

The following output is an example of detailed IGMP snooping information, and [Table 60: Output fields: IGMP snooping all](#) describes the output fields.

Sample output

```
*A:7210-2>show>service>id>snooping# all
=====
IGMP Snooping info for service 900
=====
IGMP Snooping Base info
-----
Admin State : Up
Querier      : 5.5.5.5 on SAP 1/1/14:100
-----
Sap          Oper   MRtr  Send   Max Num  Num
Id           State  Port  Queries Groups   Groups
-----
sap:1/1/13:100  Up    No    Disabled No Limit 0
sap:1/1/14:100  Up    No    Disabled No Limit 0
sap:1/1/17:100  Up    No    Disabled No Limit 1
sap:1/1/18:100  Up    No    Disabled No Limit 1
-----
IGMP Snooping Querier info
-----
Sap Id       : 1/1/14:100
IP Address   : 10.10.10.2
Expires      : 254s
Up Time      : 0d 00:02:42
Version      : 2

General Query Interval : 125s
Query Response Interval : 10.0s
Robust Count           : 2
-----
IGMP Snooping Multicast Routers
-----
MRouter      Sap Id           Up Time           Expires           Version
-----
10.10.10.2   1/1/14:100       0d 00:02:44      254s              2
-----
Number of mrouter: 1
-----
IGMP Snooping Proxy-reporting DB
-----
Group Address  Up Time
-----
239.6.6.6     0d 00:02:17
239.9.9.9     0d 00:02:56
-----
Number of groups: 2
-----
IGMP Snooping SAP 1/1/13:100 Port-DB
-----
Group Address  Type   From-
VPLS  Up Time   Expires
-----
Number of groups: 0
-----
IGMP Snooping SAP 1/1/14:100 Port-DB
-----
Group Address  Type   From-
VPLS  Up Time   Expires
=====
*A:7210-2>show>service>id>snooping#
```

Table 60: Output fields: IGMP snooping all

Label	Description
Admin State	The administrative state of the IGMP instance.
Querier	Displays the address of the IGMP querier on the IP subnet to which the interface is attached.
Sap Id	Displays the SAP IDs of the service ID.
Oper State	Displays the operational state of the SAP IDs of the service ID.
Mrtr Port	Specifies if the port is a multicast router port.
Send Queries	Specifies whether the send-queries command is enabled or disabled.
Max Num Groups	Specifies the maximum number of multicast groups that can be joined on this SAP.
MVR From VPLS	Specifies MVR from VPLS.
Num MVR Groups	Specifies the actual number of multicast groups that can be joined on this SAP.
MVR From VPLS Cfg Drops	Displays the from VPLS drop count.
MVR To SAP Cfg Drops	Displays the to SAP drop count.
MVR Admin State	Displays the administrative state of MVR.
MVR Policy	The MVR policy name.

mfib

Syntax

mfib [brief]

mfib [group *grp-address*]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays the multicast FIB on the VPLS service.

Parameters

brief

Displays a brief output.

group *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 61: Output fields: service 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
-----
239.4.4.4              sap:1/1/1            Local    Fwd
-----
Number of entries: 1
=====

*A:7210-2# show service id 5 mfib
=====
Multicast FIB, Service 5
=====
Source Address  Group Address        Sap Id           Svc Id   Fwd/Blk
-----
*               239.5.34.1          sap:lag-1:5.*    Local    Fwd
10.5.5.1        239.5.34.1          sap:lag-1:5.*    Local    Blk
                239.5.34.1          sap:lag-2:5.*    Local    Fwd
*               239.5.34.2          sap:lag-1:5.*    Local    Fwd
10.5.5.1        239.5.34.2          sap:lag-1:5.*    Local    Blk
                239.5.34.2          sap:lag-2:5.*    Local    Fwd
*               239.5.34.3          sap:lag-1:5.*    Local    Fwd
10.5.5.1        239.5.34.3          sap:lag-1:5.*    Local    Blk
                239.5.34.3          sap:lag-2:5.*    Local    Fwd
*               239.5.34.4          sap:lag-1:5.*    Local    Fwd
10.5.5.1        239.5.34.4          sap:lag-1:5.*    Local    Blk
                239.5.34.4          sap:lag-2:5.*    Local    Fwd
*               239.5.34.5          sap:lag-1:5.*    Local    Fwd
10.5.5.1        239.5.34.5          sap:lag-1:5.*    Local    Blk
                239.5.34.5          sap:lag-2:5.*    Local    Fwd
-----
Number of entries: 10
=====
*A:7210-2
```

Table 61: Output fields: service MFIB

Label	Description
Group Address	IPv4 multicast group address.
SAP ID	Indicates the SAP to which the corresponding multicast stream will be forwarded or blocked.
Forwarding/Blocking	Indicates whether the corresponding multicast stream will be blocked or forwarded.
Number of Entries	Specifies the number of entries in the MFIB.
Forwarded Packets	Indicates the number of multicast packets forwarded for the corresponding source or group.
Forwarded Octets	Indicates the number of octets forwarded for the corresponding source or group.
Svc ID	Indicates the service to which the corresponding multicast stream will be forwarded or blocked. Local means that the multicast stream will be forwarded or blocked to a SAP local to the service.

mroute

Syntax

mroute [detail]

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

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 62: Output fields: IGMP-snooping mroute](#) describes the output fields.

Sample output

```
*A:7210-2 show service id 3 igmp-snooping mrouter
=====
IGMP Snooping Multicast Routers for service 3
=====
MRouter      Sap Id      Up Time      Expires      Version
-----
10.20.1.2    lag-2:3.*   0d 02:30:39  197s         3
-----
Number of mrouter: 1
=====
*A:SNPG-1#

*A:7210-2# show service id 3 igmp-snooping mrouter detail
=====
IGMP Snooping Multicast Routers for service 3
=====
MRouter 10.20.1.2
-----
Sap Id      : lag-2:3.*
Expires     : 181s
Up Time     : 0d 02:30:55
Version     : 3

General Query Interval : 125s
Query Response Interval : 10.0s
Robust Count   : 2
-----
Number of mrouter: 1
=====
*A:7210-2#
```

Table 62: Output fields: IGMP-snooping mrouter

Label	Description
MRouter	Specifies the multicast router port.
Sap/Sdp Id	Specifies the SAP and SDP ID multicast router ports.
Up Time	Displays the length of time the mrouter has been up.
Expires	Displays the amount of time left before the query interval expires.
Version	Displays the configured version of IGMP running on this interface.
General Query Interval	The frequency at which host-query packets are transmitted.
Query Response Interval	The time to wait to receive a response to the host-query message from the host.
Robust Count	Specifies the value used to calculate several IGMP message intervals.

Label	Description
Number of Mrouters	Displays 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

Description

This command displays Multicast VPLS Registration (MVR) information.

port-db

Syntax

port-db sap *sap-id* [detail]

port-db sap *sap-id* group *grp-ip-address*

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

group *grp-address*

Displays IGMP snooping statistics matching the specified group address.

source ip-address

Displays IGMP snooping statistics matching one particular source within the multicast group.

Output

The following output is an example of IGMP snooping port database information, and [Table 63: Output fields: IGMP snooping port DB](#) describes the output fields.

Sample output

```
*A:7210-2>show>service>id>snooping# port-db sap 1/1/18:100
=====
IGMP Snooping SAP 1/1/18:100 Port-DB for service 900
=====
Group Address      Type      From-VPLS  Up Time      Expires
-----
239.6.6.6          dynamic  local      0d 00:12:35  260s
-----
Number of groups: 1
=====
*A:7210-2>show>service>id>snooping#
*A:7210-2# show service id 5 igmp-snooping port-db sap lag-1:5.*
=====
IGMP Snooping SAP lag-1:5.* Port-DB for service 5
=====
Group Address      Mode      Type      Up Time      Expires      Num
                               Src
-----
239.5.34.1         exclude  dynamic  0d 00:23:38  216s        1
239.5.34.2         exclude  dynamic  0d 00:23:38  216s        1
239.5.34.3         exclude  dynamic  0d 00:23:38  216s        1
239.5.34.4         exclude  dynamic  0d 00:23:38  216s        1
239.5.34.5         exclude  dynamic  0d 00:23:38  216s        1
-----
Number of groups: 5
=====
*A:7210-2#
=====
*A:7210-2>show>service>id>snooping# port-db sap 1/1/18:100 detail
=====
IGMP Snooping SAP 1/1/18:100 Port-DB for service 900
-----
IGMP Group 239.6.6.6
-----
Type                : dynamic
Up Time             : 0d 00:12:39      Expires            : 259s
Compat Mode         : IGMP Version 2
V1 Host Expires     : 0s              V2 Host Expires   : 259s
-----
Number of groups: 1
=====
*A:7210-2>show>service>id>snooping#
```

Table 63: Output fields: IGMP snooping port DB

Label	Description
Group Address	The IP multicast group address for which this entry contains information.
Mode	<p>Specifies the type of membership reports received on the interface for the group.</p> <p>In 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.</p> <p>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.</p>
Type	<p>Indicates how this group entry was learned.</p> <p>If this group entry was learned by IGMP, the value is set to dynamic.</p> <p>For statically configured groups, the value is set to static.</p>
Compatibility mode	<p>Specifies the IGMP mode. This is used so routers are compatible with earlier version routers. IGMPv3 hosts must operate</p> <p>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 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.</p>
V1 host expires	The time remaining until the local router will assume that there are no longer any IGMP Version 1 members on the IP subnet attached to this interface. Upon hearing any IGMPv1 membership report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv2 leave messages for this group that it receives on this interface.
V2 host expires	The time remaining until the local router will assume that there are no longer any IGMP Version 2 members on the IP subnet attached to this interface. Upon hearing any IGMPv2 membership report, this value is reset to the group membership timer. While this time remaining is non-zero, the local router ignores any IGMPv3 leave messages for this group that it receives on this interface.

Label	Description
Source address	The source address for which this entry contains information.
Up Time	The time since the source group entry was created.
Expires	The amount of time remaining before this entry will be aged out.
Number of sources	Indicates the number of IGMP group and source specific queries received on this SAP.
Forwarding/Blocking	Indicates whether this entry is on the forward list or block list.
Number of groups	Indicates the number of groups configured for this SAP.

proxy-db

Syntax

proxy-db [detail]

proxy-db group *grp-ip-address*

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

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 IGMP snooping proxy reporting database information, and [Table 64: Output fields: IGMP snooping proxy-database](#) describes the output fields.

Sample output

```
*A:7210-2 show service id 5 igmp-snooping proxy-db
=====
IGMP Snooping Proxy-reporting DB for service 5
=====
Group Address      Mode      Up Time      Num Sources
```

```

-----
239.5.34.1      exclude 0d 00:25:54  0
239.5.34.2      exclude 0d 00:25:54  0
239.5.34.3      exclude 0d 00:25:54  0
239.5.34.4      exclude 0d 00:25:54  0
239.5.34.5      exclude 0d 00:25:54  0
-----
Number of groups: 5
=====
*A:7210-2

*A:7210-2>show>service>id>snooping# proxy-db detail
=====
IGMP Snooping Proxy-reporting DB for service 900
-----
IGMP Group 239.6.6.6
-----
Up Time : 0d 00:15:17
-----
Source Address  Up Time
-----
No sources.
-----
IGMP Group 239.9.9.9
-----
Up Time : 0d 00:15:56
-----
Source Address  Up Time
-----
No sources.
-----
Number of groups: 2
=====
*A:7210-2>show>service>id>snooping#
    
```

Table 64: Output fields: IGMP snooping proxy-database

Label	Description
Group Address	The IP multicast group address for which this entry contains information.
Mode	Specifies 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 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.
Num Sources	Indicates the number of IGMP group and source specific queries received on this interface.
Number of groups	Number of IGMP groups.

Label	Description
Source Address	The source address for which this entry contains information.

querier

Syntax

querier

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

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 65: Output fields: IGMP snooping queriers](#) describes the output fields.

Sample output

```
*A:7210-2# show service id 3 igmp-snooping querier
=====
IGMP Snooping Querier info for service 3
=====
Sap Id           : lag-2:3.*
IP Address       : 10.20.1.2
Expires          : 226s
Up Time          : 0d 02:30:11
Version          : 3

General Query Interval : 125s
Query Response Interval : 10.0s
Robust Count           : 2
=====
*A:7210-2#
```

Table 65: Output fields: IGMP snooping queriers

Label	Description
SAP Id	Specifies the SAP ID of the service.
IP address	Specifies the IP address of the querier.
Expires	The time left, in seconds, that the query will expire.

Label	Description
Up time	The length of time the query has been enabled.
Version	The configured version of IGMP.
General Query Interval	The frequency at which host-query packets are transmitted.
Query Response Interval	The time to wait to receive a response to the host-query message from the host.
Robust Count	Specifies the value used to calculate several IGMP message intervals.

static

Syntax

static [**sap** *sap-id*]

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

Output

The following is an example of static IGMP snooping information, and [Table 66: Output fields: IGMP snooping static](#) describes the output fields.

Sample output

```
*A:7210-2# show service id 4093 igmp-snooping static
=====
IGMP Snooping Static Source Groups for service 4093
=====
-----
IGMP Snooping Static Source Groups for SAP 1/1/5:4093
-----
Source          Group
-----
```

```

10.93.93.1      239.93.39.1
10.93.93.1      239.93.39.2
10.93.93.1      239.93.39.3
10.93.93.1      239.93.39.4
10.93.93.1      239.93.39.5
-----
Static (*,G)/(S,G) entries: 5
-----
IGMP Snooping Static Source Groups for SAP lag-3:4093
-----
Source          Group
-----
10.93.93.1      239.93.39.1
10.93.93.1      239.93.39.2
10.93.93.1      239.93.39.3
10.93.93.1      239.93.39.4
10.93.93.1      239.93.39.5
-----
Static (*,G)/(S,G) entries: 5
=====
*A:7210-2#
    
```

Table 66: Output fields: IGMP snooping 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*]

Context

show>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

Output

The following output is an example of IGMP snooping statistics, and [Table 67: Output fields: IGMP-snooping statistics](#) describes the output fields.

Sample output

```
*A:7210-2# show service id 5 igmp-snooping statistics
=====
IGMP Snooping Statistics for service 5
=====
Message Type           Received      Transmitted   Forwarded
-----
General Queries        0             228           0
Group Queries          0             0             0
Group-Source Queries  0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V3 Reports              282           0             0
V2 Leaves              0             0             0
Unknown Type           0             N/A           0
-----
Drop Statistics
-----
Bad Length              : 0
Bad IP Checksum         : 0
Bad IGMP Checksum      : 0
Bad Encoding            : 0
No Router Alert        : 0
Zero Source IP         : 0
Wrong Version           : 0

Send Query Cfg Drops   : 0
Import Policy Drops    : 0
Exceeded Max Num Groups : 0
Exceeded Max Num Sources : 0
=====
*A:7210-2#
```

Table 67: Output fields: IGMP-snooping statistics

Label	Description
Message Type	The column heading for IGMP or MLD snooping messages.
General Queries	The number of general query messages received, transmitted, and forwarded.
Group Queries	The number of group query messages received, transmitted, and forwarded.
Group-Source Queries	The number of group-source query messages received, transmitted, and forwarded.
V1 Reports	The number of IGMPv1 or MLDv1 report messages received, transmitted, and forwarded.
V2 Reports	The number of IGMPv2 or MLDv2 report messages received, transmitted, and forwarded.

Label	Description
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.
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.

5.7.2.4 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

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.

fdb

Syntax

fdb {**all** | **mac** *ieee-address* | **sap** *sap-id*}

Context

clear>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears FDB entries for the service.

Parameters

all

Clears all FDB entries.

mac *ieee-address*

Clears only FDB entries in the FDB table with the specified 48-bit MAC address. The MAC address can be expressed in the form *aa:bb:cc:dd:ee:ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

sap

Syntax

sap *sap-id* {**all** | **counters** | **stp**}

Context

clear>service>statistics

Platforms

Supported on all 7210 SAS platforms as described in this document

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

Description

This command clears all traffic queue counters associated with the service ID.

```
stp
```

Syntax

```
stp
```

Context

```
clear>service>statistics>id
```

Platforms

Supported on all 7210 SAS platforms as described in this document

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

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. 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

Description

Commands in this context clear IGMP snooping data.

port-db

Syntax

port-db [**sap** *sap-id*] [**group** *grp-address*]

Context

clear>service>id>igmp-snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

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

Description

This command clears the information about the IGMP snooping queriers for the VPLS service.

statistics

Syntax

statistics { **all** | **sap** *sap-id* }

Context

clear>service>id>snooping

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command clears IGMP snooping statistics.

Parameters

all

Clears all statistics for the service ID.

sap *sap-id*

Clears statistics for the specified SAP ID.

5.7.2.5 VPLS debug commands

id

Syntax

id *service-id*

Context

debug>service

Platforms

Supported on all 7210 SAS platforms as described in this document

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}

Context

debug>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

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.

sap

Syntax

[no] sap *sap-id*

Context

debug>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

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

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

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

Description

This command enables STP debugging for received and transmitted BPDUs.

exception

Syntax

[no] exception

Context

debug>service>id>stp

Platforms

Supported on all 7210 SAS platforms as described in this document

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

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

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

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

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

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.

6 Internet Enhanced Service on 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

This chapter provides information about Internet Enhanced Services, the process overview, and implementation notes.

6.1 IES service overview



Note:

- IES is used for in-band management of the node on the 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T.
- IES service (standalone, without being associated with a routed VPLS) is not supported on the 7210 SAS-K 2F1C2T.


Internet Enhanced Service (IES) is a routed connectivity service where the subscriber communicates with an IP router interface to send and receive Internet traffic. An IES has one or more logical IP routing interfaces each with a SAP which acts as the access point to the subscriber's network.

IES allows IP interfaces to participate in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber be unique between other provider addressing schemes and potentially the entire Internet. While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be reserved for service IP provisioning, and be administered by a separate, but subordinate address authority.

IP interfaces defined within the context of an IES service must have a SAP associated as the uplink access point to the subscriber network. Multiple IES services are created to segregate subscriber owned IP interfaces.

The following figure shows the Internet enhanced service.

Figure 79: 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.

On 7210 SAS-D and 7210 SAS-Dxp, 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.

6.2 IES features

This section describes various general service features and any special capabilities or considerations as they relate to IES services.

6.2.1 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:

- ICMP Options

6.2.2 IPv6 support for IES IP interfaces associated with access-uplink SAPs (on 7210 SAS-D and 7210 SAS-Dxp)

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-D and 7210 SAS-Dxp.

On 7210 SAS-D and 7210 SAS-Dxp, 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 **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 *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide*. A separate route table is used for IPv6 /128-bit prefix route lookup. A limited number of IPv6 /128 prefixes route lookup entries is supported on 7210 SAS-D. The software enables lookups in this table by default (that is, no user configuration is required to enable IPv6 /128-bit route lookup).



Note:

On the 7210 SAS-D and 7210 SAS-Dxp, IPv6 interfaces are created without allocating IPv6 route entries.

Following features and restrictions is applicable for IPv6 IES IP interfaces:

- IPv6 interfaces supports only static routing.
- Only port-based ingress QoS policies are supported.
- IPv6 filter policies can be used on SAP ingress and egress.

- Routing protocols, such as OSPFv3, and others are not supported.
- A limited number of IPv6 /128 prefixes route lookup entries is supported.

6.3 SAPs

6.3.1 Encapsulations

For 7210 SAS-D and 7210 SAS-Dxp, the following Access SAP encapsulation is supported on IES services in both network mode and access-uplink mode:

- Ethernet null
- Ethernet dot1q
- Ethernet QinQ

For 7210 SAS-D and 7210 SAS-Dxp, Ethernet QinQ (access-uplink QinQ SAP) encapsulation is supported.

6.3.2 CPE connectivity check



Note:

This capability is only supported on the 7210 SAS-D and 7210 SAS-Dxp.

Static routes are used within many IES services. Unlike dynamic routing protocols, there is no way to change the state of routes based on availability information for the associated CPE. CPE connectivity check adds flexibility so that unavailable destinations will be removed from the service provider's routing tables dynamically and minimize wasted bandwidth.

The availability of the far-end static route is monitored through periodic polling. The polling period is configured. If the poll fails a specified number of sequential polls, the static route is marked as inactive.

An ICMP ping mechanism is used to test the connectivity. If the connectivity check fails and the static route is deactivated, the router will continue to send polls and re-activate any restored routes.

6.3.3 QoS policies

When applied to 7210 SAS IES services, service ingress QoS policies only create the unicast meters defined in the policy. The multi-point meters are not created on the service. With IES services, service egress QoS policies function as with other services where the class-based queues are created as defined in the policy.

On 7210 SAS ingress, only meters are supported on all the platforms.



Note:

- QoS policies only create the unicast meters defined in the policy if PIM is not configured on the associated IP interface; if PIM is configured, the multipoint meters are applied as well.
- On the 7210 SAS-K 2F1C2T, both unicast and multicast queues are configured.

In access-uplink mode, IES IP interface associated with an access SAP supports use of service ingress QoS policies. IES IP interface associated with an access-uplink SAP does not support use of service ingress QoS policies. IES IP interfaces associated with an access-uplink SAP share the port based ingress and egress QoS policies.

Note that both MAC and IPv4 criteria can be used in the QoS policies for traffic classification in an IES.

6.3.4 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.

6.3.5 Filter policies

In access-uplink mode, only IP filter policies can be applied to IES service when either access SAP or access-uplink SAP is associated with the service.

6.4 Configuring an IES service with CLI

This section provides information to configure IES services using the command line interface.

6.4.1 Basic configuration

The most basic IES service configuration has the following entities:

- customer ID (see [Configuring customer 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 IES service on ALA-48 configuration output 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#
```

Example

The following is a sample IES service configuration output 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#
```

Example

The following is a sample basic IES service configuration output for IPv6, along with the use of **max-ipv6-routes**.

The following is a sample of allocation of IPv6 routes on the node.

```
*A:7210SAS>config>system>res-prof# info
-----
max-ipv6-routes 1000
-----
```



Note:

The node must be rebooted after the preceding change.

```
*A:7210SAS>config>service# info
-----
ies 1000 customer 50 vpn 1000 create
description "to inband-mgmt"
interface "to-mgmt" create
ipv6
address 10::1/24
sap 1/1/10:100 create
exit
exit
no shutdown
-----
*A:7210SAS>config>service#
```

6.4.2 Common configuration tasks

This section provides a brief overview of the tasks that must be performed to configure IES services and provides the CLI commands:

1. Associate an IES service with a customer ID.
2. Associate customer ID with the service.
3. Assign an IP address.

4. Create an interface.
5. Define SAP parameters on the interface:
 - Select nodes and ports.
 - Optional - select filter policies (configured in the **config>filter** context).
6. Enable service.

6.4.3 Configuring IES components

6.4.3.1 Configuring an IES service

Use the following syntax to create an IES service.

Example

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#
```

6.4.3.2 Configuring IES interface parameters

Example

The following is a sample IES configuration output 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#
```

6.4.3.3 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

The following is a sample 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#
```

6.4.3.4 Configuring SAP parameters

A SAP is a combination of a port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique within a router.

When configuring IES access SAP parameters, a default QoS policy is applied to each SAP ingress. Additional QoS policies must be configured in the **config>qos** context.

Filter policies are configured in the config>filter context and must be explicitly applied to a SAP. There are no default filter policies.

For 7210 SAS-D and 7210 SAS-Dxp, SAP ingress Qos policy is supported only for access SAPs. It is not supported for access-uplink SAP.

Access-uplink SAPs (on 7210 SAS-D and 7210 SAS-Dxp) use the port-based ingress and egress QoS policies.

Example

The following is a sample 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#
```


6.4.4 Service management tasks

This section describes the service management tasks.

6.4.4.1 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 and SAP information) and then enter the new information.

Example

The following is a sample modified service configuration output.

```
*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#
```

6.4.4.2 Deleting an IES service

An IES service cannot be deleted until SAPs and interfaces are shut down and deleted, and the service is shut down 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
```

6.4.4.3 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
```

6.4.4.4 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
```

6.5 IES services command reference

6.5.1 Command hierarchies

- [Global IES commands for 7210 SAS-D and 7210 SAS-Dxp](#)
- [IES interface commands](#)
- [IES interface routed VPLS commands for 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T](#)
- [IES interface SAP commands for 7210 SAS-D and 7210 SAS-Dxp](#)
- [IES interface SAP statistics commands for 7210 SAS-D and 7210 SAS-Dxp](#)
- [IES interface SAP QoS and filter commands for 7210 SAS-D and 7210 SAS-Dxp](#)
- [IES interface IPv6 commands \(applicable only to access-uplink SAPs on 7210 SAS-D and 7210 SAS-Dxp\)](#)
- [Show commands](#)

6.5.1.1 Global IES commands for 7210 SAS-D and 7210 SAS-Dxp

```
config  
- service  
  - ies service-id [customer customer-id] [create] [vpn vpn-id]  
  - no ies service-id  
    - description description-string  
    - no description  
    - interface  
    - no interface  
    - service-name service-name  
    - no service-name  
    - [no] shutdown
```

6.5.1.2 IES interface commands

```
config
```

```

- service
  - ies service-id [customer customer-id] [create] [vpn vpn-id]
  - [no] interface ip-int-name [create]
    - address {ip-address/mask | ip-address netmask} [broadcast {all-ones | host-
ones}]
    - no address {ip-address/mask | ip-address netmask}
    - arp-timeout seconds
    - no arp-timeout
    - delayed-enable seconds
    - no delayed-enable
    - description description-string
    - no description
    - dhcp
      - description description-string
      - no description
      - gi-address ip-address [src-ip-addr]
      - no gi-address
      - [no] option
        - action {replace | drop | keep}
        - no action
        - [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 server
        - server server1 [server2 ... (up to 8 max)]
        - [no] shutdown
        - [no] trusted
      - icmp
        - mask-reply
        - 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
      - ipv6
        - [no] urpf-check
      - [no] loopback
      - [no] shutdown
      - static-arp ip-address ieee-address
      - no static-arp ip-address [ieee-address]
      - [no] urpf-check

```

6.5.1.3 IES interface routed VPLS commands for 7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

```

config
- service
  - ies service-id [customer customer-id] [vpn vpn-id]
  - interface ip-interface-name [create]
  - no interface-interface-name
  - vpls service-name

```

```
- no vpls
- ingress
- v4-routed-override-filter ip-filter-id
- no v4-routed-override-filter
```

6.5.1.4 IES interface SAP commands for 7210 SAS-D and 7210 SAS-Dxp

```
config
- service
- ies service-id [customer customer-id] [create]
- [no] interface ip-int-name
- [no] sap sap-id [create]
- accounting-policy acct-policy-id
- no accounting-policy
- [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-in-kbytes
- cbs size [kbits | bytes | kbytes]
- no cbs
- mbs size-in-kbits
- mbs size [kbits | bytes | kbytes]
- no mbs
- no mode mode
- no mode
- rate cir cir-rate [pir pir-rate]
- [no] shutdown
```

6.5.1.5 IES interface SAP statistics commands for 7210 SAS-D and 7210 SAS-Dxp

```
config
- service
- ies service-id [customer customer-id] [create]
- [no] interface ip-int-name
- [no] sap sap-id [create]
- statistics
- egress
- [no] forwarded-count
- ingress
- counter-mode {in-out-profile-count | forward-drop-count}
- [no] received-count
```

6.5.1.6 IES interface SAP QoS and filter commands for 7210 SAS-D and 7210 SAS-Dxp

```
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] [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
    
```

6.5.1.7 IES interface IPv6 commands (applicable only to access-uplink SAPs on 7210 SAS-D and 7210 SAS-Dxp)

```

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] neighbor ipv6-address mac-address
    
```

6.5.1.8 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
- sap-using [ingress] qos-policy qos-policy-id
- service-using [ies] [customer customer-id]
- id service-id
- all
- arp [ip-address] | [mac ieee-address] | [sap sap-id] | [interface ip-int-name]
- base
- dhcp
- statistics [sap sap-id] | [sdp sdp-id:vc-id] | [interface interface-name]
- summary [interface interface-name | saps]
- interface [ip-address | ip-int-name] [detail | summary]
    
```

6.5.2 Command descriptions

6.5.2.1 IES service configuration commands

6.5.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

```
config>service>ies
```

```
config>service>ies>if
```

```
config>service>ies>if>dhcp
```

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

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 are 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 Internet Enhanced Service (IES) 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 the following are true:

- an IES service is operational and an associated interface is shut down
- the IES service is administratively shutdown and brought back up
- the interface shutdown remains in administrative shutdown state

A service is regarded as operational if one IP Interface is operational.

IES IP Interfaces

When the IP interface is shut down, it enters the administratively and operationally down states. For a SAP bound to the IP interface, no packets are transmitted out the SAP, and all packets received on the SAP are dropped while incrementing the packet discard counter.

description

Syntax

description *description-string*

no description

Context

config>service>ies

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

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

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.

6.5.2.1.2 IES global commands

ies

Syntax

ies *service-id* **customer** *customer-id* [**create**] [**vpn** *vpn-id*]

no ies *service-id*

Context

config>service

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

Description

This command creates or edits an IES 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.

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 to associate the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. After 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.

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.

More than one IP interface may be created within a single IES service ID.

By default, no IES service instances exist until they are explicitly created.

The **no** form of this command deletes the IES service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces defined within the service ID have been shut down and deleted.

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-id

Specifies the VPN ID assigned to the service.

Values 1 to 2147483647

service-name

Syntax

service-name *service-name*

no service-name

Context

config>service>epipe

config>service>ies

config>service>vpls

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command configures a service name that can be used in other configuration commands and show commands that reference the service.

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, up to 64 characters. Service names may not begin with an integer (0-9).

6.5.2.1.3 IES interface IPv6 commands

ipv6

Syntax

[no] ipv6

Context

config>service>ies>if

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

Commands in this context configure IPv6 for an IES interface.



Note:

IES interface IPv6 commands are only supported for access-uplink SAPs on the 7210 SAS-D and 7210 SAS-Dxp.

address

Syntax

address *ipv6-address/prefix-length* [**eui-64**]

no address *ipv6-address/prefix-length*

Context

config>service>ies>if>ipv6

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command assigns an IPv6 address to the IES interface.

Parameters

ipv6-address/prefix-length

Specifies the IPv6 address on the interface.

Values

<i>ipv6-address/prefix:</i>	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
<i>prefix-length:</i>	1 to 128 (7210 SAS-D) 0 to 64 (7210 SAS-Dxp)	

eui-64

Keyword that specifies that a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from the MAC address on Ethernet interfaces. For interfaces without a MAC address, for example ATM interfaces, the Base MAC address of the chassis is used.

icmp6

Syntax

icmp6

Context

config>service>ies>if>ipv6

Platforms

7210 SAS-D, 7210 SAS-Dxp

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

Platforms

7210 SAS-D, 7210 SAS-Dxp

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 param-problem

Context

config>service>ies>if>ipv6>icmp6

Platforms

7210 SAS-D, 7210 SAS-Dxp

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

Platforms

7210 SAS-D, 7210 SAS-Dxp

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 disables the sending of ICMPv6 "redirect" messages.

Default

100 10

Parameters

number

Specifies the number of version 6 redirects 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 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
```

Platforms

7210 SAS-D, 7210 SAS-Dxp

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 the sending of 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
```

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command specifies that ICMPv6 host and network "unreachables" messages are generated by this interface.

The **no** form of this command disables the sending of ICMPv6 host and network "unreachables" messages.

Default

100 10

Parameters

number

Specifies the number of destination "unreachables" ICMPv6 messages are 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 "unreachables" 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

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures the IPv6 link local address.

Parameters

ipv6-address

Specifies the IPv6 link local address.

preferred

Specifies that the IPv6 address is the preferred IPv6 address for this interface. A preferred address is an address assigned to an interface whose use by upper layer protocols is unrestricted. A preferred addresses may be used as the source (or destination) address of packets sent from (or to) the interface.

neighbor

Syntax

neighbor *ipv6-address mac-address*

no neighbor *ipv6-address*

Context

config>service>ies>if>ipv6

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures IPv6-to-MAC address mapping on the IES interface.

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:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D
prefix-length: 1 to 128 (7210 SAS-D)
1 to 64 (7210 SAS-Dxp)

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.

6.5.2.1.4 IES interface commands

interface

Syntax

interface *ip-int-name*[create]

no interface *ip-int-name*

Context

config>service>ies

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

Description



Note:

On 7210 SAS-K 2F1C2T, an IES interface must always be associated with RVPLS. It cannot be used "standalone".

This command creates a logical IP routing interface for an IES. 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.

The **no** form of this command removes IP the 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 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 will be 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.

address

Syntax

address {*ip-address/mask* | *ip-address netmask*} [**broadcast** {**all-ones** | **host-ones**}]

address [*ip-address mask* | *ip-address netmask*]

no address

Context

config>service>ies>if

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

Description

This command assigns an IP address and IP subnet to an IES IP router interface. Only one IP address can be associated with an IP interface. An IP address must be assigned to each IES IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7210 SAS.

The IP address for the interface can be entered in either the Classless Inter-Domain Routing (CIDR) 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 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 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.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.

mask

Specifies the subnet mask length when the IP prefix is specified in CIDR notation. In the 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. 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 of 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

Keyword to configure the broadcast format.

Values all-ones, host-ones

arp-timeout

Syntax

arp-timeout *seconds*

no arp-timeout

Context

config>service>ies>if

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

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 reverts to the default value.

Default

14400

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 are not aged.

Values 0 to 65535

allow-directed-broadcasts

Syntax

[no] **allow-directed-broadcasts**

Context

config>service>ies>if

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

Description

This command enables the forwarding of directed broadcasts out of the IP interface. A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The **allow-directed-broadcasts** command on an IP interface enables or disables the transmission of packets destined for the subnet broadcast address of the egress IP interface.

When enabled, a frame destined 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

delayed-enable

Syntax

delayed-enable *seconds* [**init-only**]

no delayed-enable

Context

config>service>ies>if

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

Description

This command delays making an 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 (for example, after a reboot). The **delayed-enable** timer can be specified to ensure that the state has time to be synchronized. The optional **init-only** parameter 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

Keyword to specify to use this timer only after a reboot.

urpf-check

Syntax

[no] **urpf-check**

Context

config>service>ies>interface

config>service>ies>interface>ipv6

Description

This command enables the unicast RPF (uRPF) check on this interface.
The **no** form of this command disables the uRPF check on this interface.

Default

no urpf-check

ip-mtu

Syntax

ip-mtu *octets*
no ip-mtu

Context

config>service>ies>if

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures the IP maximum transmit unit (MTU) (packet size) for the interface.

The MTU that is advertised from the IES size is the following:

MINIMUM ((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))

By default, for an 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

loopback

Syntax

[no] **loopback**

Context

```
config>service>ies>if
```

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

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.

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.

static-arp

Syntax

```
static-arp ip-address ieee-address
```

```
no static-arp ip-address [ieee-address]
```

Context

```
config>service>ies>if
```

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T.

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 be configured only 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-address

Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff, where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

vpls

Syntax

vpls *service-name*

Context

```
config>service  
config>service>ies>if
```

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

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 associated with the name and the **allow-ip-int-binding flag** is set, the IP interface will be 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 will be ignored and will not be attached to the IP interface.

If the service name is applied to a VPLS service after the service name is bound to an IP interface and the VPLS service **allow-ip-int-binding** flag is set at the time the name is applied, the VPLS service is automatically resolved to the IP interface if the interface is administratively up or when the interface is placed in the administratively up state.

If the service name is applied to a VPLS service without the **allow-ip-int-binding** flag set, the system does not attempt to resolve the applied service name to an existing IP interface bound to the name. To rectify this condition, the flag must first be set and then the IP interface must enter or reenter the administratively up state.

While the specified service name may be assigned to only one service context in the system, it is possible to bind the same service name to more than one IP interface. If two or more IP interfaces are bound to the same service name, the first IP interface to enter the administratively up state (if currently administratively down) or to reenter the administratively up state (if currently administratively up) when a VPLS service is configured with the name and has the **allow-ip-int-binding** flag set will be attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

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 will attempt to resolve to an **allow-ip-int-binding** enabled VPLS service associated with the name. The specified name is expressed as an ASCII string comprised of up to 32 characters. It does not need to already be associated with a service and the system does not check to ensure that multiple IP interfaces are not bound to the same name.

ingress

Syntax

ingress

Context

config>service>ies>if>vpls

Platforms

7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

Description

The ingress node in this context under the VPLS binding defines the routed IP filter ID optional filter overrides.

v4-routed-override-filter

Syntax

v4-routed-override-filter *ip-filter-id*

no v4-routed-override-filter

Context

config>service>ies>if>vpls>ingress

Platforms

7210 SAS-D, 7210 SAS-Dxp, and 7210 SAS-K 2F1C2T

Description

The v4-routed-override-filter command is used to specify an IP filter ID that is applied to all ingress packets entering the VPLS service. The filter overrides any existing ingress IP filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and when it is not defined or it is removed, the IP routed packets uses the any existing ingress IP filter on the VPLS virtual port.

The **no** form of 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 uses 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 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

6.5.2.1.5 IES interface DHCP commands

dhcp

Syntax

dhcp

Context

config>service>ies>if

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

Commands in this context configure DHCP parameters.

gi-address

Syntax

gi-address *ip-address* [**src-ip-addr**]

no gi-address

Context

config>service>ies>if>dhcp

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures the gateway interface address for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. When the router functions as a DHCP relay, the GI address is needed to distinguish between the different subscriber interfaces and potentially between the group interfaces defined.

By default, the GI address used in the relayed DHCP packet is the primary IP address of a normal IES interface. Specifying the GI address allows the user to choose a secondary address.

Default

no gi-address

Parameters

ip-address

Specifies the host IP address to be used for DHCP relay packets.

src-ip-address

Keyword that specifies that the GI address is to be used as the source IP address for DHCP relay packets.

option

Syntax

[no] option

Context

config>service>ies>if>dhcp

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command enables DHCP Option 82 (Relay Agent Information Option) parameter processing and enters the context for configuring Option 82 suboptions.

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

Default

no option

action

Syntax

action {replace | drop | keep}

no action

Context

```
config>service>ies>if>dhcp>option
```

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures the processing required when the 7210 SAS receives a DHCP request that already has a Relay Agent Information Option (Option 82) field in the packet.

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

Parameters

replace

In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (toward the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop

Keyword to specify that the DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.

keep

Keyword to specify 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 sent on 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 occurs only when the incoming message already has 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
```

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command sends an ASCII-encoded tuple in the **circuit-id** suboption of the DHCP packet.

To send a tuple in the circuit ID, the **action replace** command must be configured in the same context.
If disabled, the **circuit-id** suboption of the DHCP packet is left empty.
The **no** form of this command reverts to the default value.

Default

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*, separated by "|".

ifindex

Specifies that the interface index is used. (The If Index of a router interface can be displayed using the **show router interface detail** command)

sap-id

Specifies that the SAP identifier is used.

vlan-ascii-tuple

Specifies that the format includes VLAN-id and dot1p bits in addition to what is included in *ascii-tuple*. The format is supported on dot1q-encapsulated ports only. Therefore, when the Option 82 bits are stripped, dot1p bits are copied to the Ethernet header of an outgoing packet.

remote-id

Syntax

```
[no] remote-id [mac | string string]
```

Context

```
config>service>ies>if>dhcp>option
```

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command specifies what information goes into the **remote-id** suboption in the DHCP Relay packet.
If disabled, the **remote-id** suboption of the DHCP packet is left empty.
The **no** form of this command reverts to the default value.

Default

no remote-id

Parameters

mac

Keyword that specifies the MAC address of the remote end is encoded in the suboption.

string *string*

Specifies the remote-id.

vendor-specific-option

Syntax

[no] vendor-specific-option

Context

config>service>ies>if>dhcp>option

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures the vendor-specific suboption of the DHCP relay packet.

client-mac-address

Syntax

[no] client-mac-address

Context

config>service>ies>if>dhcp>option>vendor

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command enables the sending of the MAC address in the vendor-specific suboption of the DHCP relay packet.

The **no** form of this command disables the sending of the MAC address in the vendor-specific suboption of the DHCP relay packet.

sap-id

Syntax

[no] sap-id

Context

```
config>service>ies>if>dhcp>option>vendor
```

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command enables the sending of 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>ies>if>dhcp>option>vendor
```

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command enables the sending of 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>ies>if>dhcp>option>vendor
```

Platforms

7210 SAS-D, 7210 SAS-Dxp

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

Specifies a string that 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>ies>if>dhcp>option>vendor

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command specifies whether the system ID is encoded in the vendor-specific suboption of Option 82.

server

Syntax

server *server1* [*server2* ... (up to 8 max)]

Context

config>service>ies>if>dhcp

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command specifies a list of servers where requests will be forwarded. The list of servers can be entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers, the request is forwarded to all the servers in the list.

There can be a maximum of 8 DHCP servers configured.

Default

no server

Parameters

server

Specifies the DHCP server IP address.

trusted

Syntax

[no] trusted

Context

config>service>ies>if>dhcp

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command enables relaying of untrusted packets.

The **no** form of this command disables the relay.

Default

no trusted

6.5.2.1.6 IES interface ICMP commands

icmp

Syntax

icmp

Context

config>service>ies>if

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

Description

This command configures the rate for 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 are 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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

Description

This command configures the rate 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

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T

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

6.5.2.1.7 IES interface SAP commands

sap

Syntax

sap *sap-id* [**create**]

no sap *sap-id*

Context

config>service>ies>if

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command creates a SAP within a service. A SAP is a combination of port and encapsulation parameters that 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 be associated with only a single service. A SAP can be defined only 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 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 **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.

Special Cases

IES

A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Attempts to create a second SAP on an IP interface will fail and generate an error; the original SAP will not be affected.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) 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 or disabled in the **environment>create** context.

accounting-policy

Syntax

accounting-policy *acct-policy-id*

no accounting-policy [*acct-policy-id*]

Context

```
config>service>ies>if>sap
```

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures 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

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>ies>if>sap

Platforms

7210 SAS-D, 7210 SAS-Dxp

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 will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued, the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default

no collect-stats

6.5.2.1.8 IES filter and QoS commands

egress

Syntax

egress

Context

config>service>ies>if>sap

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T.

Description

Commands in this context apply egress policies.

ingress

Syntax

ingress

Context

config>service>ies>if>sap

Platforms

7210 SAS-D, 7210 SAS-Dxp, 7210 SAS-K 2F1C2T.

Description

Commands in this context apply ingress policies.

aggregate-meter-rate

Syntax

aggregate-meter-rate *rate-in-kbps* [**burst** *burst-in-kbits*] [**enable-stats**]

no aggregate-meter-rate

Context

config>service>ies>if>sap>egress

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures the access SAP egress aggregate policer. The rate of the SAP egress aggregate policer must be specified. The user can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the traffic sent out of the SAP and determines the final disposition of the packet, which is either forwarded or dropped.

The user can optionally associate a set of two counters to count total forwarded packets and octets and total dropped packets and octets. When this counter is enabled, the amount of resources required increases by twice the amount of resources taken up when the counter is not used. If the **enable-stats** keyword is specified during the creation of the meter, the counter is allocated by the software, if available. To free up the counter and relinquish its use, use the **no aggregate-meter-rate** command and then recreate the meter using the **aggregate-meter rate** command.

If egress frame-based accounting is used, the SAP egress aggregate meter rate accounts for the Ethernet frame overhead. The system accounts for 12 bytes of IFG and 8 bytes of start delimiter. Frame-based accounting does not affect the count of octets maintained by the counter (if in use).



Note:

- Before enabling this command for a SAP, resources must be allocated to this feature from the egress-internal-tcam resource pool using the **config system resource-profile egress-internal-tcam egress-sap-aggregate-meter** command. See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C 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 in which 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 4000000 | max

Default max

burst-in-kbits

Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 to 16384

Default 512

enable-stats

Specifies whether the counter to count forwarded and dropped packets and octets is allocated. If this keyword is used while configuring the meter, the counter is allocated.

aggregate-meter-rate

Syntax

aggregate-meter-rate *rate-in-kbps* [**burst** *burst-in-kbits*]

no aggregate-meter-rate

Context

config>service>ies>if>sap>ingress

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command configures the SAP ingress aggregate policer. The rate of the SAP ingress aggregate policer must be specified. 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.

The sum of CIR of the individual FCs configured under the SAP cannot exceed the PIR rate configured for the SAP. Although the 7210 SAS software does not block this configuration, it is not recommended.

The following table lists the final disposition of the packet based on the operating rate of the per-FC policer and the per-SAP aggregate policer:

Table 68: 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 ¹²	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, the per-FC policer can be configured only 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.

¹² This configuration is not recommended.

Default

no aggregate-meter-rate

Parameters

rate-in-kbps

Specifies the rate in kilobits per second.

Values 0 to 20000000 | max

Default max

burst-in-kilobits

Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 4 to 2146959

Default 512

filter

Syntax

filter ipv6 *ipv6-filter-id*

filter ip *ip-filter-id*

no filter

Context

config>service>ies>if>sap>egress

config>service>ies>if>sap>ingress

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command associates a filter policy with an ingress or egress SAP. Filter policies control the forwarding and dropping of packets based on the matching criteria.

The **filter** command associates 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 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 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

Keyword indicating the filter policy is an IP filter.

ipv6

Keyword indicating the filter policy is an IPv6 filter.

ip-filter-id

Specifies the ID for the IP filter policy. Allowed values are an integer 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*

no qos

Context

config>service>ies>sap>ingress

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command associates a Quality of Service (QoS) policy with an ingress SAP or IP interface.

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 ingress policies. The **qos** command only allows ingress policies to be associated on SAP ingress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specific type returns an error.

By default, if no specific QoS policy is associated with the SAP for ingress, 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 policy ID to associate with SAP on ingress. The policy ID must already exist.

Values 1 to 65535

meter-override

Syntax

[no] **meter-override**

Context

config>service>ies>if>sap>ingress

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command creates 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 removes any existing meter overrides.

Default

no meter-overrides

meter

Syntax

meter *meter-id* [**create**]

no meter *meter-id*

Context

config>service>ies>if>sap>ingress>meter-override

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command creates a CLI node for overrides to a specific meter created on the SAP through SAP-ingress QoS policies.

The **no** form of this command removes any existing overrides for the specified meter-id.

Parameters

meter-id

Specifies the meter ID and is required when executing the **meter** command in the **meter-overrides** context. The *meter-id* must exist within the SAP-ingress QoS policy applied to the SAP. If the meter is not currently used by any forwarding class or forwarding type mappings, the meter will not actually exist on the SAP. This does not preclude creating an override context for the meter ID.

create

Keyword that 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>ies>if>sap>ingress>meter-override>meter

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command overrides 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

Keyword that 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 the EIR value. See the *7210 SAS-D, Dxp Quality of Service Guide* and *7210 SAS-K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Quality of Service Guide* for a description and relevant notes about meter modes.

cir

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

adaptation-rule

Specifies the criteria to compute the operational CIR and PIR values for this meter, while maintaining a minimum offset.

Values **max** — Keyword that is mutually exclusive with the **min** and **closest** options. When **max** (maximum) is defined, the operational PIR for the meter is equal to or less than the administrative rate specified using the **meter-override** command.

min — Keyword that is mutually exclusive with the **max** and **closest** options. When **min** (minimum) is defined, the operational PIR for the queue is equal to or greater than the administrative rate specified using the **meter-override** command.

closest — Keyword that 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>ies>if>sap>ingress>meter-override>meter

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command overrides the default committed burst size (CBS) for the meter. The **cbs** command 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, the packets are marked as in-profile by the meter to indicate that the traffic is complying with meter configured parameters.

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

Default

32 kbits

Parameters

size

Specifies the value in kbits, bytes, or kilobytes.

Values	For 7210 SAS-D: kbits — 4 to 16384, default bytes — 512 to 2097152, default kbytes — 1 to 2048, default For 7210 SAS-Dxp: kbits — 4 to 2146959, default bytes — 512 to 274810752, default kbytes — 1 to 268369, default
---------------	--

mbs

Syntax

mbs *size* [kbits | bytes | kbytes]

no mbs

Context

config>service>ies>if>sap>ingress>meter-override>meter

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command provides a mechanism to override the default maximum burst size (MBS) for the meter. The **mbs** command specifies the maximum burst size that can be transmitted by the source while still complying with the PIR. If the transmitted burst is lower than the MBS value, the packets are marked as in-profile by the meter to indicate that the traffic is complying with meter configured parameters.

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

Default

32 kbits

Parameters

size

Specifies the value in kbits, bytes, or kilobytes.

Values For 7210 SAS-D:
kbits — 4 to 16384, default
bytes — 512 to 2097152, default
kbytes — 1 to 2048, default
For 7210 SAS-Dxp:
kbits — 4 to 2146959, default
bytes — 512 to 274810752, default
kbytes — 1 to 268369, default

mode

Syntax

mode *mode*

no mode

Context

config>service>ies>if>sap>ingress>meter-override>meter

Platforms

7210 SAS-D, 7210 SAS-Dxp

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 restores the policy defined metering and profiling mode to a meter.

Parameters

mode

Specifies the rate mode of the meter-override.

Values trtcm1, trtcm2, srtcm

rate

Syntax

rate *cir* *cir-rate* [**pir** *pir-rate*]

no rate

Context

config>service>ies>if>sap>ingress>meter-override>meter

Platforms

7210 SAS-D, 7210 SAS-Dxp

Description

This command overrides the SAP-ingress QoS policy configured rate parameters for the specified meter ID.

The **no** form of this command restores 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. See the *7210 SAS-D, Dxp Quality of Service Guide* and *7210 SAS-K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Quality of Service Guide* for more information.

The actual PIR rate is dependent on the queue **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 0 to 4000000, max (7210 SAS-D)
0 to 20000000, max (7210 SAS-Dxp)

Default max

cir-rate

Specifies to override 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 to 4000000, max (7210 SAS-D)
0 to 20000000, max (7210 SAS-Dxp)

Default 0

6.5.2.1.9 IES interface SAP Statistics commands

statistics

Syntax

statistics

Context

config>service>>ies>interface>sap

Platforms

7210 SAS-D

Description

Commands in this context configure the counters associated with SAP ingress and egress.

egress

Syntax

egress

Context

config>service>>ies>interface>sap>statistics

Platforms

7210 SAS-D

Description

Commands in this context configure the egress SAP statistics counter and set the mode of the counter. This counter counts the number of packets forwarded through the SAP.

ingress

Syntax

ingress

Context

config>service>>ies>interface>sap>statistics

Platforms

7210 SAS-D

Description

Commands in this context configure the ingress SAP statistics counter.

By default, SAP ingress counters are associated with a SAP and cannot be disabled.

The IES service supports a counter that counts the total packets or octets received on the SAP.

forwarded-count

Syntax

[no] **forwarded-count**

Context

config>service>ies>if>sap>statistics>egress

Platforms

7210 SAS-D

Description

This command associates a counter with the SAP. The counter counts the number of packets forwarded through the SAP.

A limited number of such counters are available for use with access SAPs and access-uplink SAPs.

Use this command before enabling applicable accounting record collection on the SAP to associate a counter with the SAP.

The **no** form of this command disables the packet count.

counter-mode

Syntax

counter-mode {in-out-profile-count | forward-drop-count}

Context

config>service>ies>interface>sap>statistics>ingress

Platforms

7210 SAS-D

Description

This command sets the mode of ingress counters associated with the SAP to either octets or packets. On IES SAPs, collect stats cannot be enabled so the mode of the counter can be changed without any reference. Changing the mode of the counter results in loss of previously collected counts and resets the counter.

The **no** form of this command is not supported.

Default

in-out-profile-count

Parameters

in-out-profile-count

Keyword to specify 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.

forward-drop-count

Keyword to specify 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.

received-count

Syntax

[no] received-count

Context

config>service>ies>if>sap>statistics>ingress

Platforms

7210 SAS-D

Description

This command associates a counter with a SAP. It counts the number of packets and octets received on the SAP (ingress).

A limited number of such counters are available for use with access-uplink SAPs.

Use this command before enabling applicable accounting record collection on the SAP.

The **no** form of this command disables the counter.

7 Internet Enhanced Service on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

This chapter provides information about Internet Enhanced Services support on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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's network.

IES allows IP interfaces to participate in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber be unique between other provider addressing schemes and potentially the entire Internet. While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be reserved for service IP provisioning, and be administered by a separate, but subordinate address authority.

IP interfaces defined within the context of an IES service must have a SAP associated as the access point to the subscriber network. Multiple IES services are created to segregate subscriber owned IP interfaces.

The following figure shows a visual representation of IES.

Figure 80: 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.

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:

- ICMP Options
- VRRP - for IES services with more than one IP interface (available only in network mode)

In network mode, configuration options found on core IP interfaces not supported on IES IP interfaces are:

- NTP broadcast receipt

7.3 SAPs

7.3.1 IPv6 support for IES IP interfaces

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 the IPv6 Provider Edge Router over MPLS (6PE) functionality. For more information about the 6PE, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*, "IPv6 Provider Edge Router over MPLS (6PE)".

IPv4 and IPv6 route table lookup entries are shared.

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). A limited amount of IPv6 /128 prefixes route lookup entries is supported.

7.3.2 Encapsulations

The following SAP encapsulation is supported on IES services:

- Ethernet null
- Ethernet dot1q
- Ethernet QinQ
- Ethernet QinQ (access-uplink QinQ SAP)

7.3.3 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:

- 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.

7.3.3.1 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's routing tables dynamically and minimize wasted bandwidth.

The availability of the far-end static route is monitored through periodic polling. The polling period is configured. If the poll fails a specified number of sequential polls, the static route is marked as inactive.

An ICMP ping mechanism is used to test the connectivity. If the connectivity check fails and the static route is deactivated, the router will continue to send polls and re-activate any routes that are restored.

7.3.4 QoS policies

When applied to 7210 SAS IES services, service ingress QoS policies only create the unicast queues defined in the policy. The multi-point queues are not created on the service. With IES services, service egress QoS policies function as with other services where the class-based queues are created as defined in the policy.

Note that MAC, IPv4, and IPv6 criteria can be used in the QoS policies for traffic classification in an IES.

7.3.4.1 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 the network side.
- IP DSCP marking of self-generated traffic is assigned by software on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

7.3.5 Filter policies

IP filter policies with IPv4 or IPv6 match entries can be applied to IES services.

7.3.6 VRRP support for IES IP interfaces in network mode

Virtual Router Redundancy Protocol (VRRP) for IPv4 is defined in 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 the use of VRRP, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.

VRRP is supported for IES IPv4 interfaces in network mode only. VRRP is not supported in access-uplink mode. It is also not supported for IPv6 interfaces in network or access-uplink mode.



Note:

Only one VRRP instance for each IP interface is supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

7.4 Configuring an IES service with CLI

This section provides information to configure IES services using the command line interface.

7.4.1 Basic configuration

The most basic IES service configuration has the following entities:

- customer ID (see [Configuring customer 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-50.

```
*A:ALA-50>config>service# info
-----
ies 1000 customer 50 create
    description "to internet"
    interface "to-web" create
        address 10.1.1.1/24
        sap 1/1/10:100.* create
    exit
    exit
    no shutdown
-----
*A:ALA-50>config>service#
```

7.4.2 Common configuration tasks

This section provides a brief overview of the tasks that must be performed to configure IES services and provides the CLI commands:

1. Associate an IES service with a customer ID.
2. Associate customer ID with the service.
3. Assign an IP address.
4. Create an interface.
5. Define SAP parameters on the interface
6. Select nodes and ports.
7. Optional - select filter policies (configured in the config>filter context).
8. Enable service.

7.4.2.1 Configuring IES components

7.4.2.1.1 Configuring an IES service

Use the following syntax to create an IES service:

Example

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#
```

7.4.2.1.2 Configuring IES interface parameters

Example

The following is a sample of IP interface parameters under IES service.

```
*A:K-SASK12>config>service>ies>if# info detail
-----
no description
no enable-ingress-stats
no enable-mac-accounting
no tcp-mss
cpu-protection 254
no address
no mac
arp-timeout 14400
arp-retry-timer 50
```

```
no arp-limit
no allow-directed-broadcasts
icmp
  mask-reply
  redirects 100 10
  unreachable 100 10
  ttl-expired 100 10
exit
dhcp
  shutdown
  no description
  no option
  no server
  no trusted
  no relay-proxy
  no gi-address
  no relay-plain-bootp
exit
no authentication-policy
no ip-mtu
no delayed-enable
no multicast-network-domain
no bfd
no local-dhcp-server
no proxy-arp-policy
no local-proxy-arp
no remote-proxy-arp
no ptp-hw-assist
no qos-route-lookup
load-balancing
  no teid-load-balancing
  no egr-ip-load-balancing
  no spi-load-balancing
exit
no vas-if-type
no shutdown
-----
*A:K-SASK12>config>service>ies>if#
```

7.4.2.1.3 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

The following is a sample 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#
```

7.4.2.1.4 Configuring VRRP

Configuring VRRP parameters on an IES interface is optional. VRRP can be configured in either 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

The following is a sample IES interface VRRP owner configuration output.

```
config>service>ies> info  
#-----  
...  
  interface "vrrpowner"  
    address 10.10.10.23/24  
    vrrp 1 owner  
      backup 10.10.10.23  
      authentication-key "vh48l0V7Hs2H6lrMHg2aMJJnZStHwwy0" hash2  
    exit  
  exit  
...  
#-----  
config>service>ies#
```

7.4.3 Service management tasks

This section describes the service management tasks.

7.4.3.1 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

The following is a sample modified service.

```
*A:ALA-A>config>service>ies# info  
-----  
ies 1000 customer 50 create  
  description "This is a new description"  
  interface "to-web" create  
    address 10.1.1.1/24  
    mac 00:dc:98:1d:00:00
```

```
        sap 1/1/5:0.* create
        exit
    exit
    no shutdown
exit
-----
*A:ALA-A>config>service#
```

7.4.3.2 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
```

7.4.3.3 Disabling an IES service

Use the following syntax to shut down the IES service without deleting the service parameters.

```
config>service> ies service-id
shutdown
```

7.4.3.4 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: Re-enabling an IES service

```
config>service# ies 2000
config>service>ies# no shutdown
config>service>ies# exit
```

7.5 IES services command reference for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

7.5.1 Command hierarchies

7.5.1.1 Global commands

```
config
- service
- ies service-id [customer customer-id] [create] [vpn vpn-id]
- no ies service-id
- description description-string
- no description
- interface
- no interface
- service-name service-name
- no service-name
- [no] shutdown
```

7.5.1.2 Interface commands

```
config
- service
- ies service-id [customer customer-id] [create] [vpn vpn-id]
- [no] interface ip-int-name [create]
- address {ip-address/mask | ip-address netmask} [broadcast {all-ones | host-
ones}]
- no address {ip-address/mask | ip-address netmask}
- arp-timeout seconds
- no arp-timeout
- [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
- icmp
- [no] mask-reply
- [no] redirects [number seconds]
- [no] ttl-expired [number seconds]
- [no] unreachablees [number seconds]
```

```
- ip-mtu octets4
- no ip-mtu
- [no] loopback
- [no] local-proxy-arp
- [no] local-dhcp-server local-server-name
- [no] shutdown
- static-arp ip-address ieee-address
- no static-arp ip-address [ieee-address]
- static-arp ieee-address unnumbered
- no static-arp [ieee-address] unnumbered
- unnumbered ip-int-name | ip-address
- no unnumbered
- [no] vrrp virtual-router-id
```

7.5.1.3 VRRP commands (applicable only for network mode)

```
config
- service
- ies service-id [customer customer-id] [create] [vpn vpn-id]
- no ies service-id
- interface ip-int-name [create]
- no interface ip-int-name
- vrrp virtual-router-id [owner]
- no vrrp virtual-router-id
- authentication-key authentication-key | hash-key [hash | hash2]
- no authentication-key
- [no] backup ip-address
- [no] bfd-enable service-id interface interface-name dst-ip ip-address
- [no] bfd-enable interface interface-name dst-ip ip-address
- init-delay seconds
- no init-delay
- [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
```

7.5.1.4 Routed VPLS commands

```
config
- service
- ies service-id [customer customer-id] [vpn vpn-id]
- interface ip-interface-name [create]
- no interface ip-interface-name
- vpls service-name
- no vpls
- ingress
- v4-routed-override-filter ip-filter-id
```



```
- no v4-routed-override-filter
```

7.5.1.5 IES SAP configuration — QoS and filter commands

```
config
- service
- ies service-id [customer customer-id] [vpn vpn-id] [create]
- [no] interface ip-int-name
- [no] sap sap-id [create]
- egress
- agg-shaper-rate agg-rate
- no agg-shaper-rate
- filter ip ip-filter-id
- no filter [ip ip-filter-id]
- qos policy-id
- no qos
- ingress
- agg-shaper-rate agg-rate
- no agg-shaper-rate
- aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
- no aggregate-meter-rate
- filter ip ip-filter-id
- no filter [ip ip-filter-id]
- qos policy-id
- no qos
```

7.5.1.6 Interface IPv6 commands

```
config
- service
- ies service-id [customer customer-id] [create]
- [no] interface ip-int-name [create]
- ipv6
- no ipv6
- address ipv6-address/prefix-length [eui-64] [preferred]
- no address ipv6-address/prefix-length
- icmp6
- packet-too-big number seconds
- no packet-too-big
- param-problem number seconds
- no param-problem
- redirects number seconds
- no redirects
- time-exceeded number seconds
- no time-exceeded
- unreachable number seconds
- no unreachable
- link-local-address ipv6-address [preferred]
- local-proxy-nd
- no local-proxy-nd
- neighbor ipv6-address mac-address
- no neighbor ipv6-address
- proxy-nd-policy policy-name [policy-name...(up to 5 max)]
- no proxy-nd-policy
```

7.5.1.7 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
  - sap-using [ingress] qos-policy qos-policy-id
  - service-using [ies] [customer customer-id]
  - id service-id
    - all
    - arp [ip-address] | [mac ieee-address] | [sap sap-id] | [interface ip-int-name]
    - base
    - dhcp
      - statistics [sap sap-id] | [sdp sdp-id:vc-id] | [interface interface-name]
      - summary [interface interface-name | saps]
    - interface [ip-address | ip-int-name] [detail | summary]
```

7.5.2 Command descriptions

7.5.2.1 IES service configuration commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

7.5.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

config>service>ies
config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.



Note:

- See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide* for information about how to allocate addresses toward IP subnets using the **configure system resource-profile router max-ip-subnets** CLI command.
- Before using IPv6, resources for IPv6 routes must be allocated. See the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Basic System Configuration Guide* for information about how to use the **configure system resource-profile router max-ipv6-routes** CLI command.

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 the following are true:

- an IES service is operational and an associated interface is shut down
- the IES service is administratively shutdown and brought back up
- the interface shutdown remains in administrative shutdown state

A service is regarded as operational if one IP Interface is operational.

IES IP Interfaces

When the IP interface is shut down, it enters the administratively and operationally down states. For a SAP bound to the IP interface, no packets are transmitted out the SAP and all packets received on the SAP are dropped while incrementing the packet discard counter.

description

Syntax

description *description-string*

no description

Context

config>service>ies

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

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.

7.5.2.1.2 IES global commands

ies

Syntax

ies *service-id* **customer** *customer-id* [**create**] [**vpn** *vpn-id*]

no **ies** *service-id*

Context

config>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command creates or edits an IES service instance.

The **ies** command is used to create or maintain an IES. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

IP interfaces defined within the context of an IES service ID must have a SAP created.

When a service is created, the **customer** keyword and *customer-id* must be specified to associate the service with a customer. The *customer-id* must already exist having been 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 results in an error.

More than one IP interface may be created within a single IES service ID.

By default, no IES service instances exist until they are explicitly created.

The **no** form of this command deletes the IES service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces defined within the service ID have been shut down and deleted.

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.

Values 1 to 2147483647

service-name

Syntax

service-name *service-name*

no service-name

Context

config>service>ies

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a service name that can be used in other configuration commands and show commands that reference the service. 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, up to 64 characters, to identify the service. Service names may not begin with an integer (0-9).

7.5.2.1.3 IES interface commands

interface

Syntax

```
interface ip-int-name [create]
no interface ip-int-name
```

Context

```
config>service>ies
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command creates a logical IP routing interface for an IES. When created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.

The **interface** command is used to create and maintain IP routing interfaces within IES service IDs. The **interface** command can be executed in the context of an IES service ID. The IP interface created is associated with the service core network routing instance and default routing.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config service ies interface** (that is, the network core router instance). Interface names must not be in the dotted-decimal notation of an IP address. For example, the name "1.1.1.1" is not allowed, but "int-1.1.1.1" is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All IES IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The **no** form of this command removes 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 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 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 will be changed to maintain that IP interface. If *ip-int-name* already exists within another service ID, an error will occur and context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

address

Syntax

address {*ip-address/mask* | *ip-address netmask*} [**broadcast** {**all-ones** | **host-ones**}]

address {*ip-address/mask* | *ip-address netmask*}

no address

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command assigns an IP address and IP subnet, to an IES IP router interface. Only one IP address can be associated with an IP interface. An IP address must be assigned to each IES IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the 7210 SAS.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted-decimal notation. The show commands display CIDR notation and is stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

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 and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an administratively up state and an address is assigned, it becomes operationally up.

Parameters

ip-address

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.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.

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. 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 will be used in a logical "AND" function to derive the local subnet of the IP address. Allowed values are dotted-decimal addresses in the range 128.0.0.0 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

Specifies the broadcast format.

Values all-ones, host-ones

arp-timeout

Syntax

arp-timeout *seconds*

no arp-timeout

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

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

Default

14400 seconds

Parameters

seconds

Specifies the minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 to 65535

delayed-enable

Syntax

delayed-enable seconds [init-only]

no delayed-enable

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 (for example, after a reboot). To ensure that the state has time to be synchronized, the delayed-enable timer can be specified. The optional **init-only** parameter 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

Keyword to specify to use this timer only after a reboot.

loopback

Syntax

[no] loopback

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

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.

static-arp

Syntax

static-arp *ip-address* *ieee-address*

static-arp *ieee-address* **unnumbered**

no static-arp *ip-address* [*ieee-address*]

no static-arp [*ieee-address*] **unnumbered**

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a static address resolution protocol (ARP) entry associating an IP address or an unnumbered 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 or unnumbered 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-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

Keyword that specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. When this command is configured, it overrides any dynamic ARP.

unnumbered

Syntax

unnumbered *ip-int-name* | *ip-address*

no unnumbered

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures an IP interface as an unnumbered interface and specifies the IP address to be used for the interface.

To conserve IP addresses, unnumbered interfaces can be configured. The address used when generating packets on this interface is the configured *ip-address* parameter.

An error message is generated when an unnumbered interface is configured and an IP address already exists on this interface.

The **no** form of this command removes the IP address from the interface, effectively removing the unnumbered property. The interface must be shut down before the **no unnumbered** command is issued to delete the IP address from the interface or an error message is generated.

Default

no unnumbered

Parameters

ip-int-name | *ip-address*

Specifies the IP interface name or IP address with which to associate the unnumbered IP interface in dotted-decimal notation. The configured IP address must exist on this node. Nokia recommends using the system IP address as it is not associated with a particular interface and is therefore always reachable. The system IP address is the default if *ip-int-name* or *ip-address* is not configured.

vpls

Syntax

vpls *service-name*

Context

```
config>service  
config>service>ies>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command, within the IP interface context, binds 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 will be ignored and will not be attached to the IP interface.

If the service name is applied to a VPLS service after the service name is bound to an IP interface and the VPLS service **allow-ip-int-binding** flag is set at the time the name is applied, the VPLS service is automatically resolved to the IP interface if the interface is administratively up or when the interface is placed in the administratively up state.

If the service name is applied to a VPLS service without the **allow-ip-int-binding** flag set, the system does not attempt to resolve the applied service name to an existing IP interface bound to the name. To rectify this condition, the flag must first be set and then the IP interface must enter or reenter the administratively up state.

While the specified service name may be assigned to only one service context in the system, it is possible to bind the same service name to more than one IP interface. If two or more IP interfaces are bound to the same service name, the first IP interface to enter the administratively up state (if currently administratively down) or to reenter the administratively up state (if currently administratively up) when a VPLS service is configured with the name and has the **allow-ip-int-binding** flag set will be attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

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

Mandatory parameter that specifies the service name that the system will attempt to resolve to an **allow-ip-int-binding** enabled VPLS service associated with the name. The specified name is expressed as an ASCII string composed 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

The `ingress` node in this context under the VPLS binding defines the routed IP filter ID optional filter overrides.

v4-routed-override-filter

Syntax

v4-routed-override-filter *ip-filter-id*

no v4-routed-override-filter

Context

```
config>service>ies>if>vpls>ingress
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures 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 SAP 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 any existing ingress IP filter on the VPLS virtual port.

The **no** form of this command removes 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 a previously created IP filter policy in the **configure>filter>ip-filter** context.

Values 1 to 65535

7.5.2.1.4 IES interface ICMP commands

```
icmp
```

Syntax

```
icmp
```

Context

```
config>service>ies>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables responses to ICMP mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the rate for 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 are 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, the 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the rate 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 are issued can be controlled with the optional *number* and *time* parameters by indicating the maximum number of destination unreachable messages that can be issued on the interface for a specific time interval.

By default, the 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

Specify 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

7.5.2.1.5 IES SAP commands

```
sap
```

Syntax

```
sap sap-id [create]
```

```
no sap sap-id
```

Context

```
config>service>ies>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command creates a SAP within a service. A SAP is a combination of port and encapsulation parameters that identify the service access point on the interface and within the 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 be associated with only a single service. A SAP can only be defined on a port that has been configured as an access port and access uplink port using the **configure port port number ethernet mode access uplink** command.

If a port is shut down, all SAPs on that port become operationally down. When a service is shut down, SAPs for the service are not displayed as operationally down although all traffic traversing the service will be discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP are also deleted.

Special Cases

IES

A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Attempts to create a second SAP on an IP interface will fail and generate an error; the original SAP will not be affected.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) 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 or disabled in the **environment>create** context.

7.5.2.1.6 IES filter commands

filter

Syntax

filter ip *ip-filter-id*

no filter

Context

config>service>ies>if>sap>egress

config>service>ies>if>sap>ingress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command associates a filter policy with an ingress or egress SAP. Filter policies control the forwarding and dropping of packets based on the matching criteria.

The **filter** command associates a filter policy that has 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 is returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is that 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

Keyword indicating the filter policy is an IP filter.

ip-filter-id

Specifies the ID for the IP filter policy. Allowed values are an integer that corresponds to a previously created IP filter policy in the **configure>filter>ip-filter** context.

Values 1 to 65535

egress

Syntax

egress

Context

config>service>ies>if>sap

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context apply egress policies.

ingress

Syntax

ingress

Context

config>service>ies>if>sap

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context apply ingress policies

7.5.2.1.7 DHCP server commands

dhcp

Syntax

dhcp

Context

config>service>ies>interface

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure DHCP parameters.

option

Syntax

[no] option

Context

config>service>ies>if>dhcp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enables the context for configuring Option 82 suboptions.

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

Default

no option

action

Syntax

action [dhcp-action] {replace | drop | keep}

no action

Context

config>service>ies>if>dhcp>option

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the Relay Agent Information Option (Option 82) processing.

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

Default

the default is to keep the existing information intact

Parameters

dhcp-action

Specifies the DHCP option action.

replace

Keyword to specify 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

Keyword to specify that the DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.

keep

Keyword to specify that 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 sent on 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 occurs only when the incoming message already has 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.

gi-address

Syntax

gi-address *ip-address* [**src-ip-addr**]

no gi-address

Context

config>service>ies>if>dhcp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 the different subscriber interfaces and potentially between the group interfaces defined.

By default, the GI address used in the relayed DHCP packet is the primary IP address of a normal IES interface. Specifying the GI address allows the user to choose a secondary address.

Default

no gi-address

Parameters

ip-address

Specifies the host IP address to be used for DHCP relay packets.

src-ip-addr

Keyword that specifies the source IP address to be used for DHCP relay packets.

circuit-id

Syntax

circuit-id [*ascii-tuple* | *ifindex* | **sap-id** | **vlan-ascii-tuple**]

no circuit-id

Context

config>service>ies>if>dhcp>option

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the router to send an ASCII-encoded tuple in the **circuit-id** suboption of the DHCP packet.

To send a tuple in the circuit ID, the **action replace** command must be configured in the same context.

If disabled, the **circuit-id** suboption of the DHCP packet will be left empty.

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

Default

circuit-id

Parameters

ascii-tuple

Specifies that the ASCII-encoded concatenated tuple will be used, which consists of the *access-node-identifier*, *service-id*, and *interface-name*, separated by "|".

ifindex

Specifies that the interface index will be used. The If Index of a router interface can be displayed using the **show router interface detail** command.

sap-id

Specifies that the SAP identifier will be used.

vlan-ascii-tuple

Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in *ascii-tuple* already. The format is supported on dot1q-encapsulated ports only. Therefore, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

remote-id

Syntax

```
[no] remote-id [mac | string string]
```

Context

```
config>service>ies>if>dhcp>option
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies what information goes into the remote ID suboption in the DHCP relay packet.

If disabled, the **remote-id** suboption of the DHCP packet will be left empty.

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

Default

```
no remote-id
```

Parameters

mac

Keyword that specifies the MAC address of the remote end is encoded in the suboption.

string *string*

Specifies the remote ID.

vendor-specific-option

Syntax

[no] vendor-specific-option

Context

config>service>ies>if>dhcp>option

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the vendor-specific suboption of the DHCP relay packet.

client-mac-address

Syntax

[no] client-mac-address

Context

config>service>ies>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the sending of the MAC address in the vendor-specific suboption of the DHCP relay packet.

The **no** form of this command disables the sending of the MAC address in the vendor-specific suboption of the DHCP relay packet.

pool-name

Syntax

[no] pool-name

Context

config>service>ies>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command sends the pool name in the vendor-specific suboption of the DHCP relay packet.

The **no** form of this command disables the sending.

sap-id

Syntax

[no] sap-id

Context

config>service>ies>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the sending of 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>ies>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the sending of 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>ies>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Specifies a string that 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>ies>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether the system ID is encoded in the vendor-specific suboption of Option 82.

relay-plain-bootp

Syntax

relay-plain-bootp

no relay-plain-bootp

Context

```
config>service>ies>if>dhcp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies a list of servers where requests will be forwarded. The list of servers can be entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers, the request is forwarded to all of the servers in the list.

A maximum of 8 DHCP servers can be configured.

Default

no server

Parameters

server

Specifies the DHCP server IP address.

trusted

Syntax

```
[no] trusted
```

Context

```
config>service>ies>if>dhcp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables relaying of untrusted packets.

The **no** form of this command disables the relay.

Default

not enabled

qos

Syntax

qos *policy-id*

no qos

Context

config>service>ies>if>sap>egress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command associates a Quality of Service (QoS) policy with an ingress SAP or IP interface.

QoS egress policies are important for the enforcement of SLA agreements. The policy ID must be defined before associating the policy with a SAP. If the *policy-id* does not exist, an error is returned.

The **qos** command is used to associate ingress policies. The **qos** command only allows egress policies to be associated on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one egress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specific type returns an error.

By default, if no specific QoS policy is associated with the SAP for 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 value.

Parameters

policy-id

Specifies the egress policy ID to associate with SAP on egress. The policy ID must already exist.

Values 1 to 65535

agg-shaper-rate

Syntax

agg-shaper-rate *agg-rate*

no agg-shaper-rate

Context

config>service>ies>if>sap>ingress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues. Specify the CIR rate and the PIR rate. Users must not oversubscribe the total bandwidth available for use by ingress queues.

The **no** form of this command disables the use of the SAP aggregate rate shaper. That is, the SAP can use up the maximum bandwidth available.

Default

no agg-shaper-rate

Parameters

agg-rate

Specifies the rate in kilobits per second.

Values 50 to 3000000 | max (7210 SAS-K 2F6C4T) 64 to 20000000 | max (7210 SAS-K 3SFP+ 8C)

Default max

agg-shaper-rate

Syntax

agg-shaper-rate *agg-rate*

no agg-shaper-rate

Context

config>service>ies>if>sap>egress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and the BUM traffic across all the FCs of the SAP configured to use ingress queues. Specify the CIR rate and the PIR rate. Users must not oversubscribe the total bandwidth available for use by ingress queues.

The **no** form of this command disables the use of SAP aggregate rate shaper. That is, the SAP can use up the maximum bandwidth available.

Default

no agg-shaper-rate

Parameters

agg-rate

Specifies the rate in kilobits per second.

Values 50 to 1000000 | max (7210 SAS-K 2F6C4T)
64 to 10000000 | max (7210 SAS-K 3SFP+ 8C)

Default max

aggregate-meter-rate

Syntax

aggregate-meter-rate *rate-in-kbps* [**burst** *burst-in-kbits*]

no aggregate-meter-rate

Context

config>service>ies>if>sap>ingress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the access SAP ingress aggregate policer. The rate of the SAP ingress aggregate policer must be specified. Users can optionally specify the burst size for the SAP aggregate policer. The aggregate policer monitors the ingress traffic on different FCs using policers to rate-limit the flow and determines the final disposition of the packet. The packet is either forwarded to an identified profile or dropped.



Note:

- The sum of CIR of the individual FCs configured under the SAP cannot exceed the PIR rate configured for the SAP. Although the 7210 SAS software does not block this configuration, it is not recommended.
- The queued traffic flows are not limited by the aggregate meter. That is, only metered flows can use the aggregate meter. Queued flows can use only aggregate shaper.

The following table lists the final disposition of the packet based on the operating rate of the per-FC policer and the per-SAP aggregate policer.

Table 69: 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	Final packet color
Within CIR	Green	Within PIR	Green or In-profile
Within CIR	Green	Above PIR	Red and Dropped
Above CIR, Within PIR	Yellow	Within PIR	Yellow or Out-of-Profile
Above CIR, Within PIR	Yellow	Above PIR	Red or Dropped
Above PIR	Red	Within PIR	Red or Dropped
Above PIR	Red	Above PIR	Red or Dropped

The SAP ingress meter counters increment the packet or octet counts based on the final disposition of the packet.

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.

Values 50 to 3000000 | max (7210 SAS-K 2F1C2T and 7210 SAS-K 2F6C4T)
 64 to 20000000 | max (7210 SAS-K 3SFP+ 8C)

Default max

burst-in-kbits

Specifies the burst size for the policer in kilobits. The burst size cannot be configured without configuring the rate.

Values 1 to 16384 | default

Default 512

qos

Syntax

qos *policy-id*

no qos

Context

config>service>ies>if>sap>ingress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command associates a QoS policy with an ingress SAP or IP interface.

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 ingress policies. The **qos** command only allows ingress policies to be associated on SAP ingress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress QoS policy can be associated with a SAP at one time. Attempts to associate a second QoS policy of a specific type returns an error.

By default, if no specific QoS policy is associated with the SAP for ingress, 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 value.

Parameters

policy-id

Specifies the ingress policy ID to associate with SAP on ingress. The policy ID must already exist.

Values 1 to 65535

local-proxy-arp

Syntax

[no] local-proxy-arp

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and therefore will become the forwarding point for all traffic between hosts in that subnet. When the **local-proxy-arp** command is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

Default

no local-proxy-arp

local-dhcp-server

Syntax

local-dhcp-server *server-name* [create]

no local-dhcp-server *server-name*

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a local DHCP server. A local DHCP server can serve multiple interfaces but is limited to the routing context in which it was created.

Parameters

server-name

Specifies the name of local DHCP server.

create

Keyword used to create the server name. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

loopback

Syntax

[no] loopback

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies that the interface is a loopback interface that has no associated physical interface. As a result, the associated IES interface cannot be bound to a SAP.

You can configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap** command. The **loopback** flag cannot be set on an interface where a SAP is already defined, and a SAP cannot be defined on a loopback interface.

ip-mtu

Syntax

ip-mtu *octets*

no ip-mtu

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IP maximum transmit unit (MTU) (packet) for this interface.

Because this connects a Layer 2 to a Layer 3 service, this parameter can be adjusted under the IES interface.

The MTU that is advertised from the IES size is the following:

MINIMUM ((SdpOperPathMtu - EtherHeaderSize), (Configured ip-mtu))

By default (for Ethernet network interface) if **no ip-mtu** is configured, it is $(1568 - 14) = 1554$.

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

Default

no ip-mtu

7.5.2.1.8 IES interface VRRP commands

```
vrrp
```

Syntax

vrrp *virtual-router-id* [**owner**]

no vrrp *virtual-router-id*

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 an owner. The **vrrp** *virtual-router-id* context defines 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

owner

Keyword that defines the virtual router instance as an owner.

```
authentication-key
```

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

```
config>service>ies>if>vrrp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command 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
- shut down the virtual router instance on all backups
- execute the authentication-key command on the master to change the password key
- execute the authentication-key command and no shutdown command on each backup key

By default, the authentication data field contains the value 0 in all 16 octets.

The **no** form of this command restores the default null string to the value of key.

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 *authentication-key* parameter is expressed as a string consisting of up to eight alphanumeric characters. Spaces must be contained in quotation marks (" "). The quotation marks are not considered part of the string.

The string is case-sensitive and is left-justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet (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

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 option is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash

Keyword to specify 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

Keyword to specify 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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures virtual router IP addresses for the interface.

Default

no backup

Parameters

ip-address

Specifies the destination IP address for backup.

Values a.b.c.d

bfd-enable

Syntax

bfd-enable *service-id* **interface** *interface-name* **dst-ip** *ip-address*

no bfd-enable *service-id* **if** *if-name* **dst-ip** *ip-address*

bfd-enable interface *interface-name* **dst-ip** *ip-address*

no bfd-enable interface *interface-name* **dst-ip** *ip-address*

Context

config>service>ies>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 using the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated BGP protocol peering.

Default

no bfd-enable

init-delay

Syntax

init-delay *seconds*

no init-delay

Context

config>service>ies>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 will be silently discarded.

This 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 reverts to the default value.

Parameters

seconds

Specifies the number of seconds that will transpire before the advertisement timer expires.

Values 1 to 255

Default 1

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the non-owner master to reply to ICMP echo requests directed to the virtual router instance IP addresses. The ping request can be received on any routed interface.

Ping must not have been disabled at the management security level (either on the parental IP interface or based on the ping source host address). When ping reply is not enabled, ICMP echo requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP echo requests regardless of the setting of ping reply configuration.

The **ping-reply** command is available only 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 reverts to 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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command provides the ability to override 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 available only in the non-owner **vrrp** *virtual-router-id* nodal context. The owner may not be preempted because the priority of non-owners can never be higher than the owner. The owner always preempts all other virtual routers when it is available.

Non-owner 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 instance in-use priority.

A master non-owner virtual router allows itself to be preempted only 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 revert to the default mode.

Default

preempt

priority

Syntax

priority *priority*

no priority

Context

config>service>ies>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a specific priority value for 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 available only 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 reverts to the default value of 100.

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>ies>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 parent IP interface or based on the SSH source host address). Proper 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 available only 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 reverts to 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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 the virtual router MAC address. The standby router should forward traffic sent to the real MAC address of the standby router.

Default

no standby-forwarding

telnet-reply

Syntax

[no] telnet-reply

Context

config>service>ies>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 parent IP interface or based on the Telnet source host address). Proper 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 available only 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 revert to 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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables a non-owner master to reply to traceroute requests directed to the virtual router instance IP addresses.

This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

A non-owner backup virtual router never responds to such traceroute requests regardless of the traceroute reply status.

Default

no traceroute-reply

7.5.2.1.9 IES interface IPv6 commands

ipv6

Syntax

ipv6

no ipv6

Context

config>service>ies>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables IPv6 for an IES interface.

The **no** form of this command disables IPv6.

Default

no ipv6

address

Syntax

ipv6 *ipv6-address/prefix-length* [**eui-64**] [**preferred**]

no ipv6 *ipv6-address/prefix-length*

Context

```
config>service>ies>if>ipv6
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command assigns an address to the IPv6 interface.

The **no** form of this command deletes the specified IPv6 address.

Parameters

ipv6-address/prefix-length

Specifies the interface IP address.

Values

ipv6-address: 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

prefix-length: 1 to
128

eui-64

Keyword to specify that a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from the MAC address on Ethernet interfaces.

preferred

Keyword to specify that the IPv6 address is the preferred IPv6 address for this interface. A preferred address is an address assigned to an interface whose use by upper layer protocols is unrestricted. A preferred address maybe used as the source (or destination) address of packets sent from (or to) the interface.

icmp6

Syntax

```
icmp
```

Context

```
config>service>ies>if>ipv6
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures ICMPv6 parameters for the interface.

packet-too-big

Syntax

packet-too-big *number seconds*

no packet-too-big

Context

config>service>ies>if>ipv6>icmp6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether, and how often, ICMPv6 "packet-too-big" 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

no packet too big

Parameters

number

Specifies the number of ICMPv6 "packet-too-big" 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 ICMPv6 "packet-too-big" messages sent.

Values 1 to 60

Default 10

param-problem

Syntax

param-problem *number seconds*

no param-problem

Context

config>service>ies>if>ipv6>icmp6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether, and how often, ICMPv6 "parameter-problem" messages should be sent. When enabled, ICMPv6 "parameter-problem" messages are generated by this interface.

The **no** form of this command disables the sending of ICMPv6 "parameter-problem" messages.

Default

no param-problem

Parameters

number

Specifies the number of ICMPv6 "parameter-problem" 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 ICMPv6 "parameter-problem" messages sent.

Values 1 to 60

Default 10

redirects

Syntax

redirects *number seconds*

no redirects

Context

```
config>service>ies>if>ipv6>icmp6
```

Description

Platforms Supported: 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 should be sent.

The **no** form of this command disables the sending of ICMPv6 "redirect" messages.

Default

```
no redirects
```

Parameters

number

Specifies the number of ICMPv6 "redirect" 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 ICMPv6 "redirect" messages sent

Values 1 to 60

Default 10

time-exceeded

Syntax

```
time-exceeded number seconds
```

```
no time-exceeded
```

Context

```
config>service>ies>if>ipv6>icmp6
```

Description

Platforms Supported: 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

This command configures the rate for ICMPv6 "time-exceeded" messages.

The **no** form of this command disables the sending of ICMPv6 "time-exceeded" messages.

Default

no time-exceeded

Parameters

number

Specifies the number of ICMPv6 "time-exceeded" 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 ICMPv6 "time-exceeded" messages sent.

Values 1 to 60

Default 10

unreachables

Syntax

unreachables *number seconds*

no unreachables

Context

config>service>ies>if>ipv6>icmp6

Description

Platforms Supported: 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

This command enables and configures the rate for ICMPv6 host and network destination "unreachables" messages issued on the router interface.

The **no** form of this command disables the generation of ICMPv6 destination "unreachables" messages on the router interface.

Default

no unreachables

Parameters

number

Specifies the number of ICMPv6 "unreachables" 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 ICMPv6 "unreachables" messages sent.

Values 1 to 60

Default 10

link-local-address

Syntax

ipv6 *ipv6-address* [**preferred**]
no ipv6 *ipv6-address/prefix-length*

Context

config>service>ies>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command assigns the IPv6 link local address to the interface.

Parameters

ipv6-address

Specifies the IPv6 link local address.

Values *ipv6-address*: 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

preferred

Keyword to specify that the IPv6 link local address is the preferred IPv6 link local address for this interface. A preferred address is an address assigned to an interface whose use by upper layer protocols is unrestricted. A preferred addresses maybe used as the source (or destination) address of packets sent from (or to) the interface.

local-proxy-nd

Syntax

local-proxy-nd *number seconds*

no local-proxy-nd

Context

config>service>ies>if>ipv6

Description

Platforms Supported: 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

This command enables local proxy neighbor discovery on the interface.

The **no** form of this command disables local proxy neighbor discovery.

Default

no local-proxy-nd

neighbor

Syntax

neighbor *ipv6-address mac-address*

no neighbor *ipv6-address*

Context

config>service>ies>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures IPv6-to-MAC address mapping on the interface.

The **no** form of this command deletes IPv6-to-MAC address mapping for the specified IPv6 address.

Parameters

ipv6-address

Specifies the IPv6 address.

Values

<i>ipv6-</i>	x:x:x:x:x:x:x (eight 16-bit pieces)
<i>address:</i>	x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H

d: [0 to 255]D

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 unicast MAC addresses and non-IEEE reserved MAC addresses.

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command applies a proxy neighbor discovery policy for the interface.

The **no** form of this command disables the proxy neighbor discovery policy application.

Default

no proxy-nd-policy

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 (such as #, \$, spaces), the entire string must be enclosed within double quotes. The specified policy names must already be defined.

7.5.2.2 IES 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

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 that is the anchor point for an ingress and egress virtual scheduler hierarchy.

Output

The following output is an example of customer information, and [Table 70: 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 70: 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 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

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 sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

ingress

Keyword to specify matching an ingress policy.

egress

Keyword to specify 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

Keyword to specify 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 service SAP information, and [Table 71: 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.   Ing.   Egr.   Adm   Opr
                    QoS      Fltr   Fltr
-----
1/1/1                 6          1     none  none   Up    Down
1/1/2                 700        1     none  none   Up    Down
-----
Number of SAPs : 2
=====
```

*A: DUT-B#

Table 71: 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

Description

This command displays the services matching specific usage properties. If no optional parameters are specified, all services defined on the system are displayed.

Parameters

ies

Display 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 72: 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 72: 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**}

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 to identify the service in the service domain.

all

Displays detailed information about the service.

arp

Displays ARP entries for the service.

base

Displays basic service information.

sap

Displays SAPs associated to the service.

all

Syntax

all

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

This command displays detailed information for all aspects of the service.

Output

[Table 73: Output fields: ID all](#) describes the all service ID command output fields.

Sample output

Table 73: Output fields: 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.
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.

Label	Description
Egress Filter	The ID of the egress filter policy.
Far End	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Last Changed	The date and time of the most recent change to this customer.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	Specifies the operating status of the service.
Oper State	The current status of the service.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP Echo Request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far-end field.
Number of SDPs	The total number SDPs applied to this service ID.
Service Access Points	
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.

Label	Description
Admin MTU	The largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The SAP ingress QoS policy ID.
Egress qos-policy	The SAP egress QoS policy ID.
Ingress Filter-Id	The SAP ingress filter policy ID.
Egress Filter-Id	The SAP egress filter policy ID.
Multi Svc Site	Indicates the multi-service site that the SAP is a member.
Ingress sched-policy	Indicates the ingress QoS scheduler for the SAP.
Egress sched-policy	Indicates the egress QoS scheduler for the SAP.
Acct. Pol	Indicates the accounting policy applied to the SAP.
Collect Stats	Specifies whether accounting statistics are collected on the SAP.
SAP Statistics	
Dropped	The number of packets or octets dropped.
Offered Hi Priority	The number of high priority packets, as determined by the SAP ingress QoS policy.
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy.
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded.
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded.

arp

Syntax

arp [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

Description

Displays the ARP table for the IES instance. The ARP entries for a subscriber interface are displayed uniquely. Each MAC associated with the subscriber interface child group-interfaces are displayed with each subscriber interface ARP entry. They do not reflect actual ARP entries but are displayed along the interfaces ARP entry for easy lookup.

Parameters

ip-address

Specifies the IP address of the interface for which to display matching ARP entries.

Values

ipv4-address:	a.b.c.d (host bits must be 0)
ipv6-address:	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

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.

interface

Specifies matching service ARP entries associated with the IP interface.

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 74: 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*
192.168.1.1    32:67:01:01:00:03  Other    00h00m00s   to7x           1/1/3:10*
192.168.2.2    32:68:01:01:00:02  Dynamic  03h59m58s   to7x           1/1/3:10*
=====
*A:DUT-B#
    
```

Table 74: 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

Description

This command displays basic information about this IES service.

Output

The following output is an example of basic IES service information, and [Table 75: 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 75: Output fields: base

Label	Description
Service Id	Displays the service identifier.
VPN Id	Displays the VPN identifier.
Service Type	Displays the type of service.
Description	Displays generic information about the service.
Customer Id	Displays the customer identifier.
Last Status Change	Displays the date and time of the most recent status change.
Last Mgmt Change	Displays 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 ID (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	Specifies the actual largest service frame size (in octets) that can be transmitted through this port, without requiring the packet to be fragmented.

Label	Description
Admin	The administrative state of the SAP.
Opr	The operating state of the SAP.

interface

Syntax

interface [*ip-address* | *ip-int-name*] [**detail** | **summary**]

Context

show>service>id

Platforms

Supported on all 7210 SAS platforms as described in this document

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)
ipv6-address: 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

ip-int-name

Specifies the IP interface name for which to display information.

Values 32 characters maximum

detail

Displays detailed IP interface information.

Default IP interface summary output.

summary

Displays the summary of IP interface information.

Output

The following output is an example of service interface information, and [Table 76: Output fields: service ID interface](#) describes the output fields.

Sample output

```
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 76: 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.

Label	Description
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 Network service on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

This chapter provides information about the Virtual Private Routed Network (VPRN) service and implementation notes on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

8.1 VPRN service overview

RFC2547bis is an extension to the original RFC 2547, which details a method of distributing routing information and forwarding data to provide a Layer 3 Virtual Private Network (VPN) service to end customers.

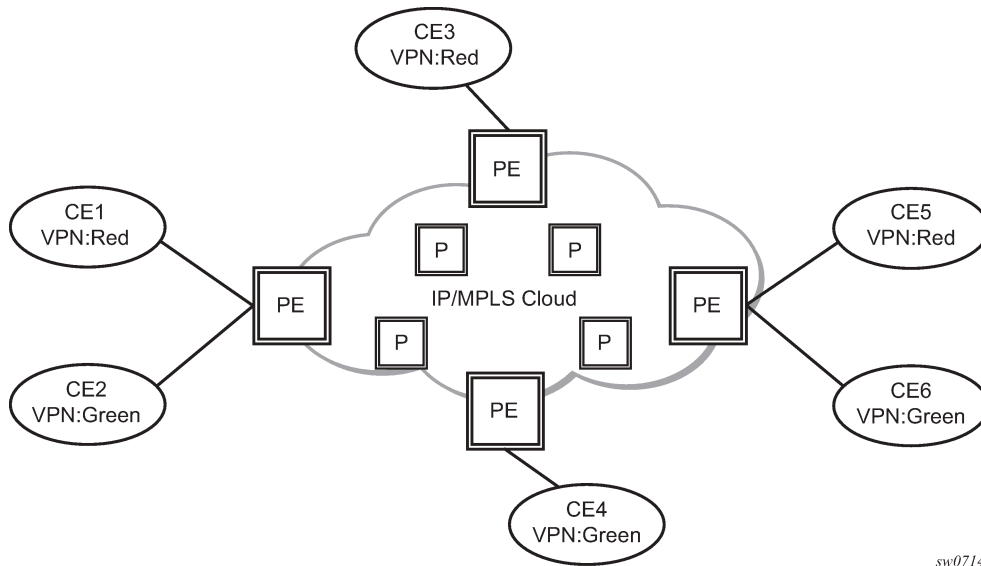
Each Virtual Private Routed Network (VPRN) consists of a set of customer sites connected to one or more PE routers. Each associated PE router maintains a separate IP forwarding table for each VPRN. Additionally, the PE routers exchange the routing information configured or learned from all customer sites via MP-BGP peering. Each route exchanged via the MP-BGP protocol includes a Route Distinguisher (RD), which identifies the VPRN association.

The service provider uses BGP to exchange the routes of a VPN among the PE routers that are attached to that VPN. This is done in a way which ensures that routes from different VPNs remain distinct and separate, even if two VPNs have an overlapping address space. The PE routers distribute routes from other CE routers in that VPN to the CE routers in a VPN. Because the CE routers do not peer with each other there is no overlay visible to the VPN's routing algorithm.

When BGP distributes a VPN route, it also distributes an MPLS label for that route. The label distributed with a VPN route depends on the configured label-mode of the VPRN that is originating the route. On 7210 SAS routers, only one method of label allocation mode is supported: label per VRF. That is, when 7210 allocates a label for a VPN route, it allocates a single label for all the VPN routes belonging to a single VRF. This does not restrict the label distribution method on the remote PE. The label distribution method configured on the remote PE must consider the scale of the PEs participating in the VPRN service.

Before a customer data packet travels across the service provider's backbone, it is encapsulated with the MPLS label that corresponds, in the customer's VPN, to the route which best matches the packet's destination address. The MPLS packet is further encapsulated with either another MPLS label header, so that it gets tunneled across the backbone to the 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. [Figure 81: Virtual Private Routed Network](#) shows a VPRN network diagram example.

Figure 81: Virtual Private Routed Network



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 to distribute VPRN routing information across the service provider's network. BGP was initially designed to distribute IPv4 routing information. Therefore, multi-protocol extensions and the use of a VPN-IPv4 address were created to extend BGP's ability to carry overlapping routing information. A VPN-IPv4 address is a 12-byte value consisting of the 8-byte route distinguisher (RD) and the 4-byte IPv4 IP address prefix. The RD must be unique within the scope of the VPRN. This allows the IP address prefixes within different VRFs to overlap.

The 7210 SAS-K 2F6C4T and 7210 SAS-K 2F6C4T provide 6PE support in the base routing instance with BGP 3017 labeled routes to interconnect IPv6 networks over an IP/MPLS infrastructure.

8.1.3 Route distinguisher

The following figure shows the route distinguisher (RD), an 8-byte value consisting of 2 major fields, the Type field and value field. The type field determines how the value field should be interpreted. The 7210 SAS implementation supports the three (3) type values as defined in the Internet draft.

Figure 82: Route distinguisher



sw0713

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 ASNs 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 ASNs is discouraged). The Assigned field contains a number assigned by the service provider.

8.1.4 Route reflector

Per RFC2547bis the use of Route Reflectors is supported in the service provider core. Multiple sets of route reflectors can be used for different types of BGP routes, including VPN-IPv4, and so on. The 7210 SAS-K 2F6C4T or 7210 SAS-K 3SFP+ 8C can only be configured a route reflector client. It cannot be used as a route reflector "server" in the network.

8.1.5 CE to PE route exchange

Routing information between the Customer Edge (CE) and Provider Edge (PE) can be exchanged by the following methods:

- static routes (IPv4 and IPv6)
- E-BGP (IPv4 and IPv6)
- OSPFv2 (IPv4)

Each protocol provides controls to limit the number of routes learned from each CE router.

8.1.5.1 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 specific routes before they are injected into the local VRF.

Route redistribution from the local VRF to CE-to-PE routing protocols is to be controlled via the route policies in each routing protocol instance, in the same manner that is used by the base router instance.

The advertisement or redistribution of routing information from the local VRF to or from the MPBGP instance is specified per VRF and is controlled by VRF route target associations or by VRF route policies.

VPN-IP routes imported into a VPRN, have the protocol type `bgp-vpn` to denote that it is an VPRN route. This can be used within the route policy match criteria.

8.1.5.2 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.

The following figure shows static route connection information. [Figure 84: Multiple hops to IP target](#) shows multiple hop information.

Figure 83: Directly connected to IP target


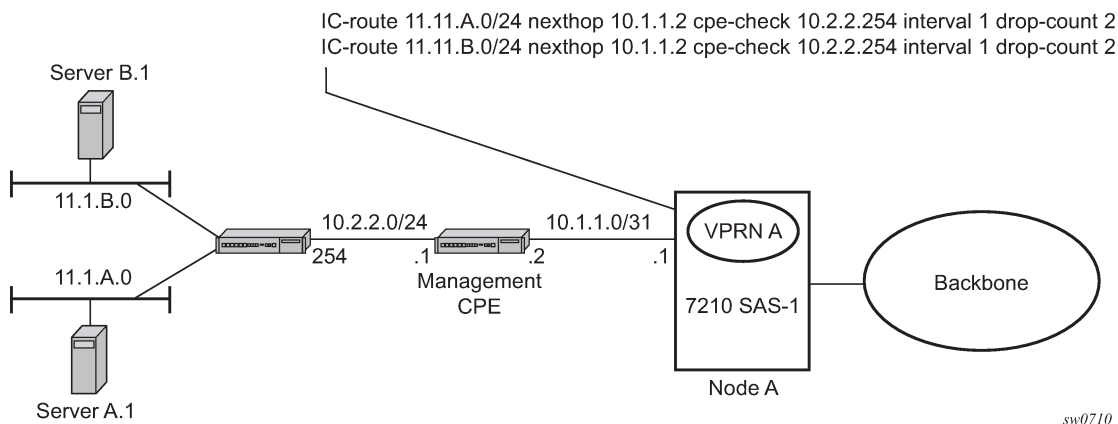


Figure 84: Multiple hops to IP target



The availability of the far-end static route is monitored through periodic polling. The polling period is configured. If the poll fails a specified number of sequential polls, the static route is marked as inactive. Either ICMP ping or unicast ARP mechanism can be used to test the connectivity. ICMP ping is preferred. If the connectivity check fails and the static route is deactivated, the 7210 SAS router will continue to send polls and reactivate any routes that are restored.

8.1.6 Constrained route distribution

This section describes constrained route distribution or RT constraint (RTC).



Note:

This feature is supported only on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

8.1.6.1 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.

8.1.6.2 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.

8.1.6.3 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, for example, in the following contexts:

- **config>service>vpls>bgp**
- **config>service>vprn**



Note:

See the *7210 SAS-K 2F6C4T, K 3SFP+ 8C 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.

8.1.6.4 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.

8.1.6.5 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.

8.2 VPRN features

This section describes various VPRN features and any 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:

- ICMP and ICMP6 options
- VRRP - for VPRN services with more than one IPv4 interface

8.2.2 SAPs

The following sections describe SAP functionality on VPRN services.

8.2.2.1 IPv6 support for VPRN IP interfaces (6VPE) in network mode

VPRN IPv6 (6VPE) access interfaces can 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 need to be allocated for IPv6 addresses. This is done using the CLI command **config> system> resource-profile> router> maxipv6- routes**. This command allocates route entries for /64 IPv6 prefix route lookups. The system does not allocate any IPv6 route entries by default; the user needs to allocate resources before using IPv6. For the command to take effect the node must be rebooted after making the change.

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). A limited amount of IPv6 /128 prefixes route lookup entries is supported.

The number of IP subnets can be configured using the command **configure> system>resource-profile>router>max-ip-subnets**. Default are assigned to this parameter. Users can increase the number of subnets if they plan to more IPv6 addresses per IPv6 interface.

8.2.2.2 Encapsulations

The following SAP encapsulations are supported on the 7210 SAS VPRN service:

- Ethernet null
- Ethernet dot1q
- QinQ
- LAG

8.2.2.3 QoS policies for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

When applied to a VPRN SAP, service ingress and service egress QoS policies create the unicast and multicast queues defined in the policy (as multicast is not supported in VPRN service). Note that both Layer 2 (dot1p only) or Layer-3 criteria can be used in the QoS policies for traffic classification in an VPRN.

8.2.2.4 Filter policies

Ingress and egress IPv4 and IPv6 filter policies can be applied to VPRN SAPs.

8.2.2.5 CPU QoS for VPRN interfaces

Traffic bound to the CPU that is received on VPRN access interfaces is policed/rate-limited and is added to the 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 SAPs. The policer CIR/PIR values are set to appropriate values based on feature scaling. CIR/PIR values are not user configurable. The policers are shared by the IP application/set of IP applications traffic received on all access SAPs across all services. Queues for CPU bound IP traffic from all VPRN SAPs are software allocated. The queues are either shared by a set of IP applications or allocated to a single IP application across all services and routing instances. The queues are shaped to an appropriate rate based on feature scaling. The shaper rate is not user configurable. The queues and shaper rate are shared by traffic of the IP application/set of IP applications received on all access SAPs. The queue instances and policers used for traffic received on network port IP interfaces are different for traffic received from access port IP interfaces. Also, the network CPU queues are given higher priority than the access CPU queues. This provides better security and mitigates the risk of access traffic affecting the network side.



Note:

- The instance of queues and policers used for traffic received on network port IP interfaces is different from the ones used for traffic received from access port IP interfaces. Additionally, the network CPU queues are prioritized over access CPU queues to provide better security and mitigate the risk of access traffic affecting the network side.
- On the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C, the IP DSCP value for self-generated traffic is user-configurable.
- On the 7210 SAS-D and 7210 SAS-K 2F1C2T, IP DSCP marking of self-generated traffic is not user-configurable and is assigned by the software.

8.2.3 CE to PE routing protocols

The 7210 SAS VPRN supports the following PE to CE routing protocols:

- eBGP (IPv4 and IPv6)
- Static (IPv4 and IPv6)
- OSPFv2 (IPv4)

8.2.3.1 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 of using dynamically created LSPs where the service tunnels are automatically bound (the "auto-bind" feature) and the ability to provide certain VPN services with their own transport tunnels by explicitly binding SDPs if desired. When the auto-bind is used, all services traverse the same LSPs and do not allow alternate tunneling mechanisms or the ability to craft sets of LSP's with bandwidth reservations for specific customers as is available with explicit SDPs for the service.

8.2.3.2 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 that the VRF table is near full and an option to disable additional route learning when full or only generate an event.

8.2.4 Spoke-SDP termination

Spoke-SDP termination into a Layer 3 service is not supported on 7210 SAS platforms. On the 7210 SAS-K 2F6C4T, spoke-SDP termination is supported on a routed VPLS that is associated with a VPRN service. All commands supported in the **config>service>vpls>spoke-sdp** context are supported on a routed VPLS. See [VPLS spoke-SDP commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#) for information about the **spoke-sdp** command.

8.2.4.1 Using OSPF in IP-VPNs



Note:

OSPFv2 as PE-CE routing protocol is supported for IPv4 VPNs. OSPFv3 is not supported as a PE-CE routing protocol.

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.

8.2.5 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.

8.2.5.1 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:**

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C platforms support only option-A and option-C. They do not support option-B. It is described as follows only for the sake of completeness.

The first option, referred to as Option-A ([Figure 85: Inter-AS option-A: VRF-to-VRF model](#)), is considered inherent in any implementation. This method uses a back-to-back connection between separate VPRN instances in each AS. Therefore, 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 85: Inter-AS option-A: VRF-to-VRF model



The second option, referred to as Option-B ([Figure 86: Inter-AS option-B](#)), relies heavily on the AS Boundary Routers (ASBRs) as the interface between the autonomous systems. This approach enhances the scalability of the eBGP VRF-to-VRF solution by eliminating the need for per-VPRN configuration on the 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 86: Inter-AS option-B




This form of inter-AS VPRNs does not require instances of the VPRN to be created on the ASBR, as in option-A, therefore there is less management overhead. This is also the most common form of Inter-AS VPRNs used between different service providers as all routes advertised between autonomous systems can be controlled by route policies on the ASBRs.

The third option, referred to as Option-C (Figure 87: Option C example), allows for a higher scale of VPRNs across AS boundaries but also expands the trust model between ASNs. Therefore, 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. Therefore, 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 87: Option C example



With Option-C, knowledge from the remote AS is distributed throughout the local AS. This distribution allows for higher scalability but also requires all PEs and ASBRs involved in the Inter-AS VPRNs to participate in the exchange of inter-AS routing information. In Option-C, the ASBRs distribute reachability information for remote PE's system IP addresses only. This is done between the ASBRs by exchanging MP-eBGP labeled routes, using RFC 3107, *Carrying Label Information in BGP-4*.



Note:

The 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C do not play the role of option-C ASBR.

Distribution of VPRN routing information is handled by either direct MP-BGP peering between PEs in the different ASNs or more likely by one or more route reflectors in ASN.

8.2.5.2 VRRP support for VPRN IP interfaces in network mode

Virtual Router Redundancy Protocol (VRRP) for IPv4 is defined in 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 the use of VRRP, see the *7210 SAS-D, Dxp, K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Router Configuration Guide*.

VRRP is supported for VRRP IPv4 interfaces. It is not supported for IPv6 interfaces in network or access-uplink mode.



Note:

Only one VRRP instance for each IP interface is supported on the 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

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 customer accounts](#))
- specify interface parameters

Example

The following is a sample configuration output of a VPRN service.

```
*A:K-SASK12>config>service>vprn$ info detail
-----
      shutdown
      no description
      dhcp
      exit
```

```
dhcp6
exit
no router-id
no maximum-routes
snmp
  no access
exit
log
exit
no mc-maximum-routes
no autonomous-system
no route-distinguisher
auto-bind-tunnel
  resolution-filter
    no ldp
    no rsvp
  exit
  resolution disabled
exit
no enable-bgp-vpn-backup
no vrf-target
no static-route-hold-down
no ptp
mvpn
  provider-tunnel
  exit
  no vrf-target
exit
radius-proxy
exit
radius-server
exit
no allow-export-bgp-vpn
twamp-light
exit
network
  ingress
    no filter
    no qos
  exit
exit
no vsd-domain
```

*A:K-SASK12>config>service>vprn\$

8.3.2 Common configuration tasks

This section provides a brief overview of the tasks that must be performed to configure a VPRN service and provides the CLI commands:

1. Associate a VPRN service with a customer ID.
2. Define an autonomous system (optional).
3. Define a route distinguisher (mandatory).
4. Define VRF route-target associations or VRF import/export policies.
5. Create an interface.
6. Define SAP parameters on the interface:
 - Select 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)
7. Define BGP parameters (optional):
 - BGP must be enabled in the **config>router>bgp** context.
 8. Enable the service.

8.3.3 Configuring VPRN components

8.3.3.1 Creating a VPRN service

Use the following syntax to create a VPRN service. A route distinguisher must be defined 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

The following is a sample VPRN service configuration output.

```
*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 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#) for CLI syntax to configure VPRN parameters.

Example

The following is a sample VPRN service with configured parameters.

```
*A:ALA-1>config>service# info  
-----  
...  
vprn 1 customer 1 create  
exit  
no enable-bgp-vpn-backup  
no vrf-target  
no static-route-hold-down
```

```
        no ptp
        mvpn
            provider-tunnel
            exit
            no vrf-target
        exit
exit
...
-----
*A:ALA-1>config>service#
```

8.3.3.3 Configuring router interfaces

See the *7210 SAS OS Router Configuration Guide* for command descriptions and syntax information to configure router interfaces.

Example

The following is a sample router interface configuration output.

```
*A:K-SASK12>config>router# info
-----
#-----
echo "IP Configuration"
#-----
        interface "system"
            no shutdown
        exit
#-----
#
```

8.3.3.4 Configuring VPRN protocols - BGP

The ASN 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

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 OS Router configuration Guide*.

8.3.3.5 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.

Use the CLI syntax to configure VPRN BGP parameters (BGP Configuration Commands on page 691).

Example

The following is a sample VPRN BGP configuration output.

```
*A:K-SASK12>config>service>vprn>bgp$ info detail
-----
      no description
      no authentication-key
      no connect-retry
      no keepalive
      no hold-time
      no damping
      no local-preference
      no loop-detect
      no min-route-advertisement
      no aggregator-id-zero
      no preference
      no remove-private
      no multihop
      no med-out
      no cluster
      no disable-4byte-asn
      no import
```



```
no export
no local-as
no path-mtu-discovery
no router-id
no disable-fast-external-failover
no disable-communities
no advertise-inactive
no enable-peer-tracking
no auth-keychain
no rapid-withdrawal
no split-horizon
no backup-path
best-path-selection
  no always-compare-med
  no deterministic-med
  no as-path-ignore
  no ignore-nh-metric
  no ignore-router-id
exit
next-hop-resolution
  no policy
exit
no peer-tracking-policy
error-handling
  no update-fault-tolerance
exit
no damp-peer-oscillations
rib-management
  ipv4
    no leak-import
  exit
exit
no shutdown
-----
*A:K-SASK12>config>service>vprn>bgp$
```

8.3.3.6 Configuring a VPRN interface

Interface names associate an IP address to the interface, and then associate the IP interface with a physical port. The logical interface can associate attributes like an IP address, port, Link Aggregation Group (LAG) or the system.

There are no default interfaces.

Note that you can configure a VPRN interface as a loop-back interface by issuing the loop-back command instead of the sap sap-id command. The loop-back flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loop-back interface.

Example

The following is a sample VPRN interface configuration output.

```
*A:K-SASK12>config>service>vprn>if$ info detail
-----
no description
no enable-ingress-stats
no enable-mac-accounting
no tcp-mss
cpu-protection 254
no address
no mac
```

```
arp-timeout 14400
arp-retry-timer 50
no arp-limit
no allow-directed-broadcasts
icmp
    mask-reply
    redirects 100 10
    unreachable 100 10
    ttl-expired 100 10
exit
dhcp
    shutdown
    no description
    no option
    no server
    no trusted
    no relay-proxy
    no gi-address
    no relay-plain-bootp
exit
no authentication-policy
no ip-mtu
no delayed-enable
no bfd
no local-dhcp-server
no proxy-arp-policy
no local-proxy-arp
no remote-proxy-arp
no ptp-hw-assist
no qos-route-lookup
load-balancing
    no teid-load-balancing
    no egr-ip-load-balancing
    no spi-load-balancing
exit
no vas-if-type
no shutdown
-----
*A:K-SASK12>config>service>vprn>if$
```

8.3.3.7 Configuring a VPRN interface SAP

A SAP is a combination of a port and encapsulation parameters which identifies the service access point on the interface and within the 7210 SAS. Each SAP must be unique within a router. A SAP cannot be defined if the interface loop-back command is enabled.

When configuring VPRN interface SAP parameters, a default QoS policy is applied to each ingress and egress SAP. Additional QoS policies and scheduler policies must be configured in the **config>qos** context. Filter policies are configured in the **config>filter** context and must be explicitly applied to a SAP. There are no default filter policies.

Example

The following is a sample VPRN interface SAP configuration output.

```
exit
vprn 1435 customer 1 create
shutdown
interface "test" create
exit
```

```
        bgp
          no shutdown
        exit
        ospf
          no shutdown
        exit
      exit
-----
*A:K-SASK12>config>service#
```

8.3.3.8 Configuring VRRP

Configuring VRRP parameters on a VPRN interface is optional. VRRP can be configured in either 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

The following is a sample VRRP interface VRRP owner configuration output.

```
A:K-SASK12>config>service>vprn>if$ info
#-----
...
  interface "vrrpowner"
    address 10.10.10.23/24
    vrrp 1 owner
      backup 10.10.10.23
      authentication-key "vh48l0V7Hs2H6lrMHg2aMJJnZStHwwy0" hash2
    exit
  exit
...
#-----
config>service>vrrp#
```

8.3.3.9 Configuring IPv6 parameters for VPRN BGP

Use the following syntax to configure IPv6 parameters for VPRN BGP.

```
config>service# vprn service-id [customer customer-id]
  bgp
    family ipv6
      group name
        family ipv6
        neighbor ipv6-address
        family ipv6
```

Example

```
A:SAS>config>service# vprn 20
A:SAS>config>service>vprn$ bgp
A:SAS>config>service>vprn>bgp$ family ipv6
A:SAS>config>service>vprn>bgp>family$ group BGP1
A:SAS>config>service>vprn>bgp>family>group$ family ipv6
A:SAS>config>service>vprn>bgp>family>group>family$ neighbor 2001:db8:a0b:12f0::1
```

```
A:SAS>config>service>vprn>bgp>family>group>family> neighbor$ family ipv6
A:SAS>config>service>vprn>bgp>family>group>family> neighbor$ exit
A:SAS>config>service>vprn>bgp>family>group>family$ exit
A:SAS>config>service>vprn>bgp>family>group$ exit
A:SAS>config>service>vprn>bgp>family$ exit
A:SAS>config>service>vprn>bgp$ exit
```

8.3.3.10 Configuring VPRN IPv6 neighbor discovery parameters

Use the following syntax to configure IPv6 neighbor discovery parameters for a VPRN service.

```
config>service# vprn service-id [customer customer-id]
  ipv6
    reachable-time seconds
    stale-time seconds
```

Example

```
A:SAS>config>service# vprn 20
A:SAS>config>service>vprn$ ipv6
A:SAS>config>service>vprn>ipv6# reachable-time 30
A:SAS>config>service>vprn>ipv6# stale-time 14400
A:SAS>config>service>vprn>ipv6# exit
A:SAS>config>service>vprn# exit
```

8.3.3.11 Configuring OSPF for VPRN

Use the following syntax to configure OSPF in the VPRN context.

```
config>service>vprn>ospf#
```

See [OSPF configuration commands](#) for the CLI syntax to configure VPRN OSPF parameters.

Example

The following is a sample VPRN OSPF configuration output.

```
*A:K-SASK12>config>service# info
-----
  vprn 2 customer 1 create
    interface "ospf_interface" create
    exit
  ospf
    area 0.0.0.0
      interface "ospf_interface"
        no shutdown
        exit
      exit
    exit
-----
*A:K-SASK12>service#
-----
```

8.3.4 Service management tasks

This section describes the service management tasks.

8.3.4.1 Modifying VPRN service parameters

Use the syntax to modify VPRN parameters ([VPRN services command reference for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C](#)).

Example

The following is a sample VPRN service creation output.

```
vprn 2 customer 1 create
  shutdown
  no description
  dhcp
  exit
  dhcp6
  exit
  no router-id
  no maximum-routes
  no maximum-ipv6-routes
  snmp
    no access
  exit
  log
  exit
  no mc-maximum-routes
  no autonomous-system
  no route-distinguisher
  auto-bind-tunnel
    resolution-filter
      no ldp
      no rsvp
    exit
  resolution disabled
  exit
  no enable-bgp-vpn-backup
  no vrf-target
  interface "test" create
    no description
    no enable-ingress-stats
    no enable-mac-accounting
    no tcp-mss
    cpu-protection 254
    no address
    no mac
    arp-timeout 14400
    arp-retry-timer 50
    no arp-limit
    no allow-directed-broadcasts
    icmp
      mask-reply
      redirects 100 10
      unreachable 100 10
      ttl-expired 100 10
    exit
  dhcp
    shutdown
    no description
```

```
no option
no server
no trusted
no relay-proxy
no gi-address
no relay-plain-bootp
exit
```

8.3.4.2 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
```

8.3.4.3 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
```

```
vprn 2 customer 1 create
shutdown
no description
dhcp
exit
dhcp6
exit
no router-id
no maximum-routes
no maximum-ipv6-routes
snmp
no access
exit
log
exit
no mc-maximum-routes
```

```
no autonomous-system
no route-distinguisher
auto-bind-tunnel
  resolution-filter
    no ldp
    no rsvp
  exit
resolution disabled
exit
no enable-bgp-vpn-backup
no vrf-target
interface "test" create
  no description
  no enable-ingress-stats
  no enable-mac-accounting
  no tcp-mss
  cpu-protection 254
  no address
  no mac
  arp-timeout 14400
  arp-retry-timer 50
  no arp-limit
  no allow-directed-broadcasts
  icmp
    mask-reply
    redirects 100 10
    unreachable 100 10
    ttl-expired 100 10
  exit
  dhcp
    shutdown
    no description
    no option
    no server
    no trusted
    no relay-proxy
    no gi-address
    no relay-plain-bootp
  exit
#
```

8.3.4.4 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 for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

8.4.1 Command hierarchies

- [VPRN service configuration commands](#)

- DHCP server commands
- IPoE commands
- IPv6 interface commands (supported only on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
- IPv6 router advertisement commands (supported only on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)
- Interface VRRP commands (for network mode only)
- Routed VPLS commands for VPRN service
- BGP configuration commands
- OSPF configuration commands
- Show commands
- Clear commands
- Debug commands

8.4.1.1 VPRN service configuration commands

```

config
  - service
    - vprn service-id [customer customer-id]
    - no vprn service-id
      - aggregate ip-prefix /ip-prefix-length [summary-only] [as-set] [aggregator as-
number: ip-address]
      - no aggregate ip-prefix /ip-prefix-length
      - auto-bind-tunnel
        - resolution {any | filter | disabled}
        - resolution-filter
          - [no] ldp
          - [no] rsvp
          - [no] sr-isis
          - [no] sr-ospf
      - description description-string
      - no description
      - ecmp max-ecmp-routes
      - no ecmp
      - [no] ignore-nh-metric
      - 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
      - [no] shutdown
      - snmp
        - community community community-name [hash | hash2] [access-permissions]
  [version SNMP-version]
    - no community community-name
  - source-address
    - application app [ip-int-name | ip-address]
    - no application app
  - [no] spoke-sdp sdp-id
    - [no] description description string
  - [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
  [preference preference] [metric metric] [tag tag] [enable | disable] {next-hop ip-int-

```



```

name | ip-address | {cpe-check cpe-ip-address [interval seconds] [drop-count count] [log]}]
{prefix-list prefix-list-name [all | none]}]
- [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
[preference preference] [metric metric] [tag tag] [enable | disable] indirect ip-address [cpe-
check cpe-ip-address [interval seconds][drop-count count] [log]] {prefix-list prefix-list-name
[all | none]}]
- [no] static-route {ip-prefix/prefix-length | ip-prefix
netmask} [preference preference] [metric metric] [tag tag] [enable | disable] black-hole
{prefix-list prefix-list-name [all | none]}]
- [no] static-route {ip-prefix/prefix-length | ip-prefix netmask}
[preference preference][metric metric] [tag tag] [enable | disable] next-hop ip-int-name | ip-
address [bfd-enable | {cpe-check cpe-ip-address [interval seconds] [drop-count count] [log]}]
{prefix-list prefix-list-name [all|none]}][description description]
- vrf-export policy-name [policy-name...(up to 5 max)]
- no vrf-export
- vrf-import policy-name [policy-name...(up to 5 max)]
- no vrf-import
- vrf-target {ext-community | export ext-community | import ext-community}
- no vrf-target
- [no] shutdown

config
- service
- vprn service-id [customer customer-id]
- no vprn service-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
    
```



Note:

See the 7210 SAS-K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Quality of Service Guide, section “Self-generated Traffic commands for 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C” for descriptions of the **config>service>vprn>sgt-qos** commands.

8.4.1.2 DHCP server commands

```

config
- service
- vprn service-id [customer customer-id]
- no vprn service-id
- dhcp
- local-dhcp-server server-name [create]
- no local-dhcp-server server-name
- description description-string
- no description
- [no] force-renews
- pool pool-name [create]
- no pool pool-name
- description description-string
- no description
- max-lease-time [days days] [hrs hours] [min minutes] [sec seconds]
- no max-lease-time
- min-lease-time [days days] [hrs hours] [min minutes] [sec seconds]
- no min-lease-time
- minimum-free minimum-free [percent] [event-when-depleted]
- no minimum-free
- [no] nak-non-matching-subnet
- offer-time [min minutes] [sec seconds]
    
```

```

- no offer-time
- options
  - custom-option option-number address [ip-address...(up to 4 max)]
  - custom-option option-number hex hex-string
  - custom-option option-number string ascii-string
  - no custom-option option-number
  - dns-server [ip-address...(up to 4 max)]
  - domain-name domain-name
  - no domain-name
  - lease-rebind-time [days days] [hrs hours] [min minutes]

[sec seconds]
- no lease-rebind-time
- lease-renew-time [days days] [hrs hours] [min minutes]

[sec seconds]
- no lease-renew-time
- lease-time [days days] [hrs hours] [min minutes] [sec seconds]
- no lease-time
- subnet {ip-address/mask|ip-address netmask} [create]
- no subnet {ip-address/mask|ip-address netmask}
  - [no] address-range start-ip-address end-ip-address
  - [no] exclude-addresses start-ip-address [end-ip-address]
  - maximum-declined maximum-declined
  - no maximum-declined
  - minimum-free minimum-free
  - no minimum-free
  - options
    - custom-option option-number address [ip-address...(up to 4
max)]
    - custom-option option-number hex hex-string
    - custom-option option-number string ascii-string
    - no custom-option option-number
    - default-router ip-address [ip-address...(up to 4 max)]
    - no default-router
    - subnet-mask ip-address
    - no subnet-mask
- [no] use-gi-address
- user-db local-user-db-name
- no user-db
    
```

8.4.1.3 IPoE commands

```

config
- subscriber-mgmt
  - local-user-db local-user-db-name [create]
  - no local-user-db local-user-db-name
  - description description-string
  - no description
  - ipoe
    - host host-name [create]
    - no host host-name
      - address gi-address [scope scope]
      - address ip-address
      - address pool pool-name [secondary-pool sec-pool-name]
[delimiter delimiter]
- address use-pool-from-client [delimiter delimiter]
- no address
- gi-address ip-address
- no gi-address
- host-identification
  - circuit-id string ascii-string
  - circuit-id hex hex-string
    
```

```

- no circuit-id
- encap-tag-range start-tag start-tag end-tag end-tag
- no encap-tag-range
- mac ieee-address
- no mac
- option60 hex-string
- no option60
- remote-id hex hex-string
- remote-id string ascii-string
- no remote-id
- sap-id sap-id
- no sap-id
- service-id service-id
- no service-id
- string string
- no string
- system-id system-id
- no system-id
- shutdown
- no shutdown
- options
- custom-option option-number address [ip-address...(up to 4 max)]
- custom-option option-number hex hex-string
- custom-option option-number string ascii-string
- no custom-option option-number
- default-router ip-address [ip-address (up to 4 max)]
- no default-router
- dns-server ip-address [ip-address...(up to 4 max)]
- no dns-server
- domain-name domain-name
- no domain-name
- lease-rebind-time [days days] [hrs hours] [min minutes] [sec seconds]
- no lease-rebind-time
- lease-renew-time [days days] [hrs hours] [min minutes] [sec seconds]
- no lease-renew-time
- lease-time [days days] [hrs hours] [min minutes] [sec seconds]
- no lease-time
- subnet-mask ip-address
- no subnet-mask
- mask type ipoe-match-type {[prefix-string prefix-string | prefix-
length prefix-length] [suffix-string suffix-string | suffix-length suffix-length]}
- no mask type ipoe-match-type
- match-list ipoe-match-type-1 [ipoe-match-type-2 (up to 4 max)]
- no match-list
- shutdown
- no shutdown
    
```

8.4.1.4 Interface commands

```

config
- service
- vprn service-id [customer customer-id] [create]
- no vprn service-id
- [no] interface ip-int-name
- address ip-address[/mask] [netmask] [broadcast {all-ones | host-ones}]
- no address
- [no] allow-directed-broadcasts
- arp-timeout [seconds]
- no arp-timeout
- delayed-enable seconds
- no delayed-enable
    
```

```
- description description-string
- no description [description-string]
- dhcp
  - description description-string
  - no description
  - gi-address ip-address [src-ip-addr]
  - no gi-address
  - [no] option
    - action {replace | drop | keep}
    - no action
    - circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tuple]
    - no circuit-id
    - remote-id [mac | string string]
    - [no] vendor-specific-option
      - [no] client-mac-address
      - [no] pool-name
      - [no] sap-id
      - [no] service-id
      - string text
      - no string
      - [no] system-id
  - server server1 [server2... (up to 8 max)]
  - no server
  - [no] shutdown
  - [no] trusted
- icmp
  - [no] mask-reply
  - redirects number seconds
  - no redirects [number seconds]
  - ttl-expired number seconds
  - no ttl-expired [number seconds]
  - unreachables number seconds
  - no unreachables [number seconds]
- [no] local-proxy-arp
- [no] loopback
- [no] remote-proxy-arp policy-name [policy-name ... (up to 5 max)]
- [no] remote-proxy-arp
- [no] shutdown
- static-arp ip-address ieee-address
- no static-arp ip-address [ieee-address]
- static-arp ieee-address unnumbered
- no static-arp [ieee-address] unnumbered
- unnumbered ip-int-name | ip-address
- no unnumbered
- [no] vrrp virtual-router-id
```

8.4.1.5 IPv6 interface commands (supported only on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

```
config
- service
  - vprn service-id [customer customer-id] [create]
    - [no] interface ip-int-name [create]
      - [no] urpf-check
      - ipv6
        - [no] urpf-check
        - no ipv6
        - address ipv6-address/prefix-length [eui-64] [preferred]
        - no address ipv6-address/prefix-length
      - icmp
        - packet-too-big number seconds
```

```
- no packet-too-big
- param-problem number seconds
- no param-problem
- redirects number seconds
- no redirects
- time-exceeded number seconds
- no time-exceeded
- unreachable number seconds
- no unreachable
- link-local-address ipv6-address [preferred]
- local-proxy-nd
- no local-proxy-nd
- neighbor ipv6-address mac-address
- no neighbor ipv6-address
- proxy-nd-policy policy-name [policy-name...(up to 5 max)]
- no proxy-nd-policy
- [no] urpf-check
- urpf-check
- [no] ignore-default
- ipv6
- [no] ignore-default
```

8.4.1.6 IPv6 router advertisement commands (supported only on 7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C)

```
config
- router [router-name]
- router-advertisement
- no router-advertisement
- interface ip-int-name
- no interface ip-int-name
- current-hop-limit number
- no current-hop-limit
- managed-configuration
- no managed-configuration
- max-advertisement-interval seconds
- no max-advertisement-interval
- max-advertisement-interval seconds
- no min-advertisement-interval
- mtu
- no mtu
- other-stateful-configuration
- no other-stateful-configuration
- prefix ipv6-prefix/prefix-length
- no prefix
- autonomous
- no autonomous
- on-link
- 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
- shutdown
```

8.4.1.7 Interface VRRP commands (for network mode only)

```
config
- service
- vprn service-id [customer customer-id] [create]
- no vprn service-id
- interface ip-int-name [create]
- no interface ip-int-name
- vrrp virtual-router-id [owner]
- no vrrp virtual-router-id
- authentication-key authentication-key | hash-key [hash | hash2]
- no authentication-key
- [no] backup ip-address
- [no] bfd-enable service-id interface interface-name dst-ip ip-address
- [no] bfd-enable interface interface-name dst-ip ip-address
- init-delay seconds
- no init-delay
- [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
```

8.4.1.8 Routed VPLS commands for VPRN service

```
config
- service
- vprn
- interface ip-interface-name [create]
- no interface ip-interface-name
- vpls service-name
- no vpls
- ingress
- v4-routed-override-filter ip-filter-id
- no v4-routed-override-filter
```

8.4.1.9 Interface SAP commands

```
config
- service
- vprn service-id [customer customer-id]
- no vprn service-id
- [no] interface ip-int-name [create] [tunnel]
- [no] sap sap-id
- accounting-policy acct-policy-id
- no accounting-policy [acct-policy-id]
- [no] collect-stats
```

```
- description description-string
- no description [description-string]
- egress
  - agg-shaper-rate agg-rate
  - no agg-shaper-rate
  - filter ip ip-filter-id
  - no filter [ip ip-filter-id]
  - qos policy-id
  - no qos [policy-id]
- ingress
  - meter-override
  - no meter-override
  - meter meter-id [create]
  - no meter meter-id
    - adaptation-rule [pir adaptation-rule] [cir adaptation-rule]
    - cbs size [kbits | bytes | kbytes]
    - no cbs
    - mbs size [kbits | bytes | kbytes]
    - no mbs
    - no mode
    - no mode
    - rate cir cir-rate [pir pir-rate]
- ingress
  - agg-shaper-rate agg-rate
  - no agg-shaper-rate
  - aggregate-meter-rate rate-in-kbps [burst burst-in-kbits]
  - no aggregate-meter-rate
  - filter ip ip-filter-id
  - no filter [ip ip-filter-id]
  - qos policy-id [shared-queuing | multipoint-shared]
  - no qos [policy-id]
- [no] shutdown
```

8.4.1.10 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
      - bfd-enable
      - no bfd-enable
      - best-path-selection
        - always-compare-med {zero | infinity}
        - always-compare-med strict-as {zero | infinity}
        - no always-compare-med
        - as-path-ignore [ipv4] [ipv6]
        - no as-path-ignore
        - as-path-ignore
        - no ignore-nh-metric
        - ignore-nh-metric
      - no ignore-router-id
      - ignore-nh-metric
```

```
- [no] connect-retry seconds
- [no] damping
- description description-string
- no description
- [no] disable-4byte-asn
- 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-as as-number [private]
- no local-as
- local-preference local-preference
- no local-preference
- loop-detect {drop-peer | discard-route | ignore-loop| off}
- no loop-detect
- med-out {number | igp-cost}
- no med-out
- min-as-origination seconds
- no min-as-origination
- min-route-advertisement seconds
- no min-route-advertisement
- multihop ttl-value
- no multihop
- preference preference
- no preference
- [no] rapid-withdrawal
- [no] remove-private
- router-id ip-address
- no router-id
- [no] shutdown
- [no] group name [dynamic-peer]
  - [no] advertise-inactive
  - [no] aggregator-id-zero
  - [no] as-override
  - auth-keychain name
  - authentication-key [authentication-key | hash-key] [hash | hash2]
  - no authentication-key
  - [no] bfd-enable
  - 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] passive
- peer-as as-number
- no peer-as
- preference preference
- no preference
- prefix-limit prefix-limit family limit [log-only] [threshold percentage]
- no prefix-limit
- [no] remove-private
- [no] shutdown
- type {internal | external}
- no type
- [no] neighbor ip-address
  - [no] advertise-inactive
  - [no] aggregator-id-zero
  - [no] as-override
  - auth-keychain name
  - authentication-key [authentication-key | hash-key] [hash | hash2]
  - no authentication-key
  - [no] bfd-enable
  - 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]
  - 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] passive
- peer-as as-number
- no peer-as
- preference preference
- no preference
- prefix-limit limit [log-only] [threshold percent]
- no prefix-limit
- [no] remove-private
- [no] shutdown
- type {internal | external}
- no type
```

8.4.1.11 OSPF configuration commands

```
config
- service
  - vprn service-id [customer customer-id]
  - no vprn service-id
    - [no] ospf
      - [no] area area-id
        - area-range ip-prefix/mask [advertise | not-advertise]
        - no area-range ip-prefix/mask
        - [no] blackhole-aggregate
        - [no] interface ip-int-name [secondary]
          - [no] advertise-subnet
          - authentication-key [authentication-key | hash-key] [hash | hash2]
          - no authentication-key
          - authentication-type {password | message-digest}
          - no authentication-type
          - bfd-enable [remain-down-on-failure]
          - no bfd-enable
          - dead-interval seconds
          - no dead-interval
          - hello-interval seconds
          - no hello-interval
          - interface-type {broadcast | point-to-point}
          - no interface-type
          - message-digest-key key-id md5 [key | hash-key] [hash | hash2]
          - no message-digest-key key-id
          - metric metric
          - no metric
          - mtu bytes
          - no mtu
          - [no] passive
          - priority number
          - no priority
          - retransmit-interval seconds
          - no retransmit-interval
          - [no] shutdown
          - transit-delay seconds
```

```
- no transit-delay
- [no] nssa
  - area-range ip-prefix/mask [advertise | not-advertise]
  - no area-range ip-prefix/mask
  - originate-default-route [type-7]
  - no originate-default-route
  - [no] redistribute-external
  - [no] summaries
- [no] sham-link ip-int-name ip-address
  - authentication-key [authentication-key | hash-key] [hash | hash2]
  - no authentication-key
  - authentication-type {password | message-digest}
  - no authentication-type
  - dead-interval seconds
  - no dead-interval
  - hello-interval seconds
  - no hello-interval
  - message-digest-key key-id md5 [key | hash-key] [hash | hash2]
  - no message-digest-key key-id
  - metric metric
  - no metric
  - retransmit-interval seconds
  - no retransmit-interval
  - [no] shutdown
  - transit-delay seconds
  - no transit-delay
- [no] stub
  - default-metric metric
  - no default-metric
  - [no] summaries
- [no] virtual-link router-id transit-area area-id
  - authentication-key [authentication-key | hash-key] [hash | hash2]
  - no authentication-key
  - authentication-type {password | message-digest}
  - no authentication-type
  - dead-interval seconds
  - no dead-interval
  - hello-interval seconds
  - no hello-interval
  - message-digest-key key-id md5 [key | hash-key] [hash | hash2]
  - no message-digest-key key-id
  - retransmit-interval seconds
  - no retransmit-interval
  - [no] shutdown
  - transit-delay seconds
  - no transit-delay
- [no] compatible-rfc1583 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] graceful-restart
  - [no] helper-disable
- [no] ignore-dn-bit
- overload [timeout seconds]
- no overload
- [no] overload-include-stub
- overload-on-boot [timeout seconds]
- no overload-on-boot
- preference preference
- no preference
- reference-bandwidth bandwidth-in-kbps
- no reference-bandwidth
- router-id ip-address
- no router-id
- [no] shutdown
- [no] super-backbone
- [no] suppress-dn-bit
- timers
  - [no] lsa-arrival lsa-arrival-time
  - [no] lsa-generate max-lsa-wait [lsa-initial-wait [lsa-second-wait]]
  - [no] spf-wait max-spf-wait [spf-initial-wait [spf-second-wait]]
- vpn-domain id {0005 | 0105 | 0205 | 8005}
- no vpn-domain
- vpn-tag vpn-tag
- no vpn-tag
- 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] graceful-restart
  - [no] helper-disable
- import policy-name [policy-name ... (up to 15 max)]
- no import policy-name
- overload [timeout seconds]
- no overload
- [no] overload-include-stub
- overload-on-boot [timeout seconds]
- no overload-on-boot
- preference preference
- no preference
- reference-bandwidth bandwidth-in-kbps
- reference-bandwidth [tbps Tera-bps] [gbps Giga-bps] [mbps Mega-bps] [kbps Kilo-
bps]
- no reference-bandwidth
- router-id ip-address
- no router-id
- [no] shutdown
- 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]]
```

8.4.1.12 Show commands

```
show
- service
  - egress-label start-label [end-label]
  - ingress-label start-label [[end-label]
  - id service-id
    - all
    - sap [sap-id [detail]]
    - sdp [sdp-id | far-end ip-address] [detail]
  - sap-using [sap sap-id]
  - sap-using interface [ip-address | ip-int-name]
  - sap-using [ingress | egress] filter filter-id
  - sap-using [ingress | egress] qos-policy qos-policy-id
  - sdp-using [sdp-id | far-end ip-address] [detail | keep-alive-history]
  - sdp-using [sdp-id[:vc-id]
  - service-using [vprn] [sdp sdp-id] [customer customer-id]
```

```
show
- router [vprn-service-id]
  - aggregate [family] [active]
  - arp [ip-int-name | ip-address[/mask] | mac ieee-mac-address | sumary] [local |
dynamic | static | managed]
  - bgp
    - damping [ip-prefix[/prefix-length]] [decayed|history|suppressed] [detail] [ipv4]
    - damping [ip-prefix[/prefix-length]] [decayed|history|suppressed] [detail] vpn-
ipv4
  - group [name] [detail]
  - neighbor [ip-address [detail]]
  - neighbor [as-number [detail]]
  - neighbor ip-address [family] filter1 [brief]
  - neighbor ip-address [family] filter2
  - neighbor as-number [family] filter2
  - neighbor ip-address orf filter3
  - neighbor ip-address graceful-restart
  - 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
  - route-table [ip-address[/mask]] [longer | best]} | [protocol protocol] | [summary]
  - static-arp [ip-address | ip-int-name | mac ieee-mac-address]
  - static-route [ip-prefix /mask] | [preference preference] | [next-hop ip-address]
[detail]
  - tunnel-table [ip-address[/mask]] [protocol protocol | sdp sdp-id]
  - tunnel-table [summary]
```

```
show
- router
  - sgt-qos
  - application
  - dscp-map
```

**Note:**

See the 7210 SAS-K 2F1C2T, K 2F6C4T, K 3SFP+ 8C Quality of Service Guide section “Network QoS Policy Command Reference, Show Commands” for descriptions of the **show>router>sgt-qos** commands.

8.4.1.13 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
    - interface [ip-int-name | ip-address] [icmp]
clear
- service
  - id service-id
  - spoke-sdp sdp-id:vc-id ingress-vc-label
  - statistics
    - sap sap-id {all | counters | stp}
    - sdp sdp-id keep-alive
  - id service-id
    - counters
    - spoke-sdp sdp-id:vc-id {all | counters | stp}
    - spoke-sdp
```

8.4.1.14 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

8.4.2.1 VPRN service configuration commands

8.4.2.1.1 Generic commands

shutdown

Syntax

[no] shutdown

Context

```
config>service>vprn
config>service>vprn>if
config>service>vprn>if>sap
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
config>service>vprn>ospf>area>sham-link
config>service>vprn>ospf>area>virtual-link
config>service>vprn>router-advert
config>service>vprn>spoke-sdp
config>service>vprn>if>vrrp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

If the ASN was previously changed, the BGP ASN inherits the new value.

Special Cases

Service Admin State

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

A service is regarded as operational providing that one IP Interface SAP and one SDP is operational.

VPRN BGP

This command 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 shut down operationally. If a BGP group is shut down, all member neighbor interfaces are shut down operationally. If a BGP neighbor is shut down, just that neighbor interface is operationally shutdown.

description

Syntax

description *description-string*

no description

Context

```
config>service>vprn>bgp
config>service>vprn
config>service>vprn>if
config>service>vprn>if>sap
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

description-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.

ecmp

Syntax

ecmp *max-ecmp-routes*

no ecmp

Context

config>service>vpn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables ECMP only for VPN-IPv4 routes learned through MP-BGP (not PE-CE VPN routes) and configures the number of routes for path sharing; for example, the value 2 means two equal-cost routes are used for cost sharing.

ECMP can be used only for routes learned with the same preference and the same protocol. When more ECMP routes are available at the best preference than configured using the *max-ecmp-routes* value, the following route selection criteria apply for VPN-IPv4 routes.

- All routes are flagged as BEST, and the following criteria determine which routes are submitted to the VRF and marked as USED according to the ECMP value configured in the VRF.
 - Routes with the lowest BGP peer router-ID are preferred.
 - Routes with the shortest cluster list are preferred.
 - Routes with the lowest BGP next-hop IP address are preferred.

The **no** form of this command disables ECMP path sharing. If ECMP is disabled, and multiple routes are available at the best preference and equal cost, the VPN-IPv4 route selection is based on the following criteria.

- Routes with the lowest BGP peer router-ID are preferred.
- Routes with the shortest cluster list are preferred.
- Routes with the lowest BGP next-hop IP address are preferred.

Default

no ecmp

Parameters

max-ecmp-routes

Specifies the maximum number of equal cost routes allowed on this routing table instance, expressed as a decimal integer. Setting the *max-ecmp-routes* value to 1 has the same result as entering **no ecmp**.

Values 1 to 4

8.4.2.1.2 Global commands

vprn

Syntax

```
vprn service-id [customer customer-id] [create]
no vprn service-id
```

Context

```
config>service
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 belongs to the same customer.

By default, no VPRN services instances exist until they are explicitly created.

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 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.

vprn customer

Values *service-id*: 1 to 2147483648
svc-name: 64 characters maximum

customer *customer-id*

Specifies an existing customer identification number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

aggregate

Syntax

aggregate *ip-prefix/ip-prefix-length* [**summary-only**] [**as-set**] [**aggregator** *as-number:ip-address*]
aggregate *ip-prefix/ip-prefix-length*

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures an aggregate route.

This command automatically installs an aggregate in the routing table when there are one or more component routes. A component route is any route used for forwarding that is a more specific match to the aggregate.

The use of aggregate routes can reduce the number of routes that need to be advertised to neighbor routers, leading to smaller routing table sizes.

Overlapping aggregate routes may be configured; in this case a route becomes a component of only the one aggregate route with the longest prefix match; for example, if one aggregate is configured as 10.0.0.0/16 and another as 10.0.0.0/24, route 10.0.128/17 would be aggregated into 10.0.0.0/16, and route 10.0.0.128/25 would be aggregated into 10.0.0.0/24. If multiple entries are made with the same prefix and the same mask the previous entry is overwritten.

A standard 4-byte BGP community may be associated with an aggregate route to facilitate route policy matching.

The **no** form of this command removes the aggregate.

Default

no aggregate

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

the mask associated with the network address expressed as a mask length

Values: 0 to 32

summary-only

Keyword that suppresses the advertisement of more specific component routes for the aggregate.

To remove the **summary-only** option, enter the same aggregate command without the **summary-only** keyword.

as-set

Optional keyword that is only applicable to BGP and creates an aggregate where the path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized. Use this option carefully, because it can increase the amount of route churn because of best path changes.

aggregator as-number:ip-address

Optional keyword that specifies the BGP aggregator path attribute to the aggregate route. When configuring the aggregator, a two-octet ASN used to form the aggregate route must be entered, followed by the IP address of the BGP system that created the aggregate route.

allow-ip-int-bind

Syntax

[no] **allow-ip-int-bind**

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables IP interface binding.

The **no** form of this command disables IP interface binding.

Default

no allow-ip-int-bind

auto-bind-tunnel

Syntax

auto-bind-tunnel

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure automatic binding of a VPRN service using tunnels to MP-BGP peers.

The **auto-bind-tunnel** context configures the binding of VPRN routes to tunnels. 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** command is set to **any**, any supported tunnel type in the 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 will be 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 (in the **configure>service>vprn>spoke-sdp** context), 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 the **ldp** command 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 **sr-isis** 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-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command defines the autonomous system (AS) to be used by this VPN routing/forwarding (VRF). This command defines the autonomous system to be used by this VPN routing.

The **no** form of 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

maximum-ipv6-routes

Syntax

maximum-ipv6-routes *number* [**log-only**] [**threshold percent**]

no maximum-ipv6-routes

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the maximum number of remote IPv6 routes that are allowed to be learned within a VPN VRF context. Local, host, static, and aggregate routes are not counted.

The VPRN service ID must be in a shutdown state to modify the **maximum-ipv6-routes** command parameters.

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 has already been reached. Protocols resubmit the routes that were initially rejected.

This command only applies to PE-CE protocols (on the 7210 SAS, BGP and OSPF are supported for PE-CE dynamic routing protocols). There are two limits to consider: one is the value configured as part of the **maximum-ipv6-routes** command and the second is the maximum IP FIB limit supported on the node.

When a route is received, the following comparison is completed:

- The protocol session remains up but the routes exceeding the configured value are not added to the VRF if all of the following conditions are true:
 - if the number of routes in the FIB plus the received route is greater than the value configured for **maximum-ipv6-routes**
 - if the number of routes is less than the maximum IP FIB limit
 - if **log-only** is not configured
- The protocol session remains up and the routes that exceed the limit of values configured for **maximum-ipv6-routes** are added to the VRF if all of the following conditions are true:
 - if the number of routes in the FIB plus the received route is greater than the value configured for **maximum-ipv6-routes**
 - if the number of routes is less than the maximum IP FIB limit
 - if **log-only** is configured

A log is generated to indicate that the limit has been exceeded.

- If the number of routes in the FIB plus the received route is greater than the maximum IP FIB limit (the limit for the prefix-length of the received route), regardless of whether **log-only** is configured, the protocol session is brought down.

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 shutdown.

Default

0 or disabled

Parameters

number

Specifies the maximum number of routes allowed to be learned in a VRF context.

Values 1 to 2147483647

log-only

Keyword to specify that if the maximum limit is reached, only log the event. This keyword does not disable learning new routes.

threshold percent

Specifies the percentage at which a warning log message and SNMP trap should be set. There are two warnings; the first is a mid-level warning at the threshold value set, and the second is a high-level warning at level between the maximum number of routes and the mid-level rate ($[\text{mid}+\text{max}] / 2$).

Values 0 to 100

maximum-routes

Syntax

maximum-routes *number* [**log-only**] [**threshold percent**]

no maximum-routes

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the maximum number of remote routes that are allowed to be learned within a VPN VRF context. Local, host, static, and aggregate routes are not counted.

The VPRN service ID must be in a shutdown state to modify the **maximum-routes** command parameters.

If the **log-only** keyword 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.

This command only applies to PE-CE protocols (on the 7210 SAS, BGP and OSPF are supported for PE-CE dynamic routing protocols). There are two limits to consider: one is the value configured as part of the **maximum-routes** command and the second is the maximum IP FIB limit supported on the node. When a route is received, the following comparison is completed:

- The protocol session remains up but the routes exceeding the configured value are not added to the VRF if all of the following conditions are true:
 - if the number of routes in the FIB plus the received route is greater than the value configured for **maximum-routes**
 - if the number of routes is less than the maximum IP FIB limit

- if **log-only** is not configured
- The protocol session remains up and the routes that exceed the limit of values configured for **maximum-routes** are added to the VRF if all of the following conditions are true:
 - if the number of routes in the FIB plus the received route is greater than the value configured for **maximum-routes**
 - if the number of routes is less than the maximum IP FIB limit
 - if **log-only** is configured

A log is generated to indicate that the limit has been exceeded.

- If the number of routes in the FIB plus the received route is greater than the maximum IP FIB limit (the limit for the prefix-length of the received route), regardless of whether **log-only** is configured, the protocol session is brought down.

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 shutdown.

Default

0 or disabled

Parameters

number

Specifies the maximum number of routes allowed to be learned in a VRF context.

Values 1 to 2147483647

log-only

Keyword to specify that if the maximum limit is reached, only log the event. This keyword does not disable learning new routes.

threshold *percent*

Specifies the percentage at which a warning log message and SNMP trap should be set. There are two warnings; the first is a mid-level warning at the threshold value set, and the second is a high-level warning at level between the maximum number of routes and the mid-level rate ($[\text{mid}+\text{max}] / 2$).

Values 0 to 100

route-distinguisher

Syntax

route-distinguisher [*ip-address:number* | *asn:number*]

no route-distinguisher

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command sets the identifier attached to routes the VPN belongs to. Each routing instance must have a unique (within the carrier domain) route distinguisher associated with it. A route distinguisher must be defined for a VPRN to be operationally active.

Default

no route-distinguisher

Parameters

The route distinguisher is a 6-byte value that can be specified in one of the following formats:

ip-address:number

Specifies the IP address in dotted-decimal notation. The assigned number must not be greater than 65535.

asn:number

Specifies the AS number as a 2-byte value less than or equal to 65535. The assigned number can be any 32-bit unsigned integer value.

router-id

Syntax

router-id *ip-address*

no router-id

Context

config>service>vprn

config>service>vprn>bgp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures an optional service name that 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 or 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, up to 64 characters. Service names may not begin with an integer (0 to 9).

snmp

Syntax

snmp

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure SNMP parameters for this VPRN.

community

Syntax

community *community-name* [**hash** | **hash2**] [*access-permissions*] [**version** *SNMP-version*]
no community [*community-name*]

Context

config>service>vprn>snmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the SNMP community names to be used with the associated VPRN instance. These VPRN community names associate SNMP v1/v2c requests with a particular VPRN context and to return a reply that contains VPRN-specific data or limit SNMP access to data in a specific VPRN instance.

VPRN SNMP communities configured with an access permission of "r" are automatically associated with the default access group "snmp-vprn-ro" and the "vprn-view" view (read only). VPRN SNMP communities configured with an access permission of "rw" are automatically associated with the default access group "snmp-vprn" and the "vprn-view" view (read/write).

The community in an SNMP v1/v2 request determines the SNMP context (the VPRN number for accessing SNMP tables) and not the VPRN of the incoming interface on which the request is received. When an SNMP request arrives on VPRN 5 interface "ringo" with a destination IP address equal to the "ringo" interface, but the community in the SNMP request is the community configured against VPRN 101, the SNMP request will be processed using the VPRN 101 context (the response will contain information about VPRN 101). Nokia recommends to avoid using a simple series of vprn snmp-community values that are similar to each other (for example, avoid my-vprncomm-1, my-vprn-comm-2, and so on).

By default, the SNMP community must be explicitly specified.

The **no** form of this command removes the SNMP community name from the specific VPRN context.

Parameters

community-name

Specifies the SNMP v1/v2c community name. This is a secret/confidential key used to access SNMP and specify a context (base vs vprn1 vs vprn2).

hash

Keyword that specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2

Keyword that specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is

assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

version *SNMP-version*

Specifies the SNMP version.

Values v1, v2c, both

access-permissions

Specifies the access rights to MIB objects.

Values **r** — Keyword that grants only read access to MIB objects. Creates an association of the community-name with the **snmp-vprn-ro** access group.

rw — Keyword that grants read and write access to MIB objects. Creates an association of the community-name with the **snmp-vprn** access group.

source-address

Syntax

source-address

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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>vprn>source-address

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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, etc.), the entire string must be enclosed within double quotes.

static-route

Syntax

```
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag tag] [enable | disable] {next-hop ip-int-name | ip-address | ipsec-tunnel ipsec-tunnel-name} [bfd-enable | {cpe-check cpe-ip-address [interval seconds] [drop-count count] [log]}
```

```
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag tag] [enable | disable] indirect ip-address [cpe-check cpe-ip-address [interval seconds][drop-count count] [log]
```

```
[no] static-route {ip-prefix/prefix-length | ip-prefix netmask} [preference preference] [metric metric] [tag tag] [enable | disable] black-hole
```

Context

```
config>service>vpn
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

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 is rejected.

If a **static-route** command is issued with **no cpe-check** target, but the destination prefix/netmask and next hop matches a static route that did have an associated **cpe-check**, the cpe-check test will be removed from the associated static route.

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 as are required to uniquely identify the static route must be entered.

Parameters

ip-prefix

Specifies the destination address of the aggregate route in dotted-decimal notation.

Values	<i>ipv4-prefix:</i>	a.b.c.d (host bits must be 0)
	<i>ipv4-prefix-length:</i>	0 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)

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, etc.), the entire string must be enclosed with double quotes.

ip-address

Specifies the IP address of the IP interface. The *ip-addr* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted-decimal notation.

Values ipv4-address a.b.c.d (host bits must be 0)

enable

Keyword that specifies static routes can be administratively enabled or disabled. Use the **enable** parameter to reenable 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

Keyword that specifies 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 second

drop-count count

Specifies the number of consecutive ping-replies that must be missed to declare the CPE down and to deactivate the associated static route.

Values Value range: 1 to 255

Default 3

log

Optional keyword that 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 will be applied.

The *ip-address* configured here can be either on the network side or the access side on this node. This address must be associated with a network directly connected to a network configured on this node.

indirect *ip-address*

Specifies that the route is indirect and specifies the next-hop IP address used to reach the destination.

The configured *ip-address* 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-address* configured can be either on the network or the access side and is at least one hop away from this node.

black-hole

Keyword that 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 the **next-hop** or **indirect** keywords. If an identical command is entered, with exception of 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

Keyword that 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

Keyword that associates the state of the static route to a BFD session between the local system and the configured next hop. This keyword cannot be configured if the next hop has **indirect** or a **blackhole** keywords specified.

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 itself 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the export policies to control routes exported from the local VPN VRF to other VRFs on the same or remote PE routers (via MP-BGP).

By default, no routes are exported from the VRF.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command sets the import policies to control routes imported to the local VPN 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.

By default, no routes are accepted into the VRF.

The **no** form of this command removes all route policy names from the import list.

Parameters

policy

Specifies the route policy statement name.

8.4.2.1.3 Router DHCP configuration commands

dhcp

Syntax

dhcp

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure DHCP parameters.

local-dhcp-server

Syntax

local-dhcp-server *server-name* [create]

no local-dhcp-server *server-name*

Context

config>service>vprn>dhcp

config>service>vprn>if

config>service>vprn>nw-if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command instantiates a local DHCP server. A local DHCP server can serve multiple interfaces but is limited to the routing context in which it was created.



Note:

The DHCP server is supported only for craft access for local management of the node. It must not be used as a general purpose DHCP server.

Parameters

server-name

Specifies the name of local DHCP server.

create

Keyword used to create the server name. The **create** keyword requirement can be enabled and disabled in the **environment>create** context.

force-renews

Syntax

[no] **force-renews**

Context

config>service>vprn>dhcp>server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the sending of "forcerenew" messages.

The **no** form of this command disables the sending of "forcerenew" messages.

Default

no force-renews

pool

Syntax

pool *pool-name* [create]

no pool *pool-name*

Context

config>service>vprn>dhcp>local-dhcp-server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a DHCP address pool on the router.

Parameters

pool name

Specifies the name of this IP address pool. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters.

create

Keyword used to create the entity. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

max-lease-time

Syntax

max-lease-time [*days days*] [*hrs hours*] [*min minutes*] [*sec seconds*]

no max-lease-time

Context

config>service>vprn>dhcp>local-dhcp-server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the maximum amount of time that a client can lease the IP address. The **no** form of this command reverts to the default value.

Default

max-lease-time days 10

Parameters

days

Specifies the number of days in the lease time.

Values 0 to 3650

hours

Specifies the number of hours in the lease time.

Values 0 to 23

minutes

Specifies the number of minutes in the lease time.

Values 0 to 59

seconds

Specifies the number of seconds in the lease time.

Values 0 to 59

min-lease-time

Syntax

min-lease-time [*days days*] [*hrs hours*] [*min minutes*] [*sec seconds*]

no min-lease-time

Context

config>service>vpn>dhcp>local-dhcp-server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the minimum amount of time that a client can lease the IP address.

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

Default

min-lease-time min 10

Parameters

days

Specifies the number of days in the lease time.

Values 0 to 3650

hours

Specifies the number of hours in the lease time.

Values 0 to 23

minutes

Specifies the number of minutes in the lease time.

Values 0 to 59

seconds

Specifies the number of seconds in the lease time.

Values 0 to 59

minimum-free

Syntax

minimum-free *minimum-free* [percent] [event-when-depleted]

no minimum-free

Context

config>router>dhcp>server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the desired minimum number of free addresses in this pool.

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

Default

1

Parameters

minimum-free

Specifies the minimum number of free addresses.

Values 0 to 255

percent

Keyword to specify that the value is a percentage.

event-when-depleted

Keyword that enables a system-generated event when all available addresses in the pool or subnet of the local DHCP server are depleted.

nak-non-matching-subnet

Syntax

[no] **nak-non-matching-subnet**

Context

config>router>dhcp>server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the local DHCPv4 server handling of an Option 50 request. Option 50 indicates a client request for a previously allocated message, as described in section 3.2 of RFC 2131, *Dynamic Host Configuration Protocol*.

When this command is enabled, the address allocation algorithm uses a pool to search for the address specified in Option 10. If the address is not found in the pool, the system returns a DHCP NAK message. If the message is found, the system drops the DHCP packet.

Default

no nak-non-matching-subnet

offer-time

Syntax

offer-time [min *minutes*] [sec *seconds*]

no offer-time

Context

config>service>vprn>dhcp>server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the offer time.

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

Default

1 minute

Parameters

minutes

Specifies the offer time in minutes.

Values 0 to 10

seconds

Specifies the offer time in seconds.

Values 0 to 59

options

Syntax

options

Context

config>service>vprn>dhcp>server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure pool options. The options defined in this context can be overruled if the same option is defined in the local user database.

custom-option

Syntax

custom-option *option-number* **address** [*ip-address...*(up to 4 max)]

custom-option *option-number* **hex** *hex-string*

custom-option *option-number* **string** *ascii-string*

no custom-option *option-number*

Context

config>service>vprn>dhcp>server>pool>options

config>service>vprn>dhcp>server>pool>subnet>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures specific DHCP options. The option defined in this context can be overruled if the same option is defined in the local user database.

The **no** form of the removes the option from the configuration.

Parameters

option-number

Specifies the option number that the DHCP server uses to send the identification strings to the DHCP client.

Values 1 to 254

address *ip-address*

Specifies the IP address of this host.

hex *hex-string*

Specifies the hex value of this option.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

string *ascii-string*

Specifies the value of this option up to 127 characters.

dns-server

Syntax

dns-server *ip-address* [*ip-address...*(up to 4 max)]

Context

config>service>vpn>dhcp>server>pool>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IP address of the DNS server.

Parameters

ip-address

Specifies the IP address of the DNS server. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

domain-name

Syntax

domain-name *domain-name*

no domain-name

Context

config>service>vpn>dhcp>server>pool>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the default domain for a DHCP client that the router uses to complete unqualified hostnames (without a dotted-decimal domain name).

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

Parameters

domain-name

Specifies the domain name for the client up to 127 characters.

lease-rebind-time

Syntax

lease-rebind-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no lease-rebind-time

Context

config>service>vpn>dhcp>server>pool>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the time the client transitions to a rebinding state.

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

Parameters

days

Specifies the number of days in the lease rebind time.

Values 0 to 3650

hours

Specifies the number of hours in the lease rebind time.

Values 0 to 23

minutes

Specifies the number of minutes in the lease rebind time.

Values 0 to 59

seconds

Specifies the number of seconds in the lease rebind time.

Values 0 to 59

lease-renew-time

Syntax

lease-renew-time [*days days*] [*hrs hours*] [*min minutes*] [**sec seconds**]

no lease-renew-time

Context

config>service>vpn>dhcp>server>pool>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the time the client transitions to a renew state.

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

Parameters

days

Specifies the number of days in the lease renew time.

Values 0 to 3650

hours

Specifies the number of hours in the lease renew time.

Values 0 to 23

minutes

Specifies the number of minutes in the lease renew time.

Values 0 to 59

seconds

Specifies the number of seconds in the lease renew time.

Values 0 to 59

lease-time

Syntax

lease-time [*days days*] [*hrs hours*] [*min minutes*] [**sec seconds**]

no lease-time

Context

config>service>vprn>dhcp>server>pool>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the amount of time that the DHCP server grants to the DHCP client to use a particular IP address.

The **no** form of this command removes the lease time parameters from the configuration.

Parameters

days

Specifies the number of days in the lease time.

Values 0 to 3650

hours

Specifies the number of hours in the lease time.

Values 0 to 23

minutes

Specifies the number of minutes in the lease time.

Values 0 to 59

seconds

Specifies the number of seconds in the lease time.

Values 0 to 59

subnet

Syntax

subnet {*ip-address/mask* | *ip-address netmask*} [**create**]

no subnet {*ip-address/mask* | *ip-address netmask*}

Context

config>service>vprn>dhcp>server>pool

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a subnet of IP addresses to be served from the pool. The subnet cannot include any addresses that were assigned to subscribers without those addresses specifically excluded. When the subnet is created, no IP addresses are made available until a range is defined.

Parameters

ip-address

Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d (no multicast address)

mask

Specifies the subnet mask in dotted-decimal notation.

Values 8 to 30

netmask

Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

Values a.b.c.d (any mask expressed as dotted quad)

create

Keyword used to create the entity. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

address-range

Syntax

[no] **address-range** *start-ip-address end-ip-address*

Context

config>service>vpn>dhcp>server>pool>subnet

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a range of IP addresses to be served from the pool. All IP addresses between the start and end IP addresses are included (other than specific excluded addresses).

Parameters

start-ip-address

Specifies the start address of this range to include. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

end-ip-address

Specifies the end address of this range to include. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

exclude-addresses

Syntax

[no] exclude-addresses *start-ip-address* [*end-ip-address*]

Context

config>service>vprn>dhcp>server>pool>subnet

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies a range of IP addresses that are excluded from the pool of IP addresses in this subnet.

Parameters

start-ip-address

Specifies the start address of this range to exclude. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

end-ip-address

Specifies the end address of this range to exclude. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

maximum-declined

Syntax

maximum-declined *maximum-declined*

no maximum-declined

Context

config>service>vprn>dhcp>server>pool>subnet

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the maximum number of declined addresses allowed.

Default

64

Parameters

maximum-declined

Specifies the maximum number of declined addresses allowed.

Values 0 to 4294967295

minimum-free

Syntax

minimum-free *minimum-free*

no minimum-free

Context

config>service>vprn>dhcp>server>pool>subnet

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the minimum number of free addresses in this subnet. If the actual number of free addresses in this subnet falls below this configured minimum, a notification is generated.

Default

1

Parameters

minimum-free

Specifies the minimum number of free addresses in this subnet.

Values 0 to 255

default-router

Syntax

default-router *ip-address* [*ip-address...*(up to 4 max)]

no default-router

Context

config>service>vpn>dhcp>server>pool>subnet

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IP address of the default router for a DHCP client. Up to four IP addresses can be specified.

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

Parameters

ip-address

Specifies the IP address of the default router. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

subnet-mask

Syntax

subnet-mask *ip-address*

no subnet-mask

Context

config>service>vpn>dhcp>server>pool>subnet

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the subnet-mask option to the client. The mask can either be defined (for supernetting) or taken from the pool address.

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

Parameters

ip-address

Specifies the IP address of the subnet mask. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

use-gi-address

Syntax

[no] **use-gi-address**

Context

config>service>vpn>dhcp>server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the use of gateway IP address (GIADDR) matching. If the **gi-address** flag is enabled, a pool can be used even if a subnet is not found. If the **local-user-db-name** is not used, the **gi-address** flag is used and addresses are handed out by GI only. If a user must be blocked from getting an address, the server maps to a local user database and configures the user with no address.

A pool can include multiple subnets. Because the GIADDR is shared by multiple subnets in a subscriber-interface, the pool may provide IP addresses from any of the subnets included when the GIADDR is matched to any of its subnets. This allows a pool to be created that represents a subnet.

Default

no use-gi-address

user-db

Syntax

user-db *local-user-db-name*

no user-db

Context

config>service>vpn>dhcp>server

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a local user database for authentication.

Default

no user-db

Parameters

local-user-db-name

Specifies the name of a local user database.

vrf-target

Syntax

vrf-target {*ext-community* | **export** *ext-community* | **import** *ext-community*}

no vrf-target

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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. 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-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

target:{*ip-address:comm-val* |*2byte-asnumber:ext-comm-val*|*4byte-asnumber:comm-val*}

ip-address: a.b.c.d

comm-val: 0 to 65535

2byte-asnumber: 1 to 65535

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.

8.4.2.1.4 IPOE commands

subscriber-mgmt

Syntax

subscriber-mgmt

Context

config

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure subscriber management entities. A subscriber is uniquely identified by a subscriber identification string. Each subscriber can have one DHCP session active at any time. Each session is referred to as a subscriber host and is identified by its IP address and MAC address.



Note:

On 7210 SAS-K 2F6C4T, the 7x50 subscriber management CLI is inherited to allow users to configure a DHCP address for specific hosts when using DHCP server for local craft access. 7210 SAS platforms do not support subscriber management. In all the commands under this section, in the context of 7210 SAS usage, DHCP subscriber refers to an IP host requesting an IP address.

local-user-db

Syntax

local-user-db *local-user-db-name* [create]

Context

config>subscr-mgmt

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a local user database.

Default

no local-user-db

Parameters

local-user-db-name

Specifies the name of a local user database.

create

Keyword to create the local user database.

ipoe

Syntax

ipoe

Context

config>subscr-mgmt>loc-user-db

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures IPoE host parameters.

mask

Syntax

mask type *dhcp-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}

no mask type *dhcp-match-type*

Context

config>subscr-mgmt>loc-user-db>ipoe

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the mask.

Parameters

dhcp-match-type

Specifies up to four matching types to identify a host.

Values DHCP: circuit-id, option60, remote-id, sap-id, string, system-id
PPP: circuit-id, remote-id, service-name, username

prefix-string prefix-string

Specifies a substring that is removed from the start of the incoming circuit ID before it is matched against the value configured in the DHCP or PPPOE circuit ID. This string can contain only printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum

prefix-length prefix-length

Specifies the number of characters to remove from the start of the incoming circuit ID before it is matched against the value configured in the DHCP circuit ID.

Values 1 to 127

suffix-string suffix-string

Specifies a substring that is removed from the end of the incoming circuit ID before it is matched against the value configured in DHCP circuit ID. This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum

suffix-length suffix-length

Specifies the number of characters to remove from the end of the incoming circuit ID before it is matched against the value configured in the DHCP circuit ID.

Values 1 to 127

host

Syntax

host *host-name* [create]

no host *host-name*

Context

config>subscr-mgmt>loc-user-db>ipoe

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command defines a DHCP subscriber.

Parameters

host-name

Specifies the host name.

create

Keyword used to create the host name. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

address

Syntax

address gi-address [**scope** *scope*]

address ip-address[/*prefix-length*]

address pool *pool-name* [**secondary-pool** *sec-pool-name*] [**delimiter** *delimiter*]

address use-pool-from-client [**delimiter** *delimiter*]

no address

Context

config>subscr-mgmt>loc-user-db>ipoe>host

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IP address definition for the host.

When the user database is from a local DHCP server, this command defines the IP address the server "offers" to the DHCP-client.

When the user database is used for PPPoE authentication, the **gi-address** parameter cannot be used. A fixed IP address then causes PPPoE to use this IP address. If no IP address is specified, the PPPoE looks for an IP address by other means (in this case, DHCP). If a pool name is specified, this pool is sent in the DHCP request so that it can be used by the DHCP server to determine which address to give to the host.

The **no** form of this command causes no IP address to be assigned to this host. In a user database referred to from a local DHCP server, creating a host without address information causes the matching client to never receive an IP address.

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

Default

no address

Parameters

gi-address

When specified, the **gi-address** of the DHCP message is taken to look for a subnet in the local DHCP server. The first available free address of the subnet is taken and "offered" to the host. When **local-user-db** is used for PPPoE authentication, this has the same result as **no address**.

ip-address

Specifies the fixed IP address to use for this host.

pool-name/sec-pool-name

Specifies the primary (and secondary) pool (in the local DHCP server) to use to search for an available address. The first available IP address from any subnet in the pool is used. When the local user database is used for PPPoE authentication, the specified pool name is sent to the DHCP server in a vendor-specific suboption under Option 82

use-pool-from-client

Specifies to use the pool-name in the Option 82 vendor-specific suboption.

delimiter delimiter

Specifies a single ASCII character as the delimiter to separate primary and secondary pool names in Option 82 VSO.

gi-address

Syntax

gi-address *ip-address*

no gi-address

Context

config>subscr-mgmt>loc-user-db>ipoe>host

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the selection of GI addresses based on the host entry in local use database.

The **gi-address** must be a valid address (associated with an interface) within the routing context that received the DHCP message on the access side.

Default

no gi-address

Parameters

ip-address

Specifies the IPv4 GI address.

auth-domain-name

Syntax

auth-domain-name *domain-name*

no auth-domain-name

Context

config>subscr-mgmt>loc-user-db>ipoe>host

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the domain name that is appended to a user name in RADIUS-authentication-request messages for the specific host.

Parameters

domain-name

Specifies the domain name to be appended to a user name in RADIUS-authentication-request message for the specific host.

host-identification

Syntax

host-identification

Context

config>subscr-mgmt>loc-user-db>ipoe>host

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure host identification parameters.

server

Syntax

server *ip-address*

no server

Context

```
config>subscr-mgmt>loc-user-db>ipoe>host
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IP address of the DHCP server to which to relay.

The **no** form of this command removes the DHCP server IP address from the configuration.

The configured DHCP server IP address must reference one of the addresses configured under the DHCP CLI context of an IES or VPRN subscriber.

Default

no server

Parameters

ip-address

Specifies the IP address of the DHCP server.

circuit-id

Syntax

circuit-id string *ascii-string*

circuit-id hex *hex-string*

no circuit-id

Context

```
config>subscr-mgmt>loc-user-db>ipoe>host>host-ident
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the circuit ID to match.

Parameters

ascii-string

Specifies the circuit ID from the Option 82.

hex-string

Specifies the circuit ID in hexadecimal format from the Option 82.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

encap-tag-range

Syntax

encap-tag-range *start-tag* *start-tag* **end-tag** *end-tag*

no encap-tag-range

Context

config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a range of encapsulation tags as the host identifications. The encapsulation tag is dot1q or qinq on an Ethernet port.

For dot1q, the start and end tag is single number, ranging from 0 to 4094; for QinQ, the start and end tag format is x.y, x or y could be "*", which indicates to ignore the inner or outer tag.

This command can be used only when **encap-tag-range** is configured as one of the match-list parameters.

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

Parameters

start-tag *start-tag*

Specifies the value of the start label in the range of SAPs allowed on this host.

Values

<i>start-tag</i>	dot1q	qtag1	
	qinq	(qtag1.qtag2 qtag1.* *.qtag2)	
		qtag1	0 to 4094
		qtag2	0 to 4094
		vpi	0 to 4095 (NNI) 0 to 255 (UNI)
		vci	1 to 65535

end-tag *end-tag*

Specifies the value of the end label in the range of SAPs allowed on this host.

Values

<i>end-tag</i>	dot1q	qtag1	
	qinq	(qtag1.qtag2 qtag1.* *.qtag2)	
		qtag1	0 to 4094

qtag2	0 to 4094
vpi	0 to 4095 (NNI)
	0 to 255 (UNI)
vci	1 to 65535

mac

Syntax

mac *ieee-address*

no mac

Context

config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the MAC address to match.

Parameters

ieee-address

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

option60

Syntax

option60 *hex-string*

no option60

Context

config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the Vendor-Identifying Vendor Option to match. Option 60 is encoded as Type-Length-Value (TLV). The *hex-string* portion of Option 60 in the received DHCP request is used for matching. Only the first 32 bytes can be defined using this command. If Option 60 included in the message is longer, those bytes are ignored.

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

Default

no option60

Parameters

hex-string

Specifies the hex value of this option.

Values 0x0 to 0xFFFFFFFF(maximum 254 hex nibbles)

remote-id

Syntax

remote-id *hex-string*

remote-id *ascii-string*

no remote-id

Context

config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the remote ID of the host.

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

Default

no remote-id

Parameters

remote-id

Specifies the remote ID.

sap-id

Syntax

sap-id *sap-id*

no sap-id

Context

config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the SAP ID from the Nokia vendor-specific suboption in Option 82 to match.

Parameters

sap-id

Specifies a SAP ID, up to 255 characters.

service-id

Syntax

service-id *service-id*

no service-id

Context

config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies an existing service ID from the Nokia vendor-specific suboption in Option 82 to match.

Parameters

service-id

Specifies an existing service ID.

Values 1 to 2147483647

string

Syntax

string *string*

no string

Context

config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the string from the Nokia vendor-specific suboption in Option 82 to match.

Parameters

string

Specifies the string, up to 255 characters.

system-id

Syntax

system-id *system-id*

no system-id

Context

config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the system ID from the Nokia vendor-specific suboption in Option 82 to match.

Parameters

system-id

Specifies the system ID, up to 255 characters.

options

Syntax

options

Context

config>subscr-mgmt>loc-user-db>ipoe>host

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure pool options. The options defined at the host level are overruled if the same option is defined at the local user database level.

custom-option

Syntax

custom-option *option-number* **address** [*ip-address...*(up to 4 max)]

custom-option *option-number* **hex** *hex-string*

custom-option *option-number* **string** *ascii-string*

no custom-option *option-number*

Context

config>subscr-mgmt>loc-user-db>ipoe>host>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures specific DHCP options. The options defined at the host level are overruled if the same option is defined at the local user database level.

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

Parameters

option-number

Specifies the option number that the DHCP server uses to send the identification strings to the DHCP client.

Values 1 to 254

address *ip-address*

Specifies the IP address of this host.

hex *hex-string*

Specifies the hex value of this option.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

string *ascii-string*

Specifies the value of this option, up to 127 characters.

default-router

Syntax

default-router *ip-address* [*ip-address...*(up to 4 max)]

no default-router

Context

config>subscr-mgmt>loc-user-db>ipoe>host>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IP address of the default router for a DHCP client. Up to four IP addresses can be specified.

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

Parameters

ip-address

Specifies the IP address of the default router. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

dns-server

Syntax

dns-server address [*ip-address...*(up to 4 max)]

no dns-server

Context

config>subscr-mgmt>loc-user-db>ipoe>host>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IP address of the DNS server.

Parameters

ipv6-address

Specifies the IPv4 address of the DNS server. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

domain-name

Syntax

domain-name *domain-name*

no domain-name

Context

config>subscr-mgmt>loc-user-db>ipoe>host>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the default domain for a DHCP client that the router uses to complete unqualified host names (without a dotted-decimal domain name).

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

Parameters

domain-name

Specifies the domain name for the client, up to 127 characters.

lease-rebind-time

Syntax

lease-rebind-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no lease-rebind-time

Context

config>subscr-mgmt>loc-user-db>ipoe>host>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the time the client transitions to a rebinding state.
The **no** form of this command removes the time from the configuration.

Parameters

days

Specifies the number of days in the lease rebind time.

Values 0 to 3650

hours

Specifies the number of hours in the lease rebind time.

Values 0 to 23

minutes

Specifies the number of minutes in the lease rebind time.

Values 0 to 59

seconds

Specifies the number of seconds in the lease rebind time.

Values 0 to 59

lease-renew-time

Syntax

lease-renew-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no lease-renew-time

Context

config>subscr-mgmt>loc-user-db>ipoe>host>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the time the client transitions to a renew state.
The **no** form of this command removes the time from the configuration.

Parameters

days

Specifies the number of days in the lease renew time.

Values 0 to 3650

hours

Specifies the number of hours in the lease renew time.

Values 0 to 23

minutes

Specifies the number of minutes in the lease renew time.

Values 0 to 59

seconds

Specifies the number of seconds in the lease renew time.

Values 0 to 59

lease-time

Syntax

lease-time [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

no lease-time

Context

config>subscr-mgmt>loc-user-db>ipoe>host>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the amount of time that the DHCP server grants to the DHCP client to use a particular IP address.

The **no** form of this command removes the lease time parameters from the configuration.

Parameters

days

Specifies the number of days in the lease time.

Values 0 to 3650

hours

Specifies the number of hours in the lease time.

Values 0 to 23

minutes

Specifies the number of minutes in the lease time.

Values 0 to 59

seconds

Specifies the number of seconds in the lease time.

Values 0 to 59

subnet-mask

Syntax

subnet-mask *ip-address*

no subnet-mask

Context

config>subscr-mgmt>loc-user-db>ipoe>host>options

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the subnet-mask option to the client. The mask can either be defined (for supernetting) or taken from the pool address.

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

Parameters

ip-address

Specifies the IP address of the subnet mask. This address must be unique within the subnet and specified in dotted-decimal notation.

Values a.b.c.d

mask

Syntax

mask type *dhcp-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}

no mask type *dhcp-match-type*

Context

config>subscr-mgmt>loc-user-db>ipoe

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the mask.

Parameters

dhcp-match-type

Specifies the data type that represents the type of matching done to identify a DHCP host.

Values circuit-id, option60, remote-id, sap-id, string, system-id

prefix-string prefix-string

Specifies a substring that is stripped of the start of the incoming circuit ID before it is matched against the value configured in the DHCP circuit ID. This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum, "*"

prefix-length prefix-length

Specifies the number of characters to remove from the start of the incoming circuit ID before it is matched against the value configured in the circuit ID.

Values 1 to 127

suffix-string suffix-string

Specifies a substring that is stripped of the end of the incoming circuit ID before it is matched against the value configured in circuit ID.

This string can contain only printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum

suffix-length suffix-length

Specifies the number of characters to remove from the end of the incoming circuit ID before it is matched against the value configured in the circuit ID.

Values 1 to 127

match-list

Syntax

match-list *match-type-1* [*match-type-2*...(up to 4 max)]

no match-list

Context

```
config>subscr-mgmt>loc-user-db>ipoe
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the type of matching done to identify a host.

Default

no match-list

Parameters

match-type-x

Specifies up to four matching types to identify a host.

- Values** circuit-id, mac, option60, remote-id, sap-id, service-id, string, system-id
- circuit-id** — Specifies the DHCP4 option (82,1) or DHCP6 option 18.
 - mac** — Specifies the MAC address of the client. Chaddr in DHCP4 and DUID in IPv6.
 - option60** — Specifies the DHCP4 option 60.
 - remote-id** — Specifies the DHCP4 option (82,2) or DHCP6 option 37
 - sap-id** — Specifies the SAP ID on which DHCPv4 packet are received. This parameter is inserted as ALU VSO (82,9,4) by the DHCPv4 relay in router. The parameter is inserted under the **vendor-specific-option** context of the DHCPv4 relay. Because the DHCP relay configuration is enabled under the **group-interface** context, the **group-interface** and **service-id** must be known before the **sap-id** can be used for local user database matching.
 - service-id** — Specifies the service ID of the ingress SAP for DHCPv4 packets. This parameter is inserted as ALU VSO (82,9,3) by the DHCPv4 relay in the router. This parameter is inserted under the **vendor-specific-option** context of the DHCPv4 relay.
 - string** — Specifies the custom string configured under the **vendor-specific-option** context of the DHCPv4 relay. The string is inserted as ALU VSO (82,9,5) by the DHCPv4 relay in router. Because the DHCP relay configuration is enabled under the **group-interface** context, the **group-interface** and **service-id** must be known before the string can be used for local user database matching.
 - system-id** — Specifies the system ID of the node name configured under the **system>name** context. This parameter is inserted as ALU VSO (82,9,1) by the DHCPv4 relay in the router. This parameter is inserted under the **vendor-specific-option** context of the DHCPv4 relay. Because the DHCP relay configuration is enabled under the **group-interface** context, the **group-interface** and **service-id** must

be known before the **system-id** can be used for local user database matching.

shutdown

Syntax

[no] shutdown

Context

config>subscr-mgmt>loc-user-db

config>subscr-mgmt>loc-user-db>ipoe>host

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

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

8.4.2.1.5 SDP commands

spoke-sdp

Syntax

[no] spoke-sdp *sdp-id*

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 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 service is created.

SDPs must be explicitly associated and bound to a service to allow far-end routers to participate in the service. Alternatively, **auto-bind** can be used. With **auto-bind**, no VPRN spoke-SDP configuration is required. When both **auto-bind** and **spoke-sdp** are configured, **spoke-sdp** takes precedence. The **spoke-sdp** configuration must be undone for **auto-bind** to take effect.

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

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. Allowed values are integers for existing SDPs.

Values 1 to 17407

vc-id

Specifies the virtual circuit identifier.

Values 1 to 4294967295

8.4.2.1.6 Interface commands

interface

Syntax

interface *ip-int-name*

no interface *ip-int-name*

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command creates a logical IP routing interface for a Virtual Private Routed Network (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 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 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 the interface and all the associated configuration. The interface must be administratively shut down before issuing the **no interface** command.

For VPRN services, the IP interface must be shut down 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. 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 will be changed to maintain that IP interface. If *ip-int-name* already exists within another service ID or is an IP interface defined within the **config router** commands, an error will occur and context will not be changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

Values 1 to 32

address

Syntax

address *ip-address/mask* | *ip-address netmask* } [**broadcast** [**all-ones** | **host-ones**]

no address

Context

```
config>service>vprn>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 in the **config>router>interface** 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 the Classless Inter-Domain Routing (CIDR) or traditional dotted-decimal notation. The show commands display CIDR notation and are stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

Use the **no** form of this command to remove the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Address	Admin state	Oper state
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable. 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 will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted-decimal notation. Allowed values are IP addresses.

Values 1.0.0.0 to 223.255.255.255 (with support of /31 subnets)
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.

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 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 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 *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. A mask of 255.255.255.255 is reserved for system IP addresses.

Values 128.0.0.0 to 255.255.255.254
a.b.c.d (no multicast/broadcast address)

broadcast

Keyword that overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) will be received by the IP interface.

Default host-ones

all-ones

Keyword to specify that the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

host-ones

Keyword to specify that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negation feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

allow-directed-broadcasts

Syntax

[no] allow-directed-broadcasts

Context

config>service>vprn>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the forwarding of directed broadcasts out of the IP interface.

A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The **allow-directed-broadcasts** command on an IP interface enables or disables the transmission of packets destined for the subnet broadcast address of the egress IP interface.

When enabled, a frame destined for the local subnet on this IP interface will be sent as a subnet broadcast out this interface. Care should be exercised when allowing directed broadcasts because it is a well-known mechanism used for denial-of-service attacks.

When disabled, directed broadcast packets discarded at this egress IP interface will be counted in the normal discard counters for the egress SAP.

By default, directed broadcasts are not allowed and will be discarded at this egress IP interface.

The **no** form of this command disables the forwarding of directed broadcasts out of the IP interface.

Default

no allow-directed-broadcasts

ip-mtu

Syntax

ip-mtu *octets*

no ip-mtu

Context

config>service>vprn>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IP maximum transmit unit (packet) for this interface.

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

Default

no ip-mtu

Parameters

octets

Specifies the MTU for this interface.

Values 512 to 9000

loopback

Syntax

[no] loopback

Context

config>service>vprn>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the associated interface as a loopback interface that has no associated physical interface. As a result, the associated interface cannot be bound to a SAP.

When using mtrace/mstat in a Layer 3 VPN context, the configuration for the VPRN should have a loopback address configured that has the same address as the core instance system address (BGP next hop).

local-proxy-arp

Syntax

[no] local-proxy-arp

Context

config>service>vprn>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and therefore becomes the forwarding point for all traffic between hosts in that subnet. When the **local-proxy-arp** command is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

Default

no local-proxy-arp

mac

Syntax

[no] mac *ieee-mac-address*

Context

config>service>vprn>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command assigns a specific MAC address to a VPRN IP interface. The default is the physical MAC address associated with the Ethernet interface on which the SAP is configured.

The **no** form of this command reverts the MAC address of the IP interface to the default value.

Parameters

ieee-mac-address

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

Allowed values are any non-broadcast, non-multicast MAC, and non-IEEE reserved MAC addresses.

proxy-arp

Syntax

[no] proxy-arp

Context

```
config>service>vprn>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables proxy ARP on the interface.

Default

```
no proxy-arp
```

proxy-arp-policy

Syntax

```
[no] proxy-arp-policy policy-name [policy-name...(up to 5 max)]
```

Context

```
config>service>vprn>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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>vprn>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

static-arp

Syntax

```
static-arp ip-address ieee-address
```

```
static-arp ieee-address unnumbered
```

```
no static-arp ip-address [ieee-address]
```

```
no static-arp [ieee-address] unnumbered
```

Context

```
config>service>vprn>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

Description

This command configures a static address resolution protocol (ARP) entry associating an IP address or an unnumbered address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can be configured only if it exists on the network attached to the IP interface.

If an entry for a particular IP address or unnumbered address already exists and a new MAC address is configured for the IP address, the existing MAC address will be replaced with the new MAC address.

The **no** form of this command removes a static ARP entry.

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. If this parameter is configured, it overrides any dynamic ARP.

unnumbered

Syntax

unnumbered *ip-int-name* | *ip-address*

no unnumbered

Context

config>service>vprn>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C.

Description

This command sets an IP interface as an unnumbered interface and specifies the IP address to be used for the interface.

To conserve IP addresses, unnumbered interfaces can be configured. The address used when generating packets on this interface is the *ip-address* parameter configured.

An error message is generated when an unnumbered interface is configured and an IP address already exists on this interface

The **no** form of this command removes the IP address from the interface, effectively removing the unnumbered property. The interface must be shut down before the **no unnumbered** command is issued to delete the IP address from the interface or an error message is generated.

Default

no unnumbered

Parameters

ip-int-name* | *ip-address

Specifies the IP interface name or IP address with which to associate the unnumbered IP interface in dotted-decimal notation. The configured IP address must exist on this node.

Nokia recommends using the system IP address, because it is not associated with a particular interface and is therefore always reachable. The system IP address is the default if *ip-int-name* or *ip-address* is not configured.

8.4.2.1.7 VPRN IPv6 interface commands

ipv6

Syntax

ipv6
no ipv6

Context

config>service>vprn>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables IPv6 for a VPRN interface.
The **no** form of this command disables IPv6.

Default

no ipv6

address

Syntax

ipv6 *ipv6-address/prefix-length* [**eui-64**] [**preferred**]
no ipv6 *ipv6-address/prefix-length*

Context

config>service>vprn>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command assigns an address to the IPv6 interface.
The **no** form of this command deletes the specified IPv6 address.

Parameters

ipv6-address/prefix-length
Specifies the interface IP address.

Values

ipv6-address: 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

Keyword to specify that a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from the MAC address on Ethernet interfaces.

preferred

Keyword to specify that the IPv6 address is the preferred IPv6 address for this interface. A preferred address is an address assigned to an interface whose use by upper layer protocols is unrestricted. A preferred address may be used as the source (or destination) address of packets sent from (or to) the interface.

icmp

Syntax

icmp

Context

config>service>vprn>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures ICMPv6 parameters for the interface.

packet-too-big

Syntax

packet-too-big *number seconds*
no packet-too-big

Context

config>service>vprn>if>ipv6>icmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether, and how often, ICMPv6 "packet-too-big" 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

no packet too big

Parameters

number

Specifies the number of ICMPv6 "packet-too-big" 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 ICMPv6 "packet-too-big" messages sent.

Values 1 to 60

Default 10

param-problem

Syntax

param-problem *number seconds*

no param-problem

Context

config>service>vprn>if>ipv6>icmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether, and how often, ICMPv6 "parameter-problem" messages should be sent. When enabled, ICMPv6 "parameter-problem" messages are generated by this interface.

The **no** form of this command disables the sending of ICMPv6 "parameter-problem" messages.

Default

no param-problem

Parameters

number

Specifies the number of ICMPv6 "parameter-problem" 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 ICMPv6 "parameter-problem" messages sent.

Values 1 to 60

Default 10

redirects

Syntax

redirects *number seconds*

no redirects

Context

config>service>vprn>if>ipv6>icmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 should be sent.

The **no** form of this command disables the sending of ICMPv6 redirect messages.

Default

no redirects

Parameters

number

Specifies the number of ICMPv6 "redirect" 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 ICMPv6 "redirect" messages sent.

Values 1 to 60

Default 10

time-exceeded

Syntax

time-exceeded *number seconds*

no time-exceeded

Context

config>service>vprn>if>ipv6>icmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the rate for ICMPv6 "time-exceeded" messages.

The **no** form of this command disables the sending of ICMPv6 "time-exceeded" messages.

Default

no time-exceeded

Parameters

number

Specifies the number of ICMPv6 "time-exceeded" 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 ICMPv6 "time-exceeded" messages sent.

Values 1 to 60

Default 10

unreachables

Syntax

unreachables *number seconds*

no unreachables

Context

config>service>vprn>if>ipv6>icmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables and configures the rate for ICMPv6 host and network destination "unreachables" messages issued on the router interface.

The **no** form of this command disables the generation of ICMPv6 destination "unreachables" messages on the router interface.

Default

no unreachables

Parameters

number

Specifies the number of ICMPv6 "unreachables" 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 ICMPv6 "unreachables" messages sent.

Values 1 to 60

Default 10

link-local-address

Syntax

ipv6 *ipv6-address* [**preferred**]

no ipv6 *ipv6-address/prefix-length*

Context

config>service>vpn>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command assigns the IPv6 link local address to the interface.

Parameters

ipv6-address

Specifies the IPv6 link local address.

Values

ipv6- x:x:x:x:x:x:x (eight 16-bit pieces)
address: x:x:x:x:x:d.d.d.d
x: [0 — FFFF]H
d: [0 — 255]D

preferred

Keyword to specify that the IPv6 link local address is the preferred IPv6 link local address for this interface. A preferred address is an address assigned to an interface whose use by upper layer protocols is unrestricted. A preferred addresses maybe used as the source (or destination) address of packets sent from (or to) the interface.

local-proxy-nd

Syntax

local-proxy-nd *number seconds*

no local-proxy-nd

Context

config>service>vpn>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables local proxy neighbor discovery on the interface.

The **no** form of this command disables local proxy neighbor discovery.

Default

no local-proxy-nd

neighbor

Syntax

neighbor *ipv6-address mac-address*

no neighbor *ipv6-address*

Context

config>service>vprn>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures IPv6-to-MAC address mapping on the interface.

The **no** form of this command deletes IPv6-to-MAC address mapping for the specified IPv6 address.

Parameters

ipv6-address

Specifies the IPv6 address.

Values

ipv6- x:x:x:x:x:x:x (eight 16-bit pieces)

address: x:x:x:x:x.d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

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 unicast MAC addresses and non-IEEE reserved MAC addresses

proxy-nd-policy

Syntax

proxy-nd-policy *policy-name* [*policy-name...*(up to 5 max)]

no proxy-nd-policy

Context

config>service>vprn>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command applies a proxy neighbor discovery policy for the interface.

The **no** form of this command disables the proxy neighbor discovery policy application.

Default

no proxy-nd-policy

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 (such as #, \$, spaces), the entire string must be enclosed within double quotes. The specified policy names must already be defined.

urpf-check

Syntax

urpf-check

Context

config>service>vprn>interface>ipv6

config>service>vprn>interface

Platforms

7210 SAS-K 3SFP+ 8Cs

Description

This command enables unicast RPF (uRPF) check on this interface.

The **no** form of this command disables unicast RPF (uRPF) check on this interface.

ignore-default

Syntax

[no] ignore-default

Context

```
config>service>vprn>urpf-check
config>router>urpf-check>ipv6
config>service>vprn>urpf-check>ipv6
```

Platforms

7210 SAS-K 3SFP+ 8C

Description

This command configures the uRPF check feature (if enabled) to ignore default routes for purposes of determining the validity of incoming packets.

The **no** form of this command considers the default route to be eligible when performing a uRPF check.

Default

no ignore-default

arp-timeout

Syntax

arp-timeout *seconds*
no arp-timeout

Context

```
config>service>vprn>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the minimum time in seconds an ARP entry learned on the IP interface will be stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host; otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

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

Default

14400 seconds

Parameters

seconds

Specifies the minimum number of seconds a learned ARP entry will be stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 to 65535

delayed-enable

Syntax

delayed-enable *seconds* [*init-only*]

no delayed-enable

Context

config>service>vprn>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command delays making an 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 (for example, after a reboot). To ensure that the state has time to be synchronized, the **delayed-enable** timer can be specified. The optional **init-only** parameter 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

Keyword that 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 or allowing MCS to sync the subscriber information. The delay occurs only immediately after a reboot.

vpls

Syntax

```
vpls service-name
```

Context

```
config>service
```

```
config>service>vprn>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command, within the IP interface context, binds 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 the **allow-ip-int-binding** flag is set, the IP interface will be attached to the VPLS service, which allows 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 will be ignored and will not be attached to the IP interface.

If the service name is applied to a VPLS service after the service name is bound to an IP interface and the VPLS service **allow-ip-int-binding** flag is set at the time the name is applied, the VPLS service is automatically resolved to the IP interface if the interface is administratively up or when the interface is placed in the administratively up state.

If the service name is applied to a VPLS service without the **allow-ip-int-binding** flag set, the system does not attempt to resolve the applied service name to an existing IP interface bound to the name. To rectify this condition, the flag must first be set and then the IP interface must enter or reenter the administratively up state.

While the specified service name may be assigned to only one service context in the system, it is possible to bind the same service name to more than one IP interface. If two or more IP interfaces are bound to the same service name, the first IP interface to enter the administratively up state (if currently administratively down) or to reenter the administratively up state (if currently administratively up) when a VPLS service is configured with the name and has the **allow-ip-int-binding** flag set will be attached to the VPLS service. Only one IP interface is allowed to attach to a VPLS service context. No error is generated for the remaining non-attached IP interfaces using the service name.

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 will attempt 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.

ingress

Syntax

ingress

Context

```
config>service>vprn>if>vpls
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command within the VPLS binding context defines the routed IP filter ID optional filter overrides.

v4-routed-override-filter

Syntax

v4-routed-override-filter *ip-filter-id*
no v4-routed-override-filter

Context

```
config>service>ies>if>vprn>ingress
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies an IP filter ID that is applied to all ingress packets entering the VPLS service. The filter overrides the existing ingress IP filter applied to SAPs or SDP bindings for packets associated with the routing IP interface. The override filter is optional and, if not defined or removed, the IP routed packets use the existing ingress IP filter on the VPLS virtual port.

The **no** form of this command removes 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 packet virtual port, when defined.

Parameters

ip-filter-id

Specifies the ID for the IP filter policy. Allowed values are integers that correspond to a previously created IP filter policy in the **configure>filter>ip-filter** context.

Values 1 to 65535

8.4.2.1.8 IPv6 router advertisement commands

router-advertisement

Syntax

[no] router-advertisement

Context

config>service>vprn

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 interface. However, the **no interface interface-name** command disables a specific interface.

Default

disabled

interface

Syntax

[no] interface *ip-int-name*

Context

config>service>vprn>router-advertisement

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures router advertisement properties on a specific interface. The interface must already exist in the **config>router>interface** context.

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>vprn>router-advert>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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. A value of zero means there is an unspecified number of hops.

Values 0 to 255

managed-configuration

Syntax

[no] managed-configuration

Context

```
config>service>vprn>router-advert>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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>vprn>router-advert>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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>vprn>router-advert>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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>vprn>router-advert>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the MTU for the nodes to use when sending packets on the link.

Default

no mtu

Parameters

mtu-bytes

Specifies the MTU for the nodes to use when sending packets on the link.

Values 1280 to 9212

other-stateful-configuration

Syntax

[no] other-stateful-configuration

Context

config>service>vprn>router-advert>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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>vprn>router-advert>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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/prefix-length

Specifies the IP prefix.

Values

<i>ipv4-prefix:</i>	a.b.c.d (host bits must be 0)
<i>ipv4-prefix-length:</i>	0 to 32
<i>ipv6-prefix:</i>	x:x:x:x:x:x:x (eight 16-bit pieces)

x:x:x:x:x:d.d.d

x: [0 to FFFF]H

d: [0 to 255]D

ipv6-prefix-length: 0 to 128

autonomous

Syntax

[no] autonomous

Context

config>service>vprn>router-advert>if>prefix

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether the prefix can be used for stateless address autoconfiguration.

Default

autonomous

on-link

Syntax

[no] on-link

Context

config>service>vprn>router-advert>if>prefix

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether the prefix can be used for on-link determination.

Default

on-link

preferred-lifetime

Syntax

[no] preferred-lifetime {seconds | infinite}

Context

config>service>vprn>router-advert>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the time that this prefix will continue to be preferred. 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 length of time, in seconds, that this prefix will be preferred.

infinite

Specifies that the prefix is always preferred. The highest value represents infinity.

Values 1 to 4294967294

valid-lifetime

Syntax

valid-lifetime {seconds | infinite}

Context

config>service>vprn>router-advert>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the time 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 time, in seconds, that this prefix will be valid.

infinite

Specifies that the prefix will always be valid. The highest value represents infinity.

Values 1 to 4294967294

reachable-time

Syntax

reachable-time *milli-seconds*

no reachable-time

Context

config>service>vprn>router-advert

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 the router should be considered reachable.

Values 0 to 3600000

retransmit-time

Syntax

retransmit-timer *milli-seconds*

no retransmit-timer

Context

config>service>vprn>router-advert

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the retransmission frequency of neighbor solicitation messages.

Default

no retransmit-time

Parameters

milli-seconds

Specifies how often the retransmission should occur.

Values 0 to 1800000

router-lifetime

Syntax

router-lifetime *seconds*

no router-lifetime

Context

config>service>vprn>router-advert

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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. 0 means that the router is not a default router on this link.

Values 0, 4 to 9000

8.4.2.1.9 DHCP server commands

dhcp

Syntax

dhcp

Context

config>service>vprn>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure DHCP parameters.

action

Syntax

action {replace | drop | keep}

no action

Context

config>service>vprn>if>dhcp>option

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the processing required when the router receives a DHCP request that already has a Relay Agent Information Option (Option 82) field in the packet.

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

In accordance with RFC 3046, *DHCP Relay Agent Information Option*, section 2.1.1, *Reforwarded DHCP Requests*, the default is to keep the existing information intact. The exception to this is if the GIADDR of the received packet is the same as the ingress address on the router. In this case, the packet is dropped and an error is logged.

Parameters

replace

Keyword that specifies that, in the upstream direction (from the user), the existing Option 82 field is replaced with the Option 82 field from the router. In the downstream direction (toward the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop

Keyword that specifies the packet is dropped, and an error is logged.

keep

Keyword that specifies 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 sent on 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 occurs only when the incoming message already has 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>vprn>if>dhcp>option

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

When enabled, the router sends the interface index (If Index) in the **circuit-id** suboption of the DHCP packet. Display the If Index of a router interface using the **show router interface detail** command. This option specifies data that must be unique to the router that is relaying the circuit.

If disabled, the **circuit-id** suboption of the DHCP packet is left empty.

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

Default

circuit-id

Parameters

ascii-tuple

Keyword to specify that the ASCII-encoded concatenated tuple will be used, which consists of the *access-node-identifier*, *service-id*, and *interface-name*, separated by "|".

ifindex

Keyword to specify that the interface index is used. Display the If Index of a router interface using the **show router interface detail** command.

sap-id

Keyword to specify that the SAP ID is used.

vlan-ascii-tuple

Keyword to specify that the format included VLAN ID and dot1p bits in addition to what is already included in the ascii-tuple. The format is supported on dot1q and qinq ports only. Therefore, when the Option 82 bits are stripped, dot1p bits are copied to the Ethernet header of an outgoing packet.

gi-address

Syntax

gi-address *ip-address* [*src-ip-addr*]

no gi-address

Context

config>service>vprn>if>dhcp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the gateway IP address (GIADDR) for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. The GIADDR is needed when the router functions as a DHCP relay to distinguish between the different subscriber interfaces and potentially between the group interfaces defined.

Default

no gi-address

Parameters

ip-address

Specifies the host IP address to be used for DHCP relay packets.

src-ip-address

Specifies the source IP address to be used for DHCP relay packets.

option

Syntax

[no] **option**

Context

```
config>service>vprn>if>dhcp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables DHCP Option 82 (Relay Agent Information Option) parameter processing and enters the context for configuring Option 82 suboptions.

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>vprn>sub-if>grp-if>dhcp>option
```

```
config>service>vprn>nw-if>dhcp>option
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

When enabled, the router sends the MAC address of the remote end (typically the DHCP client) in the **remote-id** suboption of the DHCP packet. This command identifies the host at the other end of the circuit. If disabled, the **remote-id** suboption of the DHCP packet will be left empty.

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

Default

remote-id

Parameters

mac

Keyword that specifies the MAC address of the remote end is encoded in the suboption.

string *string*

Specifies the remote ID.

vendor-specific-option

Syntax

[no] vendor-specific-option

Context

config>service>vpn>if>dhcp>option

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the Nokia vendor-specific suboption of the DHCP relay packet.

client-mac-address

Syntax

[no] client-mac-address

Context

config>service>vpn>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the sending of the MAC address in the vendor-specific suboption of the DHCP relay packet.

The **no** form of this command disables the sending of the MAC address in the vendor-specific suboption of the DHCP relay packet.

pool-name

Syntax

[no] pool-name

Context

config>service>vpn>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command sends the pool name in the vendor-specific suboption of the DHCP relay packet.

The **no** form of this command disables the sending.

sap-id

Syntax

[no] **sap-id**

Context

config>service>vprn>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the sending of 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>vprn>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the sending of 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>vprn>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the vendor-specific suboption string of the DHCP relay packet.

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

Parameters

text

Specifies a string of 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>vprn>if>dhcp>option>vendor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether the system ID is encoded in the vendor-specific suboption of Option 82.

server

Syntax

server *server1* [*server2...*(up to 8 max)]

Context

config>service>vprn>if>dhcp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies a list of servers where requests will be forwarded. The list of servers can be entered as IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers, the request is forwarded to all the servers in the list. There can be a maximum of 8 DHCP servers configured.

The flood command is applicable only in the VPLS case. There is a scenario with VPLS where the VPLS node only needs to add Option 82 information to the DHCP request to provider per-subscriber information, but it does not do full DHCP relay. In this case, the server is set to "flood". This means the DHCP request is still a broadcast and is sent through the VPLS domain. A node running at L3 further upstream can then perform the full L3 DHCP relay function.

Default

no server

Parameters

server

Specifies the DHCP server IP address.

trusted

Syntax

[no] trusted

Context

config>service>vprn>if>dhcp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command, if enabled on an IP interface, specifies that the relay agent (the SR-Series) will modify the request GIADDR to be equal to the ingress interface and forward the request. According to RFC 3046, *DHCP Relay Agent Information Option*, a DHCP request where the GIADDR is 0.0.0.0 and that contains an Option 82 field in the packet, should be discarded, unless it arrives on a "trusted" circuit.

This behavior applies only when the action in the Relay Agent Information Option is "keep". In the case where the Option 82 field is being replaced by the relay agent (action = "replace"), the original Option 82 information is lost, therefore there is no reason to enable the **trusted** command.

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

Default

not enabled

8.4.2.1.10 Interface ICMP commands

icmp

Syntax

icmp

Context

config>service>vprn>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables responses to ICMP mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance 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>vprn>if>icmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the rate for 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 are 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

tll-expired

Syntax

tll-expired *number seconds*

no ttl-expired

Context

```
config>service>vprn>if>icmp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the rate 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 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
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

This command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables are issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of destination unreachable messages that can be issued on the interface for a specific time interval.

By default, generation of ICMP destination unreachable 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 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

8.4.2.1.11 Interface SAP commands

```
sap
```

Syntax

```
sap sap-id [create]
```

```
no sap sap-id
```

Context

```
config>service>vprn>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters that identify the service access point on the interface and within the 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 be associated with only a single service. A SAP can be defined only on a port that has been configured as an access port using the **config port *port-id* ethernet mode access** command.

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 **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.

Special Cases

VPRN

A VPRN SAP must be defined on an Ethernet interface.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

port-id

Specifies the physical port ID in the *slot/mda/port* format.

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the *slot_number/MDA_number/port_number* format. For example *2/3* specifies port 3 on MDA 2 in slot.

The *port-id* must reference a valid port type. The port must be configured as an access port.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 a 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command creates the accounting policy 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.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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, by default the data is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued, the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued, the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default

no collect-stats

8.4.2.1.12 Interface SAP filter and QoS policy commands

egress

Syntax

egress

Context

config>service>vprn>if>sap

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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>vprn>if>sap

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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-shaper-rate

Syntax

agg-shaper-rate *agg-rate*

no agg-shaper-rate

Context

config>service>vprn>sap>ingress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and BUM traffic across all the FCs of the SAP configured to use ingress queues. Users can specify the CIR rate and the PIR rate. Users must not oversubscribe the total bandwidth available for use by ingress queues.

The **no** form of this command disables the SAP aggregate rate shaper. That is, the SAP can use up the maximum bandwidth available.

Default

no agg-shaper-rate

Parameters

agg-rate

Specifies the rate in kilobits per second.

Values 50 to 3000000 | max (7210 SAS-K 2F6C4T)
64 to 20000000 | max (7210 SAS-K 3SFP+ 8C)

Default max

agg-shaper-rate

Syntax

agg-shaper-rate *agg-rate*

no agg-shaper-rate

Context

config>service>vprn>sap>egress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the aggregate rate for the SAP shaper. The aggregate SAP shaper is available to limit only the unicast traffic and BUM traffic across all the FCs of the SAP configured to use ingress queues. Users can specify the CIR rate and the PIR rate. Users must not oversubscribe the total bandwidth available for use by ingress queues.

The **no** form of this command disables the SAP aggregate rate shaper. That is, the SAP can use up the maximum bandwidth available.

Default

no agg-shaper-rate

Parameters

agg-rate

Specifies the rate in kilobits per second.

Values 50 to 1000000 | max (7210 SAS-K 2F6C4T)
64 to 10000000 | max (7210 SAS-K 3SFP+ 8C)

Default max

filter

Syntax

filter ip *ip-filter-id*

filter [mac *mac-filter-id*]

no filter [ip *ip-filter-id*]

no filter [mac *mac-filter-id*]

no filter

Context

config>service>vprn>if>sap>egress

config>service>vprn>if>sap>ingress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

This command associates a filter policy with a specified *ip-filter-id* with an ingress or egress SAP. The *ip-filter-id* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation 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 that non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes configured filter ID associations 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 the IP filter policy. The filter ID must already exist within the created IP 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*

no qos

Context

config>service>vprn>if>sap>egress

config>service>vprn>if>sap>ingress

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Associates a Quality of Service (QoS) policy with an ingress or egress 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.

This 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 egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error. Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a specific type returns an error.

When an ingress QoS policy is defined on an ingress IP interface that is bound to a VPRN, the policy becomes associated with every SAP on the VPRN and augments the QoS policy that is defined on each SAP. Packets that are bridged are processed using the policy defined on the VPRN SAP; packets that are routed are processed using the policy defined in the IES IP interface-binding context.

When an egress QoS policy is associated with an IP interface that has been bound to a VPRN, the policy becomes associated with every SAP on the VPRN and augments the egress QoS policy that is defined on each SAP. Packets that are bridged are processed using the policy defined on the VPRN SAP; packets that are routed are processed using the policy defined in the IP interface-binding context.

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

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

Parameters

policy-id

Specifies the ingress or egress policy ID to associate with the SAP or IP interface on ingress or egress. The policy ID must already exist.

Values 1 to 65535

8.4.2.1.13 Interface VRRP commands

```
vrrp
```

Syntax

```
vrrp virtual-router-id [owner]  
no vrrp virtual-router-id
```

Context

```
config>service>vprn>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

The VRID can be defined as owner or non-owner.

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.

owner

Specifies the virtual router instance as an owner.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

The command 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 will be used immediately.

If a **no authentication-key** command is executed, the password authentication key reverts to the default value. The **authentication-key** command may be executed at any time.

To change the current in-use password key on multiple virtual router instances:

- identify the current master
- shut down the virtual router instance on all backups
- execute the **authentication-key** command on the master to change the password key
- execute the **authentication-key** command and **no shutdown** command on each backup

The **no** form of this command reverts the default null string to the value of the key.

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 of eight octets 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.

This parameter is expressed as a string consisting of up to eight alphanumeric characters. Spaces must be contained in quotation marks (" "). The quotation marks are not considered part of the string.

The string is case-sensitive and is left-justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet (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

Specifies the hash key. The key can be any combination of ASCII characters up to 22 characters (encrypted) for a hash key or up to 121 characters for a hash2 key. If spaces are used in the string, the entire string must be enclosed in quotation marks (" ").

This option is useful when a user must configure the parameter, but for security purposes, the actual unencrypted key value is not provided.

hash

Specifies that the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2

Specifies that the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less-encrypted **hash** form is assumed.

backup

Syntax

[no] **backup** *ip-address*

Context

config>service>vprn>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a virtual router IP address for the backup interface.

Default

no backup

Parameters

ip-address

Specifies the destination IPv4 address for the backup interface.

Values a.b.c.d

bfd-enable

Syntax

bfd-enable *service-id* **interface** *interface-name* **dst-ip** *ip-address*

no bfd-enable *service-id* **if** *if-name* **dst-ip** *ip-address*

bfd-enable **interface** *interface-name* **dst-ip** *ip-address*

no bfd-enable interface *interface-name* **dst-ip** *ip-address*

Context

```
config>service>vprn>if>vrrp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 using the **bfd-enable** command under the IP interface specified in this command.

The **no** form of this command removes BFD from the associated BGP protocol peering.

Default

no bfd-enable

Parameters

service-id

Specifies the service ID of the interface running BFD.

Values *service-id*: 1 to 2147483647

svc-name: Specifies an existing service name of up to 64 characters.

No service ID indicates a network interface.

interface-name

Specifies the name of the interface running BFD, up to 32 characters.

ip-address

Specifies the destination address to be used for the BFD session.

init-delay

Syntax

init-delay *seconds*

no init-delay

Context

```
config>service>vprn>if>vrrp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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>vprn>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command allows the virtual router instance to inherit the master VRRP router advertisement interval timer, which is used by backup routers to calculate the master down timer.

This command is available only in the non-owner nodal context and allows the current virtual router instance master to dictate the master down timer for all backup virtual routers. The **master-int-inherit** command has no effect when the virtual router instance is operating as master.

If the **master-int-inherit** command is not enabled, the locally configured **message-interval** must match the master VRRP advertisement message advertisement interval field value or the message is discarded.

The **no** form of this command reverts to the default operating condition, which requires the locally configured **message-interval** to match the received VRRP advertisement message advertisement interval field value.

Default

no master-int-inherit

message-interval

Syntax

message-interval {[seconds] [milliseconds milliseconds]}

no message-interval

Context

```
config>service>vprn>if>vrrp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the advertisement timer and indirectly sets the master down timer on the virtual router instance. The **message-interval** values must be the same for all virtual routers with the same VRID. Any VRRP advertisement message received with an Advertisement Interval field different from the virtual router instance configured **message-interval** value is silently discarded.

The **message-interval** command is available for both non-owner and owner virtual router 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 reverts to the default message interval value of 1 second.

Parameters

seconds

Specifies the time interval, in seconds, between sending advertisement messages.

Values 1 to 255

Default 1

milliseconds

Specifies the time interval, in milliseconds, between sending advertisement messages.

Values 100 to 900

ping-reply

Syntax

```
[no] ping-reply
```

Context

```
config>service>vprn>if>vrrp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the non-owner master to reply to ICMP echo requests directed at the IP address of the virtual router instance. The ping request can be received on any routed interface.

Ping must not have been disabled at the management security level (either on the parent IP interface or based on the ping source host address). When **ping-reply** is not enabled, ICMP echo requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to ICMP echo requests regardless of the **ping-reply** configuration.

The **ping-reply** command is available only for non-owner virtual routers.

The **no** form of this command reverts to the default operation of discarding all ICMP echo request messages destined for the non-owner virtual router instance IP address.

Default

no ping-reply

policy

Syntax

policy *vrrp-policy-id*

no policy

Context

config>service>vprn>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command associates a VRRP priority control policy with the virtual router instance (non-owner context only).

This command creates the context to configure a VRRP priority control policy that controls the VRRP in-use priority based on priority control events. It is a parent node for the various VRRP priority control policy commands that define the policy parameters and priority event conditions.

This command defines the initial or base value used by non-owner virtual routers. This value can be modified by assigning a VRRP priority control policy to the virtual router instance. The VRRP priority control policy can override or diminish the base priority setting to establish the actual in-use priority of the virtual router instance.

The **policy** *policy-id* command must be created first, before it can be associated with a virtual router instance.

Because VRRP priority control policies define conditions and events that must be maintained, they can be resource intensive. The number of policies is limited to 1000.

The *policy-id* does not have to be comprised consecutive integers.

The **no** form of this command deletes the specific *policy-id* from the system.

The *policy-id* must be removed first from all virtual router instances before the **no policy** command can be issued. If the *policy-id* is associated with a virtual router instance, the command fails.

Parameters

vrrp-policy-id

Specifies a VRRP priority control policy.

Values 1 to 9999

preempt

Syntax

[no] preempt

Context

config>service>vprn>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command overrides an existing non-owner master with a virtual router backup that has a higher priority. Enabling preempt mode is recommended for correct operation of the base-priority definitions on the virtual router instance. If the virtual router cannot preempt an existing non-owner master, the effect of the dynamic changing of the in-use priority is greatly diminished.

The **preempt** command is available for only non-owner VRRP virtual routers. The owner cannot be preempted because the priority of non-owners can never be higher than the owner. The owner always preempts 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 backup virtual router instance in-use priority.

A master non-owner virtual router allows itself to be preempted only when the incoming VRRP advertisement message priority field value is one of the following:

- greater than its in-use priority value
- equal to its in-use priority value, and the source IP address (primary IP address) is greater than its primary IP address

The **no** form of this command prevents a non-owner virtual router instance from preempting another, less desirable, virtual router.

Default

preempt

priority

Syntax

priority *priority*

no priority

Context

```
config>service>vprn>if>vrrp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a specific priority value for the virtual router instance. In conjunction with the optional **policy** command, the base priority is used to derive the in-use priority of the virtual router instance.

The **priority** command is available only for non-owner VRRP virtual routers. The priority of owner virtual router instance 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 resets the base priority to 100.

Parameters

priority

Specifies the base priority used by the virtual router instance. If a VRRP priority control policy is not defined, the base priority will be the in-use priority for the virtual router instance.

Values 1 to 254

Default 100

ssh-reply

Syntax

```
[no] ssh-reply
```

Context

```
config>service>vprn>if>vrrp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the non-owner master to reply to SSH requests directed at the IP address of the virtual router instance. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parent IP interface or based on the SSH source host address). Correct login and CLI command authentication are 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 available only for non-owner VRRP virtual routers.

The **no** form of this command reverts to the default operation of discarding all SSH packets destined for the non-owner virtual router instance IP address.

Default

no ssh-reply

standby-forwarding

Syntax

[no] standby-forwarding

Context

config>service>vprn>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the forwarding of packets by a standby router to the virtual router MAC address.

The **no** form of this command specifies that a standby router will not forward traffic sent to the virtual router MAC address but will forward traffic sent to the real MAC address of the standby router.

Default

no standby-forwarding

telnet-reply

Syntax

[no] telnet-reply

Context

config>service>vprn>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the non-owner master to reply to TCP port 23 Telnet requests directed at the IP address of the virtual router instance. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parent IP interface or based on the Telnet source host address). Correct login and CLI command authentication are 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 available only in non-owner VRRP virtual routers.

The **no** form of this command reverts to the default operation of discarding all Telnet packets destined for the non-owner virtual router instance IP address.

Default

no telnet-reply

traceroute-reply

Syntax

[no] **traceroute-reply**

Context

config>service>vprn>if>vrrp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables a non-owner master to reply to traceroute requests directed to the virtual router instance IP address.

This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.

A non-owner backup virtual router never responds to traceroute requests regardless of the **traceroute-reply** configuration.

Default

no traceroute-reply

8.4.2.2 VPRN interface IPv6 commands

ipv6

Syntax

ipv6
no ipv6

Context

config>service>vprn>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables IPv6 for a VPRN interface.
The **no** form of this command disables IPv6.

Default

no ipv6

address

Syntax

ipv6 *ipv6-address/prefix-length* [**eui-64**] [**preferred**]
no ipv6 *ipv6-address/prefix-length*

Context

config>service>vprn>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command assigns an address to the IPv6 interface.
The **no** form of this command deletes the specified IPv6 address.

Parameters

ipv6-address/prefix-length

Specifies the interface IP address.

Values

ipv6-address: 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

prefix-length: 1 to 128

eui-64

Keyword to specify that a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from the MAC address on Ethernet interfaces.

preferred

Keyword to specify that the IPv6 address is the preferred IPv6 address for this interface. A preferred address is an address assigned to an interface whose use by upper layer protocols is unrestricted. A preferred address maybe used as the source (or destination) address of packets sent from (or to) the interface.

icmp

Syntax

icmp

Context

config>service>vprn>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures ICMPv6 parameters for the interface.

packet-too-big

Syntax

packet-too-big *number seconds*

no packet-too-big

Context

config>service>vprn>if>ipv6>icmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether, and how often, ICMPv6 "packet-too-big" 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

no packet too big

Parameters

number

Specifies the number of ICMPv6 "packet-too-big" 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 ICMPv6 "packet-too-big" messages sent.

Values 1 to 60

Default 10

param-problem

Syntax

param-problem *number seconds*

no param-problem

Context

config>service>vprn>if>ipv6>icmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether, and how often, ICMPv6 "parameter-problem" messages should be sent. When enabled, ICMPv6 "parameter-problem" messages are generated by this interface.

The **no** form of this command disables the sending of ICMPv6 "parameter-problem" messages.

Default

no param-problem

Parameters

number

Specifies the number of ICMPv6 "parameter-problem" 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 ICMPv6 "parameter-problem" messages sent.

Values 1 to 60

Default 10

redirects

Syntax

redirects *number seconds*

no redirects

Context

config>service>vprn>if>ipv6>icmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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. The messages alert the node that a better route is available.

The **no** form of this command disables the sending of ICMPv6 redirect messages.

Default

no redirects

Parameters

number

Specifies the number of ICMPv6 "redirect" 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 ICMPv6 "redirect" messages sent.

Values 1 to 60

Default 10

time-exceeded

Syntax

time-exceeded *number seconds*

no time-exceeded

Context

config>service>vprn>if>ipv6>icmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the rate for ICMPv6 "time-exceeded" messages.

The **no** form of this command disables the sending of ICMPv6 "time-exceeded" messages.

Default

no time-exceeded

Parameters

number

Specifies the number of ICMPv6 "time-exceeded" 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 ICMPv6 "time-exceeded" messages sent.

Values 1 to 60

Default 10

unreachables

Syntax

unreachables *number seconds*

no unreachables

Context

config>service>vprn>if>ipv6>icmp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the rate for ICMPv6 host and network destination "unreachables" messages issued on the router interface.

The **no** form of this command disables the generation of ICMPv6 destination "unreachables" messages on the router interface.

Default

no unreachables

Parameters

number

Specifies the number of ICMPv6 "unreachables" 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 ICMPv6 "unreachables" messages sent.

Values 1 to 60

Default 10

link-local-address

Syntax

ipv6 *ipv6-address* [**preferred**]

no ipv6 *ipv6-address/prefix-length*

Context

```
config>service>vprn>if>ipv6
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command assigns the IPv6 link local address to the interface.

Parameters

ipv6-address

Specifies the IPv6 link local address.

Values

ipv6-address: x::x:x:x:x:x (eight 16-bit pieces)
x::x:x:x:x.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D

preferred

Keyword to specify that the IPv6 link local address is the preferred IPv6 link local address for this interface. A preferred address is an address assigned to an interface whose use by upper layer protocols is unrestricted. A preferred addresses maybe used as the source (or destination) address of packets sent from (or to) the interface.

local-proxy-nd

Syntax

```
local-proxy-nd number seconds  
no local-proxy-nd
```

Context

```
config>service>vprn>if>ipv6
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables local proxy neighbor discovery on the interface.
The **no** form of this command disables local proxy neighbor discovery.

Default

```
no local-proxy-nd
```


neighbor

Syntax

neighbor *ipv6-address mac-address*

no neighbor *ipv6-address*

Context

config>service>vprn>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures IPv6-to-MAC address mapping on the interface.

The **no** form of this command deletes IPv6-to-MAC address mapping for the specified IPv6 address.

Parameters

ipv6-address

Specifies the IPv6 address.

Values

ipv6- x:x:x:x:x:x (eight 16-bit pieces)
address: x:x:x:x:x:d.d.d.d
x: [0 to FFFF]H
d: [0 to 255]D

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 unicast MAC addresses and non-IEEE reserved MAC addresses.

proxy-nd-policy

Syntax

proxy-nd-policy *policy-name [policy-name...(up to 5 max)]*

no proxy-nd-policy

Context

config>service>vprn>if>ipv6

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command applies a proxy neighbor discovery policy for the interface.

The **no** form of this command disables the proxy neighbor discovery policy application.

Default

no proxy-nd-policy

Parameters

policy-name

Specifies an existing neighbor discovery policy name. Allowed values are any string up to 32 characters comprised of printable, 7-bit ASCII characters. If the string contains special characters (such as #, \$, spaces), the entire string must be enclosed within double quotes. The specified policy names must already be defined.

8.4.2.3 BGP commands

```
bgp
```

Syntax

```
[no] bgp
```

Context

```
config>service>vprn
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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; 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 in which 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

best-path-selection

Syntax

best-path-selection

Context

config>service>vprn>bgp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables path selection configuration.

always-compare-med

Syntax

always-compare-med {zero | infinity}

no always-compare-med

Context

config>service>vprn>bgp>best-path-selection

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures 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 command determines what MED value is inserted in the RIB-IN. If this command 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, a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.

infinity

Specifies for routes learned without a MED attribute that a value of infinity (4294967295) is used in the MED comparison. This, in effect, makes these routes the least desirable.

as-path-ignore

Syntax

```
[no] as-path-ignore
```

Context

```
config>service>vprn>bgp
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures whether the AS path is used to determine the best BGP route.

If this command is configured, 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
```

ignore-nh-metric

Syntax

```
ignore-nh-metric
```

```
no ignore-nh-metric
```

Context

```
config>service>vprn
```

```
config>service>vprn>bgp>best-path-selection
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures BGP to disregard the resolved distance to the BGP next hop in its decision process for selecting the best route to a destination.

When configured in the **config>router>bgp>bestpath-selection** context, this command applies to the comparison of two BGP routes with the same NLRI learned from base router BGP peers. When configured in the **config>service>vprn** context, this command applies to the comparison of two BGP-VPN routes for the same IP prefix imported into the VPRN from the base router BGP instance. When configured in the **config>service>vprn>bgp>best-path-selection** context, this command applies to the comparison of two BGP routes for the same IP prefix learned from VPRN BGP peers.

The **no** form of this command reverts to the default behavior, whereby BGP factors distance to the next hop into its decision process.

Default

no ignore-nh-metric

ignore-router-id

Syntax

ignore-router-id

no ignore-router-id

Context

config>service>vprn>bgp>best-path-selection

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

When this command is present and the current best path to a destination was learned from eBGP peer X with BGP identifier x and a new path is received from eBGP peer Y with BGP identifier y, the best path remains unchanged if the new path is equivalent to the current best path up to the BGP identifier comparison, even if y is less than x.

The **no** form of this command reverts to the default behavior of selecting the route with the lowest BGP identifier (y) as best.

Default

no ignore-router-id

as-override

Syntax

[no] as-override

Context

```
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command replaces all instances of the peer AS number with the local AS number in a BGP route AS_PATH.

This command breaks the BGP loop detection mechanism and should be used carefully.

By default, the override command is not enabled.

authentication-key

Syntax

authentication-key [*authentication-key* | *hash-key*] [**hash** | **hash2**]

no authentication-key

Context

```
config>service>vprn>bgp  
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8CC

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.

By default, authentication is disabled and the authentication password is empty

The **no** form of this command removes the authentication password from the configuration and effectively disables authentication.

Parameters

authentication-key

Specifies the authentication key. This 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. This 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

Keyword to specify 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

Keyword to specify the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

auth-keychain

Syntax

auth-keychain *name*

Context

```
config>service>vprn>bgp
config>service>vprn>bgp>group
config>service>vprn>bgp>group>neighbor
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

bfd-enable

Syntax

[no] bfd-enable

Context

config>router>bgp

config>router>bgp>group

config>router>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 using the **bfd-enable** command under the IP interface.

The **no** form of this command removes BFD from the associated BGP protocol peering.

Default

no bfd-enable

connect-retry

Syntax

connect-retry *seconds*

no connect-retry

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the BGP connect retry timer.

When this timer expires, BGP tries to reconnect to the configured peer. this command 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 seconds

Parameters

seconds

Specifies 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables BGP route damping for learned routes that 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 using the 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>vpn>bgp

config>service>vpn>bgp>group

config>service>vpn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command disables the use of 4-byte AS numbers. This command 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 reverts to the default behavior, 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 will strictly support IPv4 routing exchanges with that peer.

The **no** form of this removes this command from the configuration and reverts to 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures BGP to disable sending communities.

Parameters

standard

Keyword to specify standard communities that existed before VPRNs or 2547.

extended

Keyword to specify 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures BGP fast external failover.

enable-peer-tracking

Syntax

[no] **enable-peer-tracking**

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables BGP peer tracking.

Default

no enable-peer-tracking

export

Syntax

export *policy* [*policy*...]

no export

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures 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.

IF a non-existent route policy is applied to a VPRN instance, a CLI warning message is generated. 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.

By default, BGP advertises routes from other BGP routes but does not advertise any routes from other protocols unless directed by an export policy.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the IP family capability.

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

Default

no family

Parameters

ipv4

Keyword that provisions IPv4 support.

ipv6

Keyword that provisions IPv6 support.

group

Syntax

group *name* [**dynamic-peer**]

no group

Context

config>service>vprn>bgp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8CC

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 shut down 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, etc.), the entire string must be enclosed within double quotes.

dynamic-peer

Keyword that designates the specific BGP group will be used by BGP peers created dynamically based on subscriber-hosts pointing to corresponding BGP peering policy. There can be only one BGP group with this flag set in any 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>vprn>bgp>group

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures 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 shut down, 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 *ip-address*: a.b.c.d

family

Syntax

family [ipv4]

no family

Context

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8CC

Description

This command configures 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, the supported address family reverts to the default.

Default

ipv4

Parameters

ipv4

Keyword that provisions support for IPv4 routing information.

hold-time

Syntax

hold-time *seconds* [**strict**]

no hold-time

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8CC

Description

This command configures the BGP hold time.

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 command 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.

- 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.
- 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 seconds

Parameters

seconds

Specifies 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

Keyword to specify that the advertised BGP hold-time from the far-end BGP peer must be greater than or equal to the specified value.

import

Syntax

import *policy* [*policy*...]

no import

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

By default, BGP accepts all routes from configured BGP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

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>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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, the configured **keepalive** value is overridden by the **hold-time** value under the following circumstances.

- 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 **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 seconds

Parameters

seconds

Specifies the keepalive timer in seconds, expressed as a decimal integer.

Values 0 to 21845

local-address

Syntax

local-address *ip-address*

no local-address

Context

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the local IP address used by the group or neighbor when communicating with BGP peers.

Outgoing connections use the **local-address** as the source of the TCP connection when initiating connections with a peer.

When a local address is not specified, the 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.

By default, the router ID is used when communicating with iBGP peers, and the interface address is used for directly connected eBGP peers

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.

Default

no local-address

Parameters

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>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a BGP virtual autonomous system (AS) number.

In addition to the ASN configured for BGP in the **config>router>autonomous-system** context, a virtual (local) ASN is configured. The virtual ASN is added to the as-path message before the router ASN makes the virtual AS the second AS in the as-path.

This command 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 for each 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 ASN. 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 ASN. Changing the local AS at the neighbor level in an active BGP instance causes BGP to reestablish the peer relationship with the new local ASN.

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 configured virtual ASN.

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

Keyword to specify that the local-as is hidden in paths learned from the peering.

local-preference

Syntax

local-preference *local-preference*

no local-preference

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables setting the BGP local-preference attribute in incoming routes if not specified, and configures the default value for the attribute. This value is used if the BGP route arrives from a BGP peer without the **local-preference** integer set.

The specified value can be overridden by any value set via a route policy. This command 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.

By default, this command does not override the **local-preference** value set in arriving routes and analyze routes without local preference with value of 100

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 a 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.

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>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures how the BGP peer session handles loop detection in the AS path.

This command 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.

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

Keyword that sends a notification to the remote peer and drops the session.

discard-route

Keyword that discards routes received with loops in the AS path.

ignore-loop

Keyword that ignores routes with loops in the AS path but maintains peering.

off

Keyword that 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command advertises the 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 command 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

Keyword that 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 command 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 seconds

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>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the minimum interval at which a prefix can be advertised to a peer.

This command 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 the default value.

Default

30 seconds

Parameters

seconds

Specifies the minimum route advertising interval, in seconds, expressed as a decimal integer.

Values 1 to 255

multihop

Syntax

multihop *ttl-value*

no multihop

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

passive

Syntax

[no] passive

Context

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables passive mode for the BGP group or neighbor.

When in passive mode, BGP does not attempt to actively connect to the configured BGP peers but responds only when it receives a connect open request from the peer.

By default, BGP actively tries to connect to all the configured peers.

The **no** form of this command used at the group level disables passive mode where BGP actively attempts to connect to its peers.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Default

no passive

```
peer-as
```

Syntax

```
peer-as as-number
```

Context

```
config>service>vprn>bgp>group
```

```
config>service>vprn>bgp>group>neighbor
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the AS number for the remote peer. The peer ASN must be configured for each configured peer.

For eBGP peers, the peer ASN configured must be different from the AS number configured for this router under the global level, because the peer will be in a different autonomous system than this router.

For iBGP peers, the peer ASN must be the same as the AS number of this router configured under the global level.

By default, no AS numbers are defined.

This is a required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.

Parameters

as-number

Specifies the AS number, expressed as a decimal integer.

Values 1 to 65535

preference

Syntax

[no] preference *preference*

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the route preference for routes learned from the configured peers.

This command 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

prefix-limit

Syntax

prefix-limit *family limit* [log-only] [threshold *percentage*]

no prefix-limit

Context

```
config>service>vprn>bgp>group  
config>service>vprn>bgp>group>neighbor
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the maximum number of routes BGP can learn from a peer.

When the number of routes reaches a specific percentage (default is 90% of this limit), an SNMP trap is generated. When the limit is exceeded, the BGP peering is dropped and disabled.

This command only applies to BGP routes learned for different families supported by the BGP protocol on the 7210 SAS. Different IP FIB limits are supported for different IPv4 and IPv6 address prefix lengths. There are two limits to consider: one is the value configured as part of the **prefix-limit** command and the second is the maximum IP FIB limit supported on the node. These two limits impact the behavior of the **prefix-limit** command as follows. When a BGP route for the configured *family* is received, the following comparison is completed:

- BGP peering is brought down if all of the following conditions are true:
 - if the number of routes in the FIB plus the received route is greater than the value configured for **prefix-limit**
 - if the number of routes is less than the maximum IP FIB limit
 - if **log-only** is not configured
- BGP peering is remains up if all of the following conditions are true:
 - if the number of routes in the FIB plus the received route is greater than the value configured for **prefix-limit**
 - if the number of routes is less than the maximum IP FIB limit
 - if **log-only** is configured

A log is generated to report the addition of the route if the **prefix-limit** value is exceeded. Excess routes are added to the IP FIB.

- If the number of routes in the FIB plus the received route is greater than the maximum IP FIB limit, regardless of whether **log-only** is configured, the BGP peering session is brought down.

The **no** form of this command removes the **prefix-limit** command.

Default

no prefix-limit

Parameters

limit

Specifies the number of routes, expressed as a decimal integer, that can be learned from a peer.

Values 1 to 4294967295

log-only

Keyword that enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. BGP peering is not dropped.

percent

Specifies the threshold value (as a percentage) that triggers a warning message.

Default 90%

family

Specifies the address family applied for the prefix limit.

Values ipv4, ipv6

rapid-withdrawal

Syntax

[no] **rapid-withdrawal**

Context

config>service>vprn>bgp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 reverts withdrawal processing to the normal behavior.

Default

no rapid-withdrawal

remove-private

Syntax

[no] **remove-private**

Context

config>service>vprn>bgp

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command allows private ASNs to be removed from the AS path before advertising them to BGP peers.

When this command is set at the global level, it applies to all peers regardless of group or neighbor configuration. When this command is set at the group level, it applies to all peers in the group regardless of the neighbor configuration.

The OS software recognizes the set of ASNs that are defined by IANA as private. These are ASNs in the range 64512 through 65535, inclusive.

By default, private ASNs will be included in the AS path attribute.

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

type

Syntax

[no] type {internal | external}

Context

config>service>vprn>bgp>group

config>service>vprn>bgp>group>neighbor

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the BGP peer type as **internal** or **external**.

The **internal** type indicates the peer is an iBGP peer, the **external** type indicates that the peer is an eBGP peer.

By default, the OS derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, 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

Parameters

internal

Keyword that configures the peer as internal.

external

Keyword that configures the peer as external.

8.4.2.4 OSPF commands

```
ospf
```

Syntax

```
[no] ospf
```

Context

```
config>service>vprn
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures 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>vprn>ospf
```


Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures an OSPF area. An area is a collection of network segments within an AS that have been administratively grouped. 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, and address-ranges 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 0.0.0.0 to 255.255.255.255 (dotted-decimal)
0 to 4294967295 (decimal integer)

area-range

Syntax

area-range [*ip-prefix/mask*] [*ipv6-prefix/prefix-length*] [**advertise** | **not-advertise**]

no area-range *ipv6-prefix/prefix-length*

Context

config>service>vprn>ospf>area

config>service>vprn>ospf>area>nssa

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures 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

ip-prefix/mask

Specifies the IP prefix for the range in dotted-decimal notation and the subnet mask for the range, expressed as a decimal integer.

ip-prefix: 0 to 32

mask: 0 to 32)

ipv6-prefix/prefix-length

Specifies the IPv6 prefix for the range in hexadecimal notation and the prefix length for the range.

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

prefix-length: 0 to 128

advertise | not-advertise

Keyword that specify whether to advertise the summarized range of addresses into other areas. The **advertise** keyword indicates the range is advertised, and the **not-advertise** keyword indicates the range is not advertised.

Default advertise

blackhole-aggregate

Syntax

[no] blackhole-aggregate

Context

```
config>service>vpn>ospf>area
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command installs a low priority blackhole route for the entire aggregate. Existing routes that make up the aggregate will have a higher priority, and only the components of the range for which no route exists 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** command.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures 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

Keyword that 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command is similar to a virtual link with the exception that a 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>vprn>ospf>area>if

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

By default, point-to-point interfaces are advertised as subnet routes.

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>vprn>ospf>area>if

config>service>vprn>ospf>area>virtual-link

config>service>vprn>ospf>area>sham-link

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the password used by the OSPF interface, virtual link, or sham link to send and receive OSPF protocol packets on the interface when simple password authentication is configured.

All neighboring routers must use the same type of authentication and password for correct protocol communication. If the **authentication-type** is configured as the password, this key must be configured.

By default, no authentication key is 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

Keyword to specify 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

Keyword to specify 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

config>service>vprn>ospf>area>sham-link

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables authentication and specifies the type of authentication to be used on the OSPF interface, virtual link, or sham link.

Both simple **password** and **message-digest** authentication are supported.

By default, authentication is not enabled on an interface.

The **no** form of this command disables authentication on the interface.

Default

no authentication

Parameters

password

Keyword that 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

Keyword that 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

[no] **bfd-enable** [**remain-down-on-failure**]

Context

config>service>vprn>ospf>area>interface

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the use of 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 using the **bfd-enable** command under the IP interface.

BFD is not supported for IPv6 interfaces.

The **no** form of this command removes BFD from the associated IGP protocol adjacency.

Default

no bfd-enable

Parameters

remain-down-on-failure

Keyword that forces adjacency down on BFD failure.

dead-interval

Syntax

dead-interval *seconds*

no dead-interval

Context

config>service>vprn>ospf>area>if

config>service>vprn>ospf>area>sham-link

config>service>vprn>ospf>area>virtual-link

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the time 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** applies to an interface, all nodes on the subnet must have the same dead interval.

Sham link

If the **dead-interval** applies to a sham link, the interval on both endpoints of the sham link must have the same dead interval.

Virtual link

If the **dead-interval** applies to a virtual link, the interval on both endpoints of the virtual link must have the same dead interval.

Parameters

seconds

Specifies the dead interval in seconds.

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>sham-link

config>service>vprn>ospf>area>virtual-link

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the interval between OSPF hellos issued on the interface, virtual link, or sham link.

This command, in combination with the **dead-interval** command, is used to establish and maintain the adjacency. Use this command 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 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** applies to an interface, all nodes on the subnet must have the same hello interval.

Sham link

If the **hello-interval** applies to a sham link, the interval on both endpoints of the sham link must have the same hello interval.

Virtual link

If the **hello-interval** applies to a virtual link, the interval on both endpoints of the virtual 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>vprn>ospf>area>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the interface type as 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

Keyword that 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 medium, such as Ethernet.

point-to-point

Keyword that configures the interface to maintain this link as a point-to-point link.

message-digest-key

Syntax

message-digest-key *key-id* **md5** [*key* | *hash-key*] [**hash**]

no message-digest-key *key-id*

Context

config>service>vprn>ospf>area>if

config>service>vprn>ospf>area>sham-link

config>service>vprn>ospf>area>virtual-link

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

The **no** form of this command removes the message digest key identified by the *key-id*.

Parameters

key-id

Specifies the *key-id*, 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

Keyword to specify the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

metric

Syntax

metric *metric*

no metric

Context

```
config>service>vprn>ospf>area>if  
config>service>vprn>ospf>area>sham-link  
config>service>vprn>ospf>area>virtual-link
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

By default, the metric is based on **reference-bandwidth** setting and the link speed.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the OSPF interface MTU value used when negotiating an OSPF packet size used on this interface. If this command is not configured, OSPF derives the MTU value from the MTU configured (default or explicitly) in the **config>port>ethernet** context.

If this command is configured, the smaller value between the value in this command here and the MTU configured (default or explicitly) in one of the previously mentioned contexts is used.

By default, the value derived from the MTU configured in the **config>port** context is used.

To determine the actual packet size, add 14 bytes for an Ethernet packet and 18 bytes for a tagged Ethernet packet to the size of the OSPF (IP) packet MTU configured in this command.

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

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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** command allows an interface to be advertised as an OSPF interface without running the OSPF protocol.

While in passive mode, the interface will ignore ingress OSPF protocol packets and not transmit any OSPF protocol packets.

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>vprn>ospf>area>if
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the priority of the OSPF interface that is used in an election of the designated router on the subnet.

This command is used only if the interface is of the broadcast type. The router with the highest priority interface becomes the designated router. A router with a priority value of 0 is not eligible to be a 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.

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>sham-link
```

```
config>service>vprn>ospf>area>virtual-link
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies the length of time 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 interval.

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>vprn>ospf>area>if

config>service>vprn>ospf>area>sham-link

config>service>vprn>ospf>area>virtual-link

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the estimated time to transmit an LSA on the interface or sham-link.

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

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>vprn>ospf>area

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures an OSPF Not So Stubby Area (NSSA) and adds or removes the NSSA designation from the area.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is an NSSA has the capability to flood external routes that it learns throughout its area and via an ABR to the entire OSPF domain.

Existing virtual links of a non-stub area or NSSA are removed when the designation is changed to NSSA or stub.

An area can be designated as stub or NSSA but never both at the same time.

By default, an area is not configured as an NSSA.

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>vprn>ospf>area>nssa

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the type-7 parameter, which injects a type 7 LSA default route instead of a type 3 LSA into the NSSA configured with no summaries.

To return to a type 3 LSA, enter the **originate-default-route** command without the **type-7** parameter.

When configuring an NSSA with no summaries, the ABR will inject a type 3 LSA default route into the NSSA.

The **no** form of this command disables origination of a default route.

Default

no originate-default-route

Parameters

type-7

Specifies that a type 7 LSA is used for the default route.

Default type 3 LSA for the default route

redistribute-external

Syntax

[no] redistribute-external

Context

config>service>vpn>ospf>area>nssa

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables the redistribution of external routes into the NSSA or an NSSA ABR that is exporting the routes into non-NSSA areas.

NSSAs are similar to stub areas in that no external routes are imported into the area from other OSPF areas. The major difference between a stub area and an NSSA is that the NSSA has the capability to flood external routes that it learns (providing it is an ASBR) throughout its area and via an ABR to the entire OSPF domain.

The **no** form of 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables sending summary (type 3) advertisements into a stub area or NSSA on an ABR. This command is particularly useful to reduce the size of the routing and Link State Database (LSDB) tables within the stub area or NSSA. 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures an OSPF or stub area and adds or removes the stub designation from the area. External routing information is not flooded into stub areas. All routers in the stub area must be configured with the **stub** command. An OSPF area cannot be both an NSSA and a stub area. Existing virtual links of a non-stub area or non-NSSA will be removed when its designation is changed to NSSA or stub.

By default, an area is not a stub area.

The **no** form of 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>vprn>ospf>area>stub

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the metric used by the ABR for the default route into a stub area. The default metric should only be configured on an ABR of a stub area. An ABR generates a default route if the area is a **stub** area.

The **no** form of 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>vprn>ospf>area

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures a virtual link to connect ABRs to the backbone using 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 88: OSPF areas](#)), the ABRs

(routers 1 and 2 in [Figure 88: OSPF areas](#)) must be connected via a virtual link. The two ABRs will form a point-to-point like adjacency across the transit area (area 0.0.0.1 in [Figure 88: OSPF areas](#)). A virtual link can be configured only while in the area 0.0.0.0 context.

The *router-id* specified in this command must be associated with the virtual neighbor. The transit area cannot be a stub area or an NSSA.

The **no** form of 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 transit area that links the backbone area with the area that has no physical connection with the backbone.

[Figure 88: OSPF areas](#) shows an OSPF backbone area configuration. 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) the ABRs (such as routers Y and Z) must be connected via a virtual link. The two ABRs form a point-to-point like adjacency across the transit area (see area 0.0.0.4).

Figure 88: OSPF areas



compatible-rfc1583

Syntax

[no] compatible-rfc1583

Context

config>service>vprn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures 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*...up to 5 max]

no export

Context

config>service>vprn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command associates export route policies to determine which routes are exported from the route table to OSPF. Export policies are in effect only if OSPF is configured as an ASBR.

If no export policy is specified, non-OSPF routes are not exported from the routing table manager to OSPF.

If multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied. If multiple export commands are issued, the last command entered will override the previous command. A maximum of five policy names can be specified.

The **no** form of 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 (such #, \$, spaces), 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>vpn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures limits on the number of non-default AS-external-LSA entries that can be stored in the link-state database (LSDB) and specifies a wait timer before processing these 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 LSDB. Placing a limit on the non-default AS-external-LSAs in the LSDB protects the router from receiving an excessive number of external routes that consume excessive memory or CPU resources. If the number of routes reach or exceed the *limit*, the table is in an overflow state. When in an overflow state, the router will not originate any new AS-external-LSAs. In fact, it withdraws all the self-originated non-default external LSAs.

The *interval* specifies the amount of time to wait after an overflow state before regenerating and processing non-default AS-external-LSAs. The waiting period acts like a dampening period preventing the router from continuously running Shortest Path First (SPF) calculations caused by the excessive number of non-default AS-external LSAs.

The **external-db-overflow** must be set identically on all routers attached to a regular OSPF area. OSPF stub areas and 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, expressed as a decimal integer, after entering an overflow state before attempting to process non-default AS-external-LSAs.

Values 0 to 2147483647

external-preference

Syntax

external-preference *preference*

no external-preference

Context

config>service>vpn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 will be used.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is based on the default preferences, as shown in the following table.

Table 77: Route preference defaults by route type

Route type	Preference	Configurable
Direct attached	0	No
Static routes	5	Yes
OSPF internal	10	Yes ¹³
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

¹³ 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

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 in the **config>router** context.

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.

Values 1 to 255

graceful-restart

Syntax

[no] graceful-restart

Context

config>service>vprn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures graceful restart for OSPF in the VPRN context. When the control plane of a GR-capable router fails, the VPRN OSPF peers (GR helpers) temporarily preserve neighbor information, so packets continue to be forwarded through the failed GR router using the last known routes. The helper state remains until the peer completes its restart or exits if the GR timer value is exceeded.

The **no** form of this command disables graceful restart and removes all graceful restart configurations in the OSPF instance.

Default

no graceful-restart

helper-disable

Syntax

[no] helper-disable

Context

config>service>vprn>ospf>graceful-restart

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command disables the helper support for graceful restart.

When **graceful-restart** is enabled, the router can be a helper (meaning that the router is helping a neighbor to restart) or be a restarting router or both. The 7210 SAS supports only helper mode. This facilitates the graceful restart of neighbors but does not act as a restarting router (meaning that the 7210 SAS will not help the neighbors to restart).

The **no helper-disable** command enables helper support and is the default when **graceful-restart** is enabled.

Default

helper-disable

ignore-dn-bit

Syntax

[no] ignore-dn-bit

Context

config>service>vprn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 will be ignored. When disabled, the DN bit will not be ignored for OSPF LSA packets.

overload

Syntax

overload [*timeout seconds*]

no overload

Context

config>service>vprn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 continues to reach the router.

Put the IGP in an overload state by entering a timeout value. The IGP will enter the overload state until the timeout timer expires or a **no overload** command is executed.

If the **overload** command is encountered during the execution of an [overload-on-boot](#) command 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 reverts to the default value. When the **no overload** command is executed, the overload state is terminated regardless of 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>vprn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

When the router is in an overload state, the router is used only if there is no other router to reach the destination. This command configures the IGP upon bootup in the overload state until one of the following events occur:

- timeout timer expires
- manual override of the current overload state is entered with the **no overload** command

The **no overload** command does not affect the **overload-on-boot** command.

The **no** form of this command removes the **overload-on-boot** command 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>vprn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the preference for internal routes for an OSPF instance.

Different protocols should not be configured with the same preference. If this occurs, the tiebreaker is based on the default preference table, as listed in [Table 77: Route preference defaults by route type](#). If multiple routes are learned with an identical preference using the same protocol, the lowest cost route is used.

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. [Table 77: Route preference defaults by route type](#) lists the defaults for different route types.

Values 1 to 255

reference-bandwidth

Syntax

reference-bandwidth *reference-bandwidth*

no reference-bandwidth

Context

config>service>vprn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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:

$\text{cost} = \text{reference-bandwidth} / \text{bandwidth}$

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** 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

router-id

Syntax

router-id *ip-address*

no router-id

Context

config>service>vprn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the router ID for the OSPF instance.

Configuring the router ID in the base instance of OSPF overrides the router ID configured in the **config>router** context.

The default value for the base instance is inherited from the configuration in the **config>router** context. If the router ID in the **config>router** context is not configured, the following applies:

- the system uses the system interface address (which is also the loopback address)
- if a system interface address is not configured, use the last 32 bits of the chassis MAC address

When configuring a new router ID, the instance is not automatically restarted with the new router ID. The next time the instance is initialized, the new router ID is used.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for the instance, or reboot the entire router.

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

Default

0.0.0.0 (base OSPF)

Parameters

ip-address

Specifies a 32-bit unsigned integer uniquely identifying the router in the AS.

super-backbone

Syntax

[no] **super-backbone**

Context

config>service>vpn>ospf

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command specifies whether CE-PE functionality is required. The OSPF super backbone indicates the type of the LSA generated as a result of routes redistributed into OSPF. When enabled, the redistributed routes are injected as summary, external, or NSSA LSAs. When disabled, the redistributed routes are injected as either external or NSSA LSAs only.

Default

no super-backbone

suppress-dn-bit

Syntax

[no] **suppress-dn-bit**

Context

```
config>service>vprn>ospf
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 will not be set. When disabled, this instance of the OSPF router will follow the normal procedure to determine whether to set the DN bit.

Default

```
no suppress-dn-bit
```

timers

Syntax

```
timers
```

Context

```
config>service>vprn>ospf
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context configure OSPF timers. Timers control the delay between receipt of an LSA requiring an SPF calculation and the minimum time between successive SPF calculations.

Changing the timers affects CPU utilization and network reconvergence times. Lower values reduce convergence time but increase CPU utilization. Higher values reduce CPU utilization but increase reconvergence time.

lsa-arrival

Syntax

```
lsa-arrival lsa-arrival-time
```

```
no lsa-arrival
```

Context

```
config>service>vprn>ospf>timers
```

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the minimum delay that must pass between receipt of the same LSAs arriving from neighbors. Nokia recommends that the neighbor configured (**Isa-generate**) *Isa-second-wait* interval is equal to or greater than the **Isa-arrival** timer configured in this command.

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

Default

no **Isa-arrival**

Parameters

Isa-arrival-time

Specifies the timer in milliseconds. Values entered that do not match this requirement are rejected.

Values 0 to 600000

Isa-generate

Syntax

Isa-generate *max-Isa-wait* [*Isa-initial-wait* [*Isa-second-wait*]]

no Isa-generate-interval

Context

config>service>vprn>ospf>timers

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures 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 *Isa-second-wait* timer until a maximum value is reached. Configuring the **Isa-arrival** interval to equal or less than the *Isa-second-wait* interval configured in the **Isa-generate** command is recommended.

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

Default

no **Isa-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 1 or in millisecond increments. Values entered that do not match this requirement are rejected.

Values 1 to 600000

spf-wait

Syntax

spf-wait *max-spf-wait* [*spf-initial-wait* [*spf-second-wait*]]

no spf-wait

Context

config>service>vprn>ospf>timers

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command defines the maximum interval between two consecutive SPF calculations. Timers that determine when to initiate the first, second, and subsequent SPF calculations after a topology change occurs can be controlled with this command. Subsequent SPF runs (if required) will occur at exponentially increasing intervals of the *spf-second-wait* interval. For example, if the *spf-second-wait* interval is 1000, the next SPF will run after 2000 milliseconds, and the next SPF will run after 4000 milliseconds, and so on, until it reaches the **spf-wait** value. The SPF interval will stay at the **spf-wait** value until there are no more SPF runs scheduled in that interval. After a full interval without any SPF runs, the SPF interval will drop back to *spf-initial-wait*.

The timer must be entered in increments of 100 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command configures the type of the extended community attribute exchanged using BGP to carry the OSPF VPN domain ID. 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 any order.

Default

no vpn-domain

Parameters

id

Specifies the OSPF VPN domain in the "xxx.xxx.xxx" format. This is exchanged using BGP in the extended community attribute associated with a prefix.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

8.4.2.5 VPRN show commands

egress-label

Syntax

egress-label *start-label* [*end-label*]

Context

show>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays services using the range of egress labels.

If only the mandatory *start-label* parameter is specified, only services using the specified label are displayed.

If both *start-label* and *end-label* parameters are specified, the services using 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-label1* is specified, services only using *egress-label1* 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-label1* value.

Output

The following output is an example of egress label information, and [Table 78: 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 78: 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 $start-label \leq X \leq end-label$ 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 ingress label information, and [Table 79: 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 79: 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 specified label range.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

ip-address

Specifies the IP address of the interface for which to display matching SAPs.

Values 1.0.0.0 to 223.255.255.255

ip-int-name

Specifies the IP interface name for which to display matching SAPs.

ingress

Keyword that specifies matching an ingress policy.

egress

Keyword that 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 SAP information, and [Table 80: 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 80: 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.
Opr	The actual state of the SAP.

sdp

Syntax

sdp [*sdp-id* | **far-end** *ip-address*] [**detail** | **keep-alive-history**]

Context

show>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 81: 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                Oper Path MTU     : 0
Far End         : 10.10.10.104    Delivery           : MPLS
Admin State     : Up                Oper State        : Down
Flags          : SignalingSessDown TransportTunnDown
Signaling       : TLDP                VLAN VC Etype     : 0x8100
Last Status Change : 02/01/2007 09:11:39  Adv. MTU Over.    : No
Last Mgmt Change  : 02/01/2007 09:11:46

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

Associated LSP LIST :
Lsp Name       : to-104
Admin State    : Up                Oper State        : Down
Time Since Last Tran*: 01d07h36m
=====
* indicates that the corresponding row element may have been truncated.
*A:ALA-12#

*A:dut-a>show>service# sdp
=====
Services: Service Destination Points
=====
SdpId  AdmMTU  OprMTU  IP address   Adm  Opr       Del   LSP   Sig
-----
17406  0       9166   10.2.2.2    Up   Up        MPLS  L/B   TLDP
17407  0       9166   10.3.3.3    Up   Up        MPLS  L/B   TLDP
-----
Number of SDPs : 2
-----
Legend: R = RSVP, L = LDP, B = BGP, M = MPLS-TP, n/a = Not Applicable
        I = SR-ISIS, 0 = SR-OSPF
=====
*A:dut-a>show>service#
    
```

Table 81: Output fields: SDP

Label	Description
SDP Id	The SDP identifier.

Label	Description
Adm MTU	Specifies the largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Opr MTU	Specifies the actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
IP address	Specifies the IP address of the remote end of the MPLS tunnel defined by this SDP.
Adm Admin State	Specifies the state of the SDP.
Opr Oper State	Specifies the operating state of the SDP.
Flags	Specifies all the conditions that affect the operating status of this SDP.
Signal Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
Last Status Change	Specifies the time of the most recent operating status change to this SDP.
Last Mgmt Change	Specifies the time of the most recent management-initiated change to this SDP.
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Deliver Delivered	Specifies the type of delivery used by the SDP: MPLS.
Number of SDPs	Specifies the total number of SDPs displayed according to the criteria specified.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hello Timeout	Specifies the number of seconds to wait for an SDP echo response message before declaring a timeout.

Label	Description
Unmatched Replies	Specifies the number of SDP unmatched message replies.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hold Down Time	Specifies the 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 82: 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

*A:dut-a>show>service# sdp-using
=====
SDP Using
=====
SvcId      SdpId      Type  Far End      Opr  I.Label  E.Label
              State
-----
5          17406:4294967294  BgpAd 2.2.2.2     Up   131067  131066
5          17407:4294967295  BgpAd 3.3.3.3     Up   131066  131066
-----
Number of SDPs : 2
-----
*A:dut-a>show>service#
```

Table 82: 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.
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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

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 83: 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:de14# show service service-using
=====
Services
=====
ServiceId  Type      Adm  Opr      CustomerId  Last Mgmt Change
-----
```

```

1      uVPLS      Up      Up      1      10/26/2006 15:44:57
2      Epipe     Up      Down    1      10/26/2006 15:44:57
10     mVPLS     Down   Down    1      10/26/2006 15:44:57
11     mVPLS     Down   Down    1      10/26/2006 15:44:57
100    mVPLS     Up      Up      1      10/26/2006 15:44:57
101    mVPLS     Up      Up      1      10/26/2006 15:44:57
102    mVPLS     Up      Up      1      10/26/2006 15:44:57
999    uVPLS     Down   Down    1      10/26/2006 16:14:33
-----
Matching Services : 8
-----
A:de14#
    
```

Table 83: 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** | **bgp-ad** | **fdb** | **labels** | **mfib** | **sap** | **sdp** | **split-horizon-group** | **stp**}

Context

show>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Displays detailed information about the service.

arp

Displays ARP entries for the service.

base

Displays basic service information.

bgp-ad

Displays BGP AD related information for the service.

fdb

Displays FDB entries.

interface

Displays service interfaces.

labels

Displays labels being used by this service.

sap

Displays SAPs associated to the service.

sdp

Displays SDPs associated with the service.

split-horizon-group

Displays split horizon group information.

stp

Displays STP information.

all

Syntax

all

Context

show>service>id

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 84: Output fields: service ID all](#) describes the output fields.

Sample output

```
*A:duth# show service id 502 all

=====
Service Detailed Information
=====
Service Id       : 502                Vpn Id          : 0
Service Type    : VPRN
Name            : (Not Specified)
Description     : (Not Specified)
Customer Id     : 1                  Creation Origin  : manual
Last Status Change: 12/23/2014 04:57:18
Last Mgmt Change  : 12/23/2014 04:57:18
Admin State     : Up                  Oper State      : Up

Route Dist.     : 502:502            VPRN Type       : regular
AS Number       : 100                Router Id       : 8.8.8.8
ECMP            : Enabled            ECMP Max Routes : 1

Auto Bind Tunnel
Resolution      : filter
Filter Protocol : rsvp

Max IPv6 Routes : No Limit
Ignore NH Metric : Disabled
Vrf Target      : target:502:502
Vrf Import     : None
Vrf Export     : None
MVPN Vrf Target : None
MVPN Vrf Import : None
MVPN Vrf Export : None
Label mode     : vrf
BGP VPN Backup  : Disabled

SAP Count       : 0                  SDP Bind Count  : 5
Allow IP Intf Bind: Disabled
VSD Domain     : <none>

-----
ETH-CFM service specifics
-----
Tunnel Faults   : ignore

-----
VPRN service Network Specifics
-----
Ing Net QoS Policy : none
Ingress FP QGrp   : (none)           Ing FP QGrp Inst : (none)

-----
Service Destination Points(SDPs)
-----
Sdp Id 502:502  -(9.9.9.9)
-----
Description     : (Not Specified)
SDP Id         : 502:502              Type           : Spoke
Spoke Descr    : (Not Specified)
VC Type       : n/a                  VC Tag         : n/a
Admin Path MTU : 0                  Oper Path MTU  : 9186
Delivery      : MPLS
```

```
Far End : 9.9.9.9
Tunnel Far End : n/a
LSP Types : RSVP

Admin State : Up
Acct. Pol : None
Ingress Label : n/a
Ingr Mac Fltr-Id : n/a
Ingr IP Fltr-Id : n/a
Ingr IPv6 Fltr-Id : n/a
BGP IPv4 FlowSpec : Disabled
BGP IPv6 FlowSpec : Disabled
Admin ControlWord : Not Preferred
BFD Template : None
BFD-Enabled : no
Last Status Change : 12/23/2014 18:12:34
Last Mgmt Change : 12/21/2014 19:39:56
Class Fwding State : Down
Flags : None
Local Pw Bits : None
Peer Pw Bits : None
Peer Fault Ip : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None

Oper State : Up
Collect Stats : Disabled
Egress Label : n/a
Egr Mac Fltr-Id : n/a
Egr IP Fltr-Id : n/a
Egr IPv6 Fltr-Id : n/a

Oper ControlWord : False
BFD-Encap : ipv4
Signaling : n/a

Application Profile: 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. : n/a
E. Fwd. Pkts. : n/a
Extra-Tag-Drop-Pkts: n/a
I. Fwd. Octs. : n/a
E. Fwd. Octets : n/a
Extra-Tag-Drop-0c*: n/a

-----
ETH-CFM SDP-Bind specifics
-----
Squelch Levels : None

-----
RSVP/Static LSPs
-----
Associated LSP List :
Lsp Name : i2
Admin State : Up
Time Since Last Tr*: 04d06h44m
Oper State : Up

-----
Sdp Id 602:502 - (2.2.2.2)
-----
Description : (Not Specified)
SDP Id : 602:502
Spoke Descr : (Not Specified)
VC Type : n/a
Admin Path MTU : 0
Delivery : MPLS
Far End : 2.2.2.2
Tunnel Far End : n/a
Type : Spoke
VC Tag : n/a
Oper Path MTU : 9186
LSP Types : RSVP

Admin State : Up
Acct. Pol : None
Ingress Label : n/a
Oper State : Up
Collect Stats : Disabled
Egress Label : n/a
```

```
Ingr Mac Fltr-Id : n/a          Egr Mac Fltr-Id : n/a
Ingr IP Fltr-Id  : n/a          Egr IP Fltr-Id  : n/a
Ingr IPv6 Fltr-Id : n/a        Egr IPv6 Fltr-Id : n/a
BGP IPv4 FlowSpec : Disabled
BGP IPv6 FlowSpec : Disabled
Admin ControlWord : Not Preferred Oper ControlWord : False
BFD Template     : None
BFD-Enabled      : no           BFD-Encap       : ipv4
Last Status Change : 12/23/2014 18:12:24 Signaling       : n/a
Last Mgmt Change  : 12/21/2014 19:39:56
Class Fwding State : Down
Flags            : None
Local Pw Bits    : None
Peer Pw Bits     : None
Peer Fault Ip    : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None

Application Profile: None

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

Statistics           :
I. Fwd. Pkts.       : n/a            I. Fwd. Octs.     : n/a
E. Fwd. Pkts.       : n/a            E. Fwd. Octets    : n/a
Extra-Tag-Drop-Pkts: n/a            Extra-Tag-Drop-Oc*: n/a

-----
ETH-CFM SDP-Bind specifics
-----
Squelch Levels      : None

-----
RSVP/Static LSPs
-----
Associated LSP List :
Lsp Name            : b2
Admin State         : Up              Oper State         : Up
Time Since Last Tr*: 04d06h44m

-----
Sdp Id 702:502  -(7.7.7.7)
-----
Description         : (Not Specified)
SDP Id              : 702:502          Type               : Spoke
Spoke Descr        : (Not Specified)
VC Type            : n/a             VC Tag             : n/a
Admin Path MTU     : 0               Oper Path MTU      : 9186
Delivery           : MPLS
Far End            : 7.7.7.7
Tunnel Far End     : n/a             LSP Types          : RSVP

Admin State        : Up              Oper State         : Up
Acct. Pol         : None            Collect Stats      : Disabled
Ingress Label     : n/a            Egress Label      : n/a
Ingr Mac Fltr-Id : n/a            Egr Mac Fltr-Id  : n/a
Ingr IP Fltr-Id  : n/a            Egr IP Fltr-Id   : n/a
Ingr IPv6 Fltr-Id : n/a           Egr IPv6 Fltr-Id : n/a
BGP IPv4 FlowSpec : Disabled
BGP IPv6 FlowSpec : Disabled
Admin ControlWord : Not Preferred Oper ControlWord  : False
```

```

BFD Template      : None
BFD-Enabled       : no
Last Status Change : 12/23/2014 18:12:27
Last Mgmt Change  : 12/21/2014 19:39:56
Class Fwding State : Down
Flags             : None
Local Pw Bits     : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None

Application Profile: None

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

Statistics          :
I. Fwd. Pkts.       : n/a
E. Fwd. Pkts.       : n/a
Extra-Tag-Drop-Pkts: n/a
I. Fwd. Octs.       : n/a
E. Fwd. Octets      : n/a
Extra-Tag-Drop-0c* : n/a

-----
ETH-CFM SDP-Bind specifics
-----
Squelch Levels      : None

-----
RSVP/Static LSPs
-----
Associated LSP List :
Lsp Name            : g2
Admin State          : Up
Time Since Last Tr* : 04d06h44m
Oper State           : Up

-----
Sdp Id 802:502 - (99.99.99.99)
-----
Description          : (Not Specified)
SDP Id               : 802:502
Spoke Descr         : (Not Specified)
VC Type              : n/a
Admin Path MTU       : 0
Delivery             : MPLS
Far End              : 99.99.99.99
Tunnel Far End       : n/a
Type                 : Spoke
VC Tag               : n/a
Oper Path MTU        : 9182
LSP Types            : BGP

Admin State          : Up
Acct. Pol            : None
Ingress Label        : n/a
Ingr Mac Fltr-Id     : n/a
Ingr IP Fltr-Id      : n/a
Ingr IPv6 Fltr-Id    : n/a
BGP IPv4 FlowSpec    : Disabled
BGP IPv6 FlowSpec    : Disabled
Admin ControlWord    : Not Preferred
BFD Template         : None
BFD-Enabled          : no
Last Status Change   : 12/23/2014 18:12:41
Last Mgmt Change     : 12/21/2014 19:39:56
Class Fwding State   : Down
Flags                : None
Oper State           : Up
Collect Stats        : Disabled
Egress Label         : n/a
Egr Mac Fltr-Id     : n/a
Egr IP Fltr-Id      : n/a
Egr IPv6 Fltr-Id    : n/a
Oper ControlWord     : False
BFD-Encap           : ipv4
Signaling            : n/a
    
```

```
Local Pw Bits      : None
Peer Pw Bits      : None
Peer Fault Ip     : None
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None

Application Profile: None

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

Statistics          :
I. Fwd. Pkts.     : n/a                I. Fwd. Octs.    : n/a
E. Fwd. Pkts.     : n/a                E. Fwd. Octets   : n/a
Extra-Tag-Drop-Pkts: n/a            Extra-Tag-Drop-0c*: n/a

-----
ETH-CFM SDP-Bind specifics
-----
Squelch Levels    : None

-----
BGP Tunnel Information :
-----
BGP LSP Id        : 262147

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

-----
Sdp Id 902:502 - (22.22.22.22)
-----
Description       : (Not Specified)
SDP Id            : 902:502                Type              : Spoke
Spoke Descr      : (Not Specified)
VC Type          : n/a                    VC Tag            : n/a
Admin Path MTU   : 0                      Oper Path MTU     : 9182
Delivery         : MPLS
Far End          : 22.22.22.22
Tunnel Far End   : n/a                    LSP Types        : BGP

Admin State      : Up                      Oper State        : Up
Acct. Pol       : None                    Collect Stats     : Disabled
Ingress Label   : n/a                    Egress Label     : n/a
Ingr Mac Fltr-Id : n/a                  Egr Mac Fltr-Id  : n/a
Ingr IP Fltr-Id : n/a                  Egr IP Fltr-Id   : n/a
Ingr IPv6 Fltr-Id : n/a                Egr IPv6 Fltr-Id : n/a
BGP IPv4 FlowSpec : Disabled
BGP IPv6 FlowSpec : Disabled
Admin ControlWord : Not Preferred          Oper ControlWord  : False
BFD Template    : None
BFD-Enabled     : no                    BFD-Encap        : ipv4
Last Status Change : 12/23/2014 19:25:37      Signaling         : n/a
Last Mgmt Change  : 12/21/2014 19:39:56
Class Fwding State : Down
Flags           : None
Local Pw Bits   : None
Peer Pw Bits    : None
Peer Fault Ip   : None
```

```
Peer Vccv CV Bits : None
Peer Vccv CC Bits : None

Application Profile: None

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

Statistics        :
I. Fwd. Pkts.    : n/a               I. Fwd. Octs.   : n/a
E. Fwd. Pkts.    : n/a               E. Fwd. Octets  : n/a
Extra-Tag-Drop-Pkts: n/a           Extra-Tag-Drop-0c*: n/a
```

ETH-CFM SDP-Bind specifics

Squelch Levels : None

BGP Tunnel Information :

BGP LSP Id : 262387

RSVP/Static LSPs

Associated LSP List :
No LSPs Associated

Number of SDPs : 5

* indicates that the corresponding row element may have been truncated.

Service Access Points

No Sap Associations

Service Interfaces

Interface

```
If Name       : toixial
Admin State   : Up                               Oper (v4/v6)  : Up/Down
Protocols     : None
IP Addr/mask  : 108.2.1.1/24                     Address Type  : Primary
IGP Inhibit   : Disabled                         Broadcast Addr: Host-ones
HoldUp-Time   : 0                               Track Srrp Inst : 0
Description   : N/A
```

Details

```
Description   : (Not Specified)
If Index      : 2                               Virt. If Index : 2
Last Oper Chg : 12/23/2014 20:36:08             Global If Index : 48
Mon Oper Grp  : None
Srrp En Rtng  : Disabled                       Hold time     : N/A
Port Id       : rvpls
TOS Marking   : Trusted                         If Type       : VPRN
```

```
Sntp B.Cast      : False
MAC Address     : 00:00:0b:00:01:00
Ingress stats  : Disabled
ARP Timeout    : 14400s
ARP Retry Timer: 5000ms
ARP Limit      : Disabled
ARP Threshold  : Disabled
ARP Limit Log Only: Disabled
IP MTU         : (default)
IP Oper MTU    : 1500
LdpSyncTimer  : None
LSR Load Balance : system
EGR Load Balance : both
Vas If Type    : none
TEID Load Balance : Disabled
SPI Load Balance : Disabled
uRPF Chk      : disabled
uRPF Ipv6 Chk : disabled
PTP HW Assist  : Disabled
Mpls Rx Pkts  : 0
Mpls Tx Pkts  : 0
Mac Accounting : Disabled
IPv6 DAD      : Enabled
IPv6 Nbr ReachTime: 30s
IPv6 stale time : 14400s
IPv6 Nbr Limit : Disabled
IPv6 Nbr Threshold: 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    : rvpls-502
Binding Status : Up
Reason       : (Not Specified)
Egr Reclass Plcy : 0
Ing Filter   : none
Ingr IPv6 Flt : none
EVPN Tunnel  : false

-----
Admin Groups
-----
No Matching Entries
-----

Srlg Groups
-----
No Matching Entries
-----

Interface
-----
If Name      : toB-Ixia
Admin State  : Up
Oper (v4/v6) : Up/Down
```



```
Protocols      : OSPFv2
IP Addr/mask   : 200.41.2.8/24      Address Type   : Primary
IGP Inhibit   : Disabled          Broadcast Addr : Host-ones
HoldUp-Time   : 0                 Track Srrp Inst : 0
Description    : N/A
```

Details

```
Description    : (Not Specified)
If Index       : 3                 Virt. If Index : 3
Last Oper Chg  : 12/23/2014 20:00:55 Global If Index : 49
Mon Oper Grp   : None
Srrp En Rtng  : Disabled          Hold time      : N/A
Port Id       : rvpls
TOS Marking    : Trusted          If Type        : VPRN
SNTP B.Cast    : False
MAC Address    : c4:08:4a:6f:99:65 Mac Accounting  : Disabled
Ingress stats  : Disabled          IPv6 DAD       : Enabled
ARP Timeout    : 14400s           IPv6 Nbr ReachTime: 30s
ARP Retry Timer : 5000ms           IPv6 stale time : 14400s
ARP Limit      : Disabled          IPv6 Nbr Limit  : Disabled
ARP Threshold  : Disabled          IPv6 Nbr Threshold: Disabled
ARP Limit Log Only: Disabled       IPv6 Nbr Log Only : Disabled
IP MTU         : (default)
IP Oper MTU    : 1500
LdpSyncTimer   : None
LSR Load Balance : system
EGR Load Balance : both
Vas If Type    : none
TEID Load Balance : Disabled
SPI Load Balance : Disabled
uRPF Chk      : disabled
uRPF Ipv6 Chk : disabled
PTP HW Assist  : Disabled
Mpls Rx Pkts  : 0                 Mpls Rx Bytes  : 0
Mpls Tx Pkts  : 0                 Mpls Tx Bytes  : 0
```

DHCP6 Relay Details

```
Description    : (Not Specified)
Admin State    : Down              Lease Populate  : 0
Oper State     : Down              Nbr Resolution : Disabled
If-Id Option   : None              Remote Id      : Disabled
Src Addr       : Not configured
Python plcy    : (Not Specified)
```

DHCP6 Server Details

```
Admin State    : Down              Max. Lease States : 8000
```

ICMP Details

```
Redirects      : Number - 100      Time (seconds) - 10
Unreachables   : Number - 100      Time (seconds) - 10
TTL Expired    : Number - 100      Time (seconds) - 10
```

Routed VPLS Details

```
VPLS Name      : rvpls-622
Binding Status  : Up
Reason         : (Not Specified)
Egr Reclass Plcy : 0
Ing Filter     : none
Ingr IPv6 Flt  : none
EVPN Tunnel    : false
```

Admin Groups

```

-----
No Matching Entries
-----

-----
Srlg Groups
-----
No Matching Entries
-----
=====
*A:duth#
    
```

Table 84: Output fields: service ID all

Label	Description
Service Detailed Information	
Service Id	The service identifier.
VPN Id	The number that identifies the VPN.
Customer Id	The customer identifier.
Last Status Change	The date and time of the most recent change in the administrative or operating status of the service.
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer.
Admin State	The current administrative state.
Oper State	The current operational state.
Route Dist.	Displays the route distribution number.
AS Number	Displays the autonomous system number.
Router Id	Displays the router ID for this service.
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.

Label	Description
Description	Description of the split horizon group.
Last Changed	The date and time of the most recent management-initiated change to this split horizon group.
Service Destination Points (SDPs)	
SDP Id	The SDP identifier.
Type	Indicates whether this Service SDP binding is a spoke or a mesh.
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	Specifies the type of delivery used by the SDP (always 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.

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.
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.

Label	Description
Peer Pw Bits	<p>Indicates the bits set by the LDP peer when there is a fault on its side of the pseudowire. LAC failures occur on the SAP that has been configured on the pipe service, PSN bits are set by SDP-binding failures on the pipe service. The pwNotForwarding bit is set when none of the preceding failures apply, such as an MTU mismatch failure. This value is applicable only if the peer is using the pseudowire status signaling method to indicate faults.</p> <p>pwNotForwarding — Pseudowire not forwarding lacIngressFault Local — Attachment circuit RX fault lacEgressFault Local — Attachment circuit TX fault psnIngressFault Local — PSN-facing PW RX fault psnEgressFault Local — PSN-facing PW TX fault pwFwdingStandby — Pseudowire in standby mode</p>
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.

Label	Description
Egress Stats	The number of packets/octets forwarded out of this SAP.
Ingress Drop Stats	Number of packets/octets dropped by the system.
Extra-Tag Drop Stats	Number of packets received with the count of VLAN tags exceeding the count of VLAN tags implied by the SAP encapsulation.
Ingress Meter 1	The index of the ingress QoS meter of this SAP.
For. InProf	Number of in-profile packets/octets received on this SAP.
For. OutProf	Number of out-of-profile packets/octets received on this SAP.
If Name	IP interface name assigned by user.
Protocols	Protocols enabled for use on this interface.
Oper (v4/v6)	Operational status of this interface for IPv4 and IPv6.
IP Addr/mask	IPv4 address and mask assigned to this interface.
Address Type	Indicates 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.

Label	Description
Unreachables	Specifies the rate for ICMP unreachable messages.
TTL Expired	Specifies the rate for ICMP TTL messages.
MAC Address	Specifies the 48-bit IEEE 802.3 MAC address.

authentication

Syntax

authentication

Context

show>service>id

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

Commands in this context display subscriber authentication information.

statistics

Syntax

statistics [*policy name*] [*sap sap-id*]

Context

show>service>id>authentication

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays session authentication statistics for this service.

Parameters

policy name

Specifies the subscriber authentication policy statistics to display.

sap sap-id

Specifies the SAP ID statistics to display. See [Common CLI command descriptions](#) for command syntax.

Output

The following output is an example of session authentication statistics information.

Sample output

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

arp

Syntax

arp [*ip-address*] | [*mac ieee-address*] | [*sap sap-id*] | [*interface ip-int-name*] [*sdp sdp-id:vc-id*] [*summary*]

Context

show>service>id

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays the ARP table for the IES instance.

Parameters

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 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 service ARP information, and [Table 85: 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-100-190.11.1 1/1/11:0
=====
*A:ALA-12#
```

Table 85: 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays basic information about the service ID, including service type, description, SAPs, and SDPs.

Output

The following output is an example of basic service information, and [Table 86: Output fields: base](#) describes the output fields.

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 86: 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.

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 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.

statistics

Syntax

statistics [*sap sap-id*]

statistics [**sdp** *sdp-id:vc-id*]

statistics [**interface** *interface-name*]

Context

show>service>id>dhcp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 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 is an example of DHCP statistics information, and [Table 87: 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 87: Output fields: DHCP statistics

Label	Description
Received Packets	The number of packets received from the DHCP clients.

Label	Description
Transmitted Packets	The number of packets transmitted to the DHCP clients.
Received Malformed Packets	The number of corrupted or 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Specifies the IP address of the interface for which to display information.

Values 1.0.0.0 to 223.255.255.255

ip-int-name

Specifies 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 IP interface information, and [Table 88: Output fields: interface](#) describes the output fields.

Sample output

```
*A:ALA12# 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
SAP Id      : 71/1/1.2.2
TOS Marking : Untrusted
SNTP B.Cast : False
MAC Address  : Not configured.
IP MTU      : 1500
Arp Populate : Disabled
Cflowd      : None
Virt. If Index : 26
If Type      : IES
IES ID       : 88
Arp Timeout  : 14400
ICMP Mask Reply : True

Proxy ARP Details
Proxy ARP    : Enabled
Policies     : ProxyARP
Local Proxy ARP : Disabled

DHCP Details
```

```

Admin State : Up
Action      : Keep
ICMP Details
Redirects   : Number - 100
Unreachables : Number - 100
TTL Expired : Number - 100
Lease Populate : 0
Trusted      : Disabled
Time (seconds) - 10
Time (seconds) - 10
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
SAP Id      : 101/1/2:0
TOS Marking : Untrusted
SNTP B.Cast : False
MAC Address : Not configured.
Arp Populate : Disabled
Virt. If Index : 27
If Type     : IES
IES ID      : 88
Arp Timeout : 14400
Proxy ARP Details
Proxy ARP   : Disabled
Local Proxy ARP : Disabled
ICMP Details
Redirects   : Number - 100
Unreachables : Number - 100
TTL Expired : Number - 100
Time (seconds) - 10
Time (seconds) - 10
Time (seconds) - 10
-----
Interfaces : 2
=====
A:ALA-49#
    
```

Table 88: 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.

Label	Description
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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 SAP information, and [Table 89: Output fields: 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
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
-----
Sap Statistics
-----
                Packets      Octets
Forwarding Engine Stats
Dropped        : 0          0
Off. HiPrio    : 0          0
```

```

Off. LowPrio      : 0          0
Off. Uncolor     : 0          0

Queueing Stats(Ingress QoS Policy 100)
Dro. HiPrio      : 0          0
Dro. LowPrio     : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0

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

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

Ingress Queue 10 (Unicast) (Priority)
Off. HiPrio      : 0          0
Off. LoPrio      : 0          0
Dro. HiPrio      : 0          0
Dro. LoPrio      : 0          0
For. InProf      : 0          0
For. OutProf     : 0          0
'''
=====
*A:ALA-12#
    
```

Table 89: Output fields: 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, SapAdmin Down, InterfaceAdminDown, PortOperDown, PortMTUToo

Label	Description
	Small, 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.
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented.
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP.
Egress qos-policy	The egress QoS policy ID assigned to the SAP.
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP.
Egress Filter-Id	The egress filter policy ID assigned to the SAP.
Acct. Pol	The accounting policy ID assigned to the SAP.
Collect Stats	Specifies whether collect stats is enabled.
Dropped	The number of packets and octets dropped because of 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.

Label	Description
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.
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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 output is an example of SDP information, and [Table 90: Output fields: 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
Max Drop Count : 3
Hello Msg Len  : 0
Hold Down Time: 10

Statistics      :
I. Fwd. Pkts.  : 0
I. Fwd. Octs.  : 0
E. Fwd. Pkts.  : 0
MCAC Policy Name :
MCAC Max Unconst BW: no limit
MCAC In use Mand BW: 0
MCAC In use Opnl BW: 0
I. Dro. Pkts.  : 0
I. Dro. Octs.  : 0
E. Fwd. Octets : 0
MCAC Max Mand BW : no limit
MCAC Avail Mand BW: unlimited
MCAC Avail Opnl BW: unlimited

Associated LSP LIST :
Lsp Name       : A_B_1
Admin State    : Up
Time Since Last Tr*: 00h26m35s
Oper State     : Up

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                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 90: Output fields: 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.

Label	Description
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.
Last Changed	The date and time of the most recent change to the SDP.
Signaling	Specifies the signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The administrative state of the keepalive process.
Oper State	The operational state of the keepalive process.
Hello Time	Specifies how often the SDP echo request messages are transmitted on this SDP.
Max Drop Count	Specifies the maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault.
Hello Msg Len	Specifies the length of the SDP echo request messages transmitted on this SDP.
Hold Down Time	Specifies the amount of time to wait before the keepalive operating status is eligible to enter the alive state.
I. Fwd. Pkts.	Specifies the number of forwarded ingress packets.
I. Dro. Pkts.	Specifies the number of dropped ingress packets.
E. Fwd. Pkts.	Specifies the number of forwarded egress packets.
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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays aggregated routes.

Parameters

active

Keyword that filters out inactive aggregates.

Output

The following output is an example of aggregate route information, and [Table 91: 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 91: 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.

Label	Description
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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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.

mac *ieee-mac-addr*

Only displays ARP entries associated with the specified MAC address.

Output

The following output is an example of router ARP table information, and [Table 92: Output fields: router 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 92: Output fields: router 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays BGP routes that have been dampened because of 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 route or routes are listed.

When a **decayed**, **history**, or **suppressed** keyword is specified, only those types of dampened routes are listed.

Parameters

ip-prefix/mask

Displays damping information for the specified IP prefix and mask length.

ip-address

Displays damping entries 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 93: Output fields: BGP 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 : 32.32.27.203
Local Pref       : none
Age              : 00h22m09s          Last update    : 02d00h58m
    
```

```
FOM Present      : 738                FOM Last upd.   : 2039
Number of Flaps  : 2                  Flags           : ud*i
Path             : 60203 65001 19855 3356 1239 22406
Applied Policy   : default-damping-profile
-----
Network : 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 : 32.32.27.203
Local Pref     : none
Age            : 00h00m38s            Last update    : 02d01h20m
FOM Present    : 2011                FOM Last upd. : 2023
Number of Flaps : 2                  Flags         : ud*i
Path           : 60203 65001 19855 3356 3561 5551 1889
Applied Policy : default-damping-profile
-----
Network : 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 : 32.32.27.203
Local Pref     : none
Age            : 00h00m38s            Last update    : 02d01h20m
FOM Present    : 2011                FOM Last upd. : 2023
Number of Flaps : 2                  Flags         : ud*i
Path           : 60203 65001 19855 1299 702 1889
Applied Policy : default-damping-profile
-----
Network : 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 : 32.32.27.203
Local Pref     : none
Age            : 00h00m07s            Last update    : 02d01h20m
FOM Present    : 1018                FOM Last upd. : 1024
Number of Flaps : 1                  Flags         : ud*i
Path           : 60203 65001 19855 1299 702 1889
Applied Policy : default-damping-profile
-----
*A:ALA-12#

*A:ALA-12# show router 3 bgp damping 15.203.192.0/18 detail
=====
BGP Router ID : 10.0.0.14          AS : 65206    Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes 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 : 32.32.27.203
Local Pref     : none
Age            : 00h00m42s            Last update    : 02d01h20m
FOM Present    : 2003                FOM Last upd. : 2025
Number of Flaps : 2                  Flags         : ud*i
Path           : 60203 65001 19855 3356 702 1889
```

```
Applied Policy : default-damping-profile
-----
Paths : 1
=====
*A:ALA-12#
*A:ALA-12# show router 3 bgp damping suppressed detail
=====
BGP Router ID : 10.0.0.14      AS : 65206   Local AS : 65206
=====
Legend -
Status codes : u - used, s - suppressed, h - history, d - decayed, * - valid
Origin codes  : i - IGP, e - EGP, ? - incomplete, - best
=====
BGP Damped Routes (Suppressed)
=====
Network : 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 : 32.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update  : 02d01h20m
FOM Present  : 2936               FOM Last upd. : 3001
Number of Flaps : 3                Flags       : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 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 : 32.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update  : 02d01h20m
FOM Present  : 2936               FOM Last upd. : 3001
Number of Flaps : 3                Flags       : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 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 : 32.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update  : 02d01h20m
FOM Present  : 2936               FOM Last upd. : 3001
Number of Flaps : 3                Flags       : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
Network : 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 : 32.32.27.203
Local Pref   : none
Age          : 00h01m28s          Last update  : 02d01h20m
FOM Present  : 2936               FOM Last upd. : 3001
Number of Flaps : 3                Flags       : si
Path         : 60203 65001 19855 3356 702 1889
Applied Policy : default-damping-profile
-----
*A:ALA-12#
```

Table 93: Output fields: BGP damping

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured autonomous system number.
Local AS	The configured or inherited local AS for the specified peer group. If not configured, it is the same value as the AS.
Network	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 next hop for the route.
Peer AS	The AS 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 the following:

- **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 94: Output fields: BGP 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
Cluster Id       : None                Keep Alive       : 30
NLRI             : Unicast             Client Reflect   : Enabled
                                           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 94: Output fields: BGP 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 Disabled — The BGP peer group has been operationally disabled Down — The BGP peer group is operationally inactive Up — The BGP peer group is operationally active
Peer AS	The configured or inherited peer AS for the specified peer group
Local AS	The configured or inherited local AS for the specified peer group
Local Address	The configured or inherited local address for originating peering for the specified peer group
Loop Detect	The configured or inherited loop detect setting for the specified peer group
Connect Retry	The configured or inherited connect retry timer value
	Authentication
	None - No authentication is configured
	MD5 - MD5 authentication is configured
Local Pref	The configured or inherited local preference value
MED Out	The configured or inherited MED value assigned to advertised routes without a MED attribute

Label	Description
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 to 4294967285 — 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 next hop in route updates to neighbors in the peer group Enabled — BGP sends only its own IP address as the BGP next hop in route updates to neighbors in the specified peer group
Aggregator ID 0	Disabled — BGP is not configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group Enabled — BGP is configured to set the aggregator ID to 0.0.0.0 in all originated route aggregates sent to the neighbor in the peer group
Remove Private	Disabled — BGP does not remove all private ASNs from the AS path attribute in updates sent to the neighbor in the peer group Enabled — BGP removes all private ASNs 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

Label	Description
Keep Alive	The configured keepalive setting
Cluster Id	None — No cluster ID has been configured
Client Reflect	Disabled — The BGP route reflector will not reflect routes to this neighbor Enabled — The BGP route reflector is configured to reflect routes to this neighbor
NLRI	The type of NLRI information that the specified peer group can accept Unicast — IPv4 unicast routing information can be carried
Preference	The configured route preference value for the peer group
List of Peers	A list of BGP peers configured under the peer group
Total Peers	The total number of peers configured under the peer group
Established	The total number of peers that are in an established state

neighbor

Syntax

neighbor [*ip-address* [*detail*]]
neighbor [*as-address* [*detail*]]
neighbor *ip-address* [*family*] [*filter1*] [**brief**]
neighbor *ip-number* [**family**] *filter2*
neighbor *as-number* [**family**] *filter2*
neighbor *ip-address* **orf** [*filter3*]
neighbor *ip-address* **graceful-restart**

Context

show>router>bgp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays BGP neighbor information. It can be configured with or without any parameters. If it is configured without parameters, information about all BGP peers is displayed.

If it is issued with a specific IP address or ASN, information about only that specific peer or peers with the same AS is displayed.

If either **received-routes** or **advertised-routes** is specified, the routes received from or sent to the specified peer are displayed (see [Sample output — received routes](#)).



Note:

This information is not available for SNMP.

If either **history** or **suppressed** is specified, the routes learned from peers that either have a history or are suppressed (respectively) are displayed.

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 that the peer is operationally disabled and must be restarted by the operator.

Parameters

ip-address

Displays information for the specified IP address.

Values

ipv4-address: a.b.c.d (host bits must be 0)

ipv6-address: x:x:x:x:x:x:x[-*interface*]

x:x:x:x:x:d.d.d.d[-*interface*]

x: [0 to FFFF]H

d: [0 to 255]D

interface: 32 characters maximum,
mandatory for link local

addresses.

as-number

Displays information for the specified ASN.

Values 1 to 65535

family

Specifies the type of routing information to be distributed by this peer group.

Values **ipv4** — Displays only those BGP peers that have the IPv4 family enabled.

vpn-ipv4 — Displays only those BGP peers that have the VPN-IPv4 family enabled.

ipv6 — Displays only those BGP peers that have the IPv6 family enabled.

vpn-ipv6 — Displays only those BGP peers that have the VPN-IPv6 family enabled.

filter1

Displays information for the specified IP address

- Values**
- received-routes** — Displays the number of routes received from this peer.
 - advertised-routes** — Displays the number of routes advertised by this peer.
 - history** — Displays statistics for dampened routes.
 - suppressed** — Displays the number of paths from this peer that have been suppressed by damping.
 - detail** — Displays detailed information pertaining to *filter1*.

filter2

Displays information for the specified ASN.

- Values**
- history** — Display statistics for dampened routes.
 - suppressed** — Display the number of paths from this peer that have been suppressed by damping.
 - detail** — Displays detailed information pertaining to *filter2*.

brief

Displays information in a brief format. This parameter is only supported with received-routes and advertised-routes.

orf

Displays outbound route filtering for the BGP instance. ORF (Outbound Route Filtering) is used to inform a neighbor of targets (using target-list) that it is willing to receive. This mechanism helps lessen the update exchanges between neighbors and saves CPU cycles to process routes that could have been received from the neighbor only to be dropped/ ignored.

filter3

Displays path information for the specified IP address.

- Values**
- send** — Displays the number of paths sent to this peer.
 - receive** — Displays the number of paths received from this peer.

graceful-restart

Displays neighbors configured for graceful restart.

Output

The following outputs are examples of BGP neighbor information, and the associated tables describe the output fields.

- [Sample output, Table 95: Output fields: BGP neighbor](#)
- [Sample output — received routes, Table 96: Output fields: neighbor received routes](#)
- [Sample output — add-path, Table 97: 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            Keep Alive    : 30
Active Hold Time : 0        Active Keep Alive: 0
Cluster Id   : None
Preference   : 170          Num of Flaps  : 0
Recd. Prefixes : 0         Active Prefixes : 0
Recd. Paths   : 0         Suppressed Paths : 0
Input Queue   : 0          Output Queue  : 0
i/p Messages  : 0          o/p Messages  : 0
i/p Octets    : 0          o/p Octets    : 0
i/p Updates   : 0          o/p Updates   : 0
Export Policy : direct2bgp
=====
*A:ALA-12#

*A:ALA-12# show router 3 bgp neighbor detail
=====
BGP Neighbor (detail)
=====
-----
Peer : 10.0.0.15          Group : To_AS_40000
-----
Peer AS      : 65205
Peer Address : 10.0.0.15      Peer Port    : 0
Local AS     : 65206
Local Address : 10.0.0.16     Local Port    : 0
Peer Type    : External
State        : Active        Last State    : Connect
Last Event   : openFail
Last Error   : Hold Timer Expire
Connect Retry : 20           Local Pref.   : 100
Min Route Advt. : 30        Min AS Orig.  : 15
Multipath : 1            Multihop      : 5
Damping       : Disabled    Loop Detect    : Ignore
MED Out       : No MED Out  Authentication : None
Next Hop Self : Disabled    AggregatorID Zero: Disabled
Remove Private : Disabled   Passive       : Disabled
Prefix Limit  : No Limit
Hold Time     : 90            Keep Alive    : 30
Active Hold Time : 0        Active Keep Alive: 0
Cluster Id    : None        Client Reflect : Enabled
Preference    : 170          Num of Flaps  : 0
Recd. Prefixes : 0         Active Prefixes : 0
Recd. Paths   : 0         Suppressed Paths : 0
Input Queue   : 0          Output Queue  : 0
i/p Messages  : 0          o/p Messages  : 0
i/p Octets    : 0          o/p Octets    : 0
```

```
i/p Updates      : 0                o/p Updates      : 0
Export Policy    : direct2bgp
```

```
=====
*A:ALA-12#
```

Table 95: Output fields: BGP neighbor

Label	Description
Peer	The IP address of the configured BGP peer
Group	The BGP peer group to which this peer is assigned
Peer AS	The configured or inherited peer AS for the peer group
Peer Address	The configured address for the BGP peer
Peer Port	The TCP port number used on the far-end system
Local AS	The configured or inherited local AS for the peer group
Local Address	The configured or inherited local address for originating peering for the peer group
Local Port	The TCP port number used on the local system
Peer Type	External — Peer type configured as external BGP peers Internal — Peer type configured as internal BGP peers
State	Idle — The BGP peer is not accepting connections Active — BGP is listening for and accepting TCP connections from this peer Connect — BGP is attempting to establish a TCP connection from this peer Open Sent — BGP has sent an OPEN message to the peer and is waiting for an OPEN message from the peer Open Confirm — BGP has received a valid OPEN message from the peer and is awaiting a KEEPALIVE or NOTIFICATION Established — BGP has successfully established a peering and is exchanging routing information
Last State	Idle — The BGP peer is not accepting connections Active — BGP is listening for and accepting TCP connections from this peer Connect — BGP is attempting to establish a TCP connection 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 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 connectRetry — Connect retry timer expired holdTime — Hold time timer expired keepAlive — Keepalive timer expired rcvOpen — Receive an OPEN message revKeepalive — Receive an KEEPALIVE message rcvUpdate — Receive an UPDATE message rcvNotify — 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

Label	Description
Loop Detect	Ignore — The BGP neighbor is configured to ignore routes with an AS loop Drop — The BGP neighbor is configured to drop the BGP peering if an AS loop is detected Off — AS loop detection is disabled for the neighbor
MED Out	The configured or inherited MED value 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 ASNs from the AS path attribute, in updates sent to the specified neighbor Enabled — BGP removes all private ASNs 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
Prefix Limit	No Limit — No route limit assigned to the BGP peer group 1 to 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

Label	Description
	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

Sample output — 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 96: Output fields: neighbor received routes

Label	Description
BGP Router ID	Displays the local BGP router ID
AS	The configured AS 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 ? — incomplete > — best
Network	Route IP prefix and mask length for the route
Nexthop	BGP next hop 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 97: 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

Label	Description
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 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

Label	Description
	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
Hold Time	The configured hold-time setting
Keep Alive	The configured keepalive setting
Min Hold Time	The configured minimum hold-time setting
Active Hold Time	The negotiated hold time, if the BGP neighbor is in an established state
Active Keep Alive	The negotiated keepalive time, if the BGP neighbor is in an established state
Cluster Id	The configured route reflector cluster ID None — no cluster ID has been configured
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

Label	Description
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
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

Label	Description
	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
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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays a summary of BGP path attributes.

Output

The following output is an example of BGP path information, and [Table 98: Output fields: BGP 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 98: Output fields: BGP 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 next hop.
MED	The Multi-Exit Discriminator value.

Label	Description
Local Preference	The local preference value.
Refs	The number of routes using a specified set of path attributes.
ASes	The number of AS numbers in the AS path attribute.
Segments	The number of segments in the AS path attribute.
Flags	eBGP-learned — Path attributes learned by an eBGP peering.
	iBGP-Learned — Path attributes learned by an iBGP peering.
Aggregator	The route aggregator ID.

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays BGP route information.

When this command is issued without any parameters, the entire BGP routing table displays.

When this command is issued with an IP prefix/mask or IP address, the best match for the parameter displays.

Parameters

family *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 enabled 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	{ <i>ip-address:number1</i> <i>as-number1:number2</i> <i>as-number2:number3</i> }
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:{ <i>ip-address:comm-val1 as-number1:comm-val2 as-number2:comm-val1</i> }	
as-number1	0 to 65535
comm-val1	0 to 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 99: Output fields: BGP 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 99: Output fields: BGP 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.
next-hop	The BGP next hop.
From	The advertising IP address BGP neighbor.
Res. next-hop	The resolved next hop.
Local Pref.	The local preference value.
Flag	u — used
	s — suppressed
	h — history
	d — decayed
	* — valid
	i — igp
	e — egp
	? — incomplete
> — 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.

Label	Description
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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays a summary of BGP neighbor information.

If confederations are not configured, that portion of the output will not display.

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 the following:

- **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 neighbor information, and [Table 100: 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 Supressed Routes : 0          Total History Routes : 0
Total Decayed Routes   : 0
=====
BGP Summary
=====
Neighbor      AS PktRcvd PktSent InQ OutQ  Up/Down State|Recv/Actv/Sent
-----
10.0.0.1      65206      5  21849  0   0  00h01m29s 32/0/86683
10.0.0.12     65206      0    0    0   0  00h01m29s Active
10.0.0.13     65206      5  10545  0   50 00h01m29s 6/0/86683
10.0.0.15     65205      0    0    0   0  00h01m29s Active
10.0.0.16     65206      5   9636  0   50 00h01m29s 6/0/86683
10.0.27.1     2          0    0    0   0  00h01m29s Active
10.0.28.1     60203    22512  15    0   0  00h01m29s 116955/86689/9
=====
*A:ALA-12#
```

Table 100: Output fields: BGP summary

Label	Description
BGP Router ID	The local BGP router ID.
AS	The configured AS 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.

Label	Description
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 because of route damping.
Total History Routes	Total number of routes with history because of route damping.
Total Decayed Routes	Total number of decayed routes because of route damping.
Neighbor	BGP neighbor address.
AS (Neighbor)	BGP neighbor AS 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 current state of the BGP neighbor (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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays the router IP interface table, sorted by interface index.

Parameters

ip-address

Only displays the interface information associated with the specified IP address.

ip-int-name

Only displays the interface information associated with the specified IP interface name.

detail

Displays detailed IP interface information.

summary

Displays summary IP interface information for the router.

Output

The following outputs are example of interface information, and the associated tables describe the output fields.

- [Sample output — detailed, Table 101: Output fields: detailed IP interface](#)
- [Sample output — summary, Table 102: Output fields: interface summary](#)

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
              Time (seconds) - 10
TTL Expired  : Number - 100
              Time (seconds) - 10
=====
*A:ALA-12#
    
```

Table 101: Output fields: detailed IP interface

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 — 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.

Label	Description
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.
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 will issue 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 will issue in a specific period of time.
	Disabled — Indicates the IP interface does not generate ICMP destination unreachable messages.
TTL Expired	The maximum number (Number) of ICMP TTL expired messages the IP interface will issue in a specific period of time (Time (seconds)).
	Disabled — Indicates 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 102: Output fields: interface summary

Label	Description
Instance	The router instance number.
Router Name	The name of the router instance.
Interfaces	The number of IP interfaces in the router instance.

route-table

Syntax

route-table [*ip-prefix* [*/mask*] [**longer**] | [**protocol** *protocol*] | [**summary**]]

Context

show>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 route table summary information.

Output

The following output is an example of routing table information, and [Table 103: Output fields: route 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 103: Output fields: route table

Label	Description
Dest Address	The route destination address and mask.
next-hop	The next-hop IP address for the route destination.
Type	Local — The route is a local route. Remote — The route is a remote route.
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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Only displays static ARP entries associated with the specified IP address.

ip-int-name

Only displays static ARP entries associated with the specified IP interface name.

mac ieee-mac-addr

Only displays static ARP entries associated with the specified MAC address.

Output

The following output is an example of static ARP table information, and [Table 104: 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 104: 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.
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*] [**detail**]

Context

show>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays the static entries in the routing table.

If no options are present, all static routes are displayed sorted by prefix.

Parameters

ip-prefix /mask

Displays static routes only matching the specified *ip-prefix* and *mask*.

preference preference

Only displays static routes with the specified route preference.

Values 0 to 65535

next-hop ip-addr

Only displays static routes with the specified next hop IP address.

detail

Displays detailed information about the static route.

Output

The following output is an example of static route information, and [Table 105: 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
Next hop Type    : IP
Active           : Y
Preference       : 5
Tag              : 0
State            : n/a
Drop Count       : 3
CPE Echo Reply Rx : 3
CPE Down Trans  : 0
-----
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
Next hop Type    : IP
Active           : Y
Preference       : 5
Tag              : 0
-----
No. of Static Routes: 2
    
```

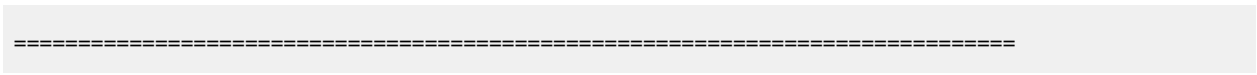


Table 105: Output fields: static route

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 next-hop for this type of route is black hole.
	ID — The static route is an indirect route, where the next-hop for this type of route is the non-directly connected next-hop.
	NH — The route is a static route with a directly connected next-hop. The next-hop 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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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 when 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 106: 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 106: Output fields: tunnel table

Label	Description
Active	Displays the number of active LDPs and SDPs.
Available	Displays the number of available LDPs and SDPs.

8.4.2.6 VPRN clear commands

interface

Syntax

interface [*ip-int-name* | *ip-addr*] [**icmp**]

Context

clear>router

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command resets the specified BGP peer or peers. This can cause existing BGP connections to be shut down 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

Specifies that BGP neighbors reevaluate all routes in the local-RIB against the configured export policies.

soft-inbound

Specifies that BGP neighbors reevaluate all routes in the RIB-In against the configured import policies.

statistics

Specifies the BGP neighbor statistics.

protocol

Syntax

protocol

Context

clear>router>bgp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command resets the entire BGP protocol. If the ASN was previously changed, the BGP ASN does not inherit the new value.

id

Syntax

`id service-id`

Context

clear>service

clear>service>statistics

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command clears commands for a specific service.

Parameters

service-id

Specifies the ID that uniquely identifies a service.

Values 1 to 2147483648

sap

Syntax

`sap sap-id {all | counters | stp}`

Context

clear>service>statistics

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command clears all spanning tree statistics for the service ID.

8.4.2.7 VPRN debug commands

id

Syntax

[no] id *service-id*

Context

debug>service

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command displays Subscriber Host Connectivity Verification (SHCV) events for a particular SAP.

Parameters

sap-id

Specifies the physical port identifier portion of the SAP definition. See [Common CLI command descriptions](#) for command syntax.

sap

Syntax

[no] **sap** *sap-id*

Context

debug>service>id

debug>service>stp

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables debugging for a particular event type.

The **no** form of this command disables debugging.

event-type

Syntax

[no] event-type {config-change | oper-status-change}

Context

debug>service>id>sap

Platforms

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

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

7210 SAS-K 2F6C4T and 7210 SAS-K 3SFP+ 8C

Description

This command enables STP debugging for port states.

The **no** form of this command disables debugging.

9 Common CLI command descriptions

This section 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

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 107: SAP-ID formats

Type	Syntax	Example
port-id	<i>slot/mda/port[.channel]</i>	1/1/5
null	[<i>port-id</i> <i>lag-id</i>]	<i>port-id</i> : 1/1/3 <i>lag-id</i> : lag-3
dot1q	[<i>port-id</i> <i>lag-id</i>]:qtag1	<i>port-id</i> :qtag1: 1/1/3:100 <i>lag-id</i> :qtag1:lag-3:102 <i>cp.conn-prof-id</i> : 1/2/1:cp.2
qinq	[<i>port-id</i> <i>lag-id</i>]:qtag1.qtag2	<i>port-id</i> :qtag1.qtag2: 1/1/3:100.10 <i>lag-id</i> :qtag1.qtag2: lag-10:

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 108: 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.
Ethernet	QinQ	qtag1: 0 to 4094 qtag2: 0 to 4094	The SAP is identified by two 802.1Q tags on the port.

Notes:

1. A 0 qtag1 value also accepts untagged packets on the dot1q port.
2. A 0 qtag1 value is allowed with some 7210 SAS platforms. See [Configuration guidelines for 7210 SAS-D and 7210 SAS-Dxp](#) and [Configuration guidelines for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C](#) for information about platforms and frame processing.
3. A 0 qtag2 value is allowed only on 7210 SAS-K platforms as described in [Configuration guidelines for 7210 SAS-K 2F1C2T, 7210 SAS-K 2F6C4T, and 7210 SAS-K 3SFP+ 8C](#).

10 Appendix: port-based split horizon

This section provides Port-Based Split Horizon configuration information.

10.1 Overview



Note:

This feature is only supported on the 7210 SAS-D and 7210 SAS-Dxp.

The port-based split horizon feature can be used to disable local switching on the 7210 SAS. A loop-free topology can be achieved using split horizon on 7210 SAS switches.

Traffic arriving on an access or an access-uplink port within a split horizon group will not be copied to other access and an access-uplink ports in the same split horizon group, but will be copied to an access-uplink ports in other split horizon groups.


Because 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

Example

The following figure shows an example of split horizon groups used to prevent communication between two access SAPs and between two access-uplink SAPs.

Figure 89: 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 "access uplink" is created to prevent any communication between SAP part of port 1/1/1 and port 1/1/3 (configured as an access-uplink) within the same VPLS.
3. VPLS 100 is created on 7210 SAS-1 with SAPs 1/1/1, 1/1/3, and SAPs on 1/1/4 and 1/1/5 as part of this VPLS. CE1, CE2 and CE3 are the customer sites.

4. With this configuration, any communication between ports 1/1/4 and 1/1/5 gets blocked, similarly communication between ports 1/1/1 and 1/1/3 gets blocked but any traffic received on ports (SAPs) part of split horizon group "access uplink" will be switched to ports (SAPs) part of split horizon group "access" and 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:

1. Create a split horizon group in the config prompt. The group name must be unique across the system.

```
7210-SAS1>config#info
#-----
echo "Split-horizon-group Configuration"
#-----
    split-horizon-group access create
        description "Block access between access Ports"
    split-horizon-group access-uplink create
        description "Block access between access-uplink Ports"
    exit
#-----
7210-SAS1>config#
```

2. Associate ports 1/1/4 and 1/1/5 with split horizon group "access" By default all ports are access ports. The default Ethernet encapsulation for access port is "null".

```
7210-SAS1>config#info
#-----
echo "Port Configuration"
#-----
    port 1/1/4
        split-horizon-group access
        ethernet
        exit
        no shutdown
    exit
    port 1/1/5
        split-horizon-group access
        ethernet
        exit
        no shutdown
    exit
#-----

7210-SAS1>config#
```

3. Configure ports 1/1/1 and 1/1/3 as access uplink and associate these ports with split horizon group "access-uplink" default Ethernet encapsulation for access uplink port is "qinq".

```
7210-SAS1>config# info
#-----
echo "Port Configuration"
#-----
    port 1/1/1
        split-horizon-group access-uplink
        ethernet
        mode access uplink
```

```

        exit
        no shutdown
    exit
    port 1/1/3
        split-horizon-group access-uplink
        ethernet
            mode access uplink
        exit
        no shutdown
    exit
#-----

7210-SAS1>config#
    
```

4. Create a VPLS instance 100.

```

#-----
echo "Service Configuration"
#-----
    service
        customer 2 create
        exit
        vpls 100 customer 2 create
        stp
            shutdown
        exit
        sap 1/1/1:100.* create
        exit
        sap 1/1/3:100.* create
        exit
        sap 1/1/4 create
        exit
        sap 1/1/5 create
        exit
        no shutdown
    exit
    ...
#-----
    
```



Note:

A split horizon on a port must be configured before creating any SAPs associated with that port.

10.2.1 Verification

The following shows the command usage to verify the split horizon configuration on a 7210 SAS.

```

7210-SAS1# show split-horizon-group
=====
Port: Split Horizon Group
=====
Name                Description
-----
access              Block access between access Ports
access-uplinkBlock access between access-uplink 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 access-uplink
```

```
=====
Port: Split Horizon Group
=====
```

Name	Description
Access-uplink uplink Ports	Block access between access-

```
-----
Associations
-----
```

Port1/1/1	10/100/Gig Ethernet SFP
Port1/1/3	10/100/Gig Ethernet SFP

```
Ports Associated : 2
=====
```

```
7210-SAS1#
```


11 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.

11.1 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 4724, Graceful Restart Mechanism for BGP (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 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

11.2 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, and Sx/S-1/10GE



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, and Sx/S-1/10GE



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

11.3 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

11.4 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

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

RFC 7490, Remote Loop-Free Alternate (LFA) Fast Reroute (FRR) is supported on K12, K30, Mxp, Sx/S-1/10GE, R6, and R12



Note:
With Segment Routing.

11.5 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-rrp-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 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 supports 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 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

11.6 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.

11.7 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.

11.8 IP — Version 6

RFC 2464, Transmission of IPv6 Packets over Ethernet Networks 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.

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, 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.

RFC 4007, IPv6 Scoped Address Architecture 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.

RFC 4193, Unique Local IPv6 Unicast Addresses 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.

RFC 4291, Internet Protocol Version 6 (IPv6) Addressing Architecture 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.

RFC 4443, Internet Control Message Protocol (ICMPv6) for the 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.

RFC 4861, Neighbor Discovery for IP version 6 (IPv6) 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.

RFC 4862, IPv6 Stateless Address Autoconfiguration (Router Only) is supported on K12, K30, T(N), Mxp, Sx/S-1/10GE, Sx-10/100GE, Sx/S-1/10GE-VC, R6, and R12

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, 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.

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

11.9 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.

11.10 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, 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, R6, and R12

11.11 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

11.12 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

11.13 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

11.14 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.

11.15 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

11.16 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

11.17 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

11.18 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, 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, R6, and R12

11.19 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

11.20 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

11.21 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.

11.22 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:
Only on 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:
Only on 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:
Only on Dxp-12p ETR. Dxp-ETR and 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:
Only on 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:

Only on 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:

Only on 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:

Only on 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:

Only on 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:

Only on 7210 SAS-Sx 10/100GE QSFP28 variant and Dxp-12p ETR.

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

11.23 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](#)